Oracle® Applications InterConnect

User's Guide

Release 3.1.3

August 3, 2000

Part No. A86039-01

ORACLE®

Oracle Applications InterConnect User's Guide, Release 3.1.3

Part No. A86039-01

Contents

4   Runtime Management Console

Glossary

Index

Send Us Your Comments

Oracle Applications InterConnect User's Guide , Release 3.1.3

Part No. A75325-04

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, please indicate the chapter, section, and page number (if available). You can send comments to us in the following ways:

> Oracle Corporation
> Oracle Server Documentation
> Attention: Oracle Applications InterConnect
> 500 Oracle Parkway
> Redwood Shores, CA 94065
> USA

If you would like a reply, please give your name, address, and telephone number below.

If you have problems with the software, please contact your local Oracle Support Services.

Preface

Welcome to the *Oracle Applications InterConnect User's Guide,* **Release 3.1.3**.

This user's guide provides information to work effectively with Oracle Applications InterConnect, including instructions on how to operate and use the product.

This preface explains how this user's guide is organized and introduces other sources of information that can help you.

Intended Audience

The target audience for Applications InterConnect is any organization that needs to integrate an Oracle product with another Oracle product or a legacy system.

Structure

This manual contains four chapters:

Chapter 1	Introduces Oracle Applications InterConnect and presents an overview of the product and the tools.
Chapter 2	Describes the design-time components and concepts of Oracle Applications Interconnect. It also explains iStudio, the tool provided to organize and improve integration tasks.
Chapter 3	Describes the runtime components and concepts of Oracle Applications InterConnect.
Chapter 4	Introduces Runtime Management Console and describes how you use it to troubleshoot errors during the execution of interconnected applications and systems.

Audience

This guide is targeted at the following types of users:

- Integration engineers, for iStudio.

- System Administrators, for the runtime component.

- SAP R/3 or other ERP vendor consultants, for the integration pack.

The audience ideally should have the following pre-requisites, which are discussed but not explained:

1. Domain knowledge of the applications that you are integrating.

2. Database concepts and SQL, PL/SQL, or SQL* Plus programming background.

Related Materials

Additional Documentation

Oracle Applications InterConnect shares business and setup information with other Oracle Applications products. Therefore, you may want to refer to other user's guides when you set up and use Oracle Applications InterConnect.

- *Oracle8i PL/SQL Programming Guide*

- *Oracle8i - The Complete Reference*

- *Oracle Message Broker (OMB) User's Guide*

Conventions

The following conventions are also used in this manual:

Convention	Meaning
. . .	Vertical ellipsis points in an example mean that information not directly related to the example has been omitted.
. . .	Horizontal ellipsis points in statements or commands mean that parts of the statement or command not directly related to the example have been omitted
boldface text	Boldface type in text indicates a term defined in the text, the glossary, or in both locations.

Convention	Meaning
< >	Angle brackets enclose user-supplied names.
[ ]	Brackets enclose optional clauses from which you can choose one or none.

Support

From on-site support to central support, our team of experienced professionals provides the help and information you need to keep Oracle Applications InterConnect working for you. This team includes your Technical Representative, Account Manager, and Oracle's large staff of consultants and support specialists with expertise in your business area, managing an Oracle8 server, and your hardware and software environment.

x

1

Introduction to Oracle Applications InterConnect

This chapter provides you an overview of Oracle Applications InterConnect, its features and components.

This chapter contains the following sections:

What is Oracle Applications InterConnect?

Oracle Applications InterConnect is a comprehensive and flexible application integration platform that enables seamless integration of enterprise software. Applications InterConnect is specifically designed to integrate Oracle products with other Oracle products or third party legacy systems.

Applications InterConnect allows de-coupled integration of applications. It eliminating the complexities of point to point solutions. In addition, OAI provides a tool (iStudio) for modeling the data in the integration scenario. iStudio eliminates the need for "hardwired" or "hard-coded" integrations. Users define their integration using iStudio which minimizes (at best eliminates) the need to write any code for the integration. The integration information is captured as metadata in a repository.

Using Applications InterConnect then becomes a simple three step process:

1. Install the InterConnect components i.e. the generic plumbing. No coding!

2. Install pre-configured, pre-tested metadata for that specific integration to make the plumbing "smart." This metadata is created using iStudio and is captured in a repository. NOTE: Oracle provides certain productized pre-created metadata as Integration Packs (e.g. CRM 3i to SAP, iProcurement to SAP, CRM 3i to Oracle ERP 10.7). If the customer needs are outside these productized packs, a new pack must be created for them by Oracle Consulting. No coding!

3. Modify the metadata through iStudio to cater to specific application customizations that the customer might have done. No coding!

Platform Features and Core Components

Oracle Applications InterConnect's key features are:

- Out-of-the-box integration

- Distributed deployment

- Loose coupling of applications

- Tools for easy customization

- Single point of contact for customer support

- Based on proven Oracle Server Technologies (Java, CORBA, Oracle 8i)

- Event-based distributed messaging system

- Guaranteed, exactly once, in-order message delivery

- Content-based routing support

- Supports each of the major messaging paradigms—Publish/Subscribe, Request/Reply, and Point-To-Point.

Out-of-the-box Integration

The integration packs described above provide the pre-configurations required to enable integration between supported applications with little or no need for customization. This acts as a plug-and-play feature for most standard integration.

Distributed Deployment

Applications InterConnect's infrastructure supports reliable, high-performance integration between applications installed either locally or distributed over a WAN. This gives you the flexibility of deploying geographically distributed applications.

Loose Coupling of Applications

Oracle Applications InterConnect is based on an asynchronous messaging architecture. This allows the applications being integrated to be loosely coupled with each other. When loosely coupled, applications can continue to function normally even when participating applications become unavailable. Further, when an application is upgraded or modified, loose coupling minimizes the impact on the other applications.

Synchronous messages is also supported using the request/reply paradigm.

Tools for Easy customization

Oracle Applications InterConnect includes iStudio, a GUI-based integration specification tool that allows the product developer, consultant or IT professional to visually review and modify the pre-configured integration specifications. iStudio thus supports the easy definition of integration customization, providing the ability to extend existing integration to support additional applications. The configuration information you create in iStudio is stored in the InterConnect Repository for use by

Guaranteed, Exactly Once, In-order Message Delivery

Oracle Applications Interconnect runtime synchronizes all interconnected applications, and coordinates and verifies the receipt, transformation and delivery of each message in the system to insure inter-system integrity.

Content-based Routing Support

Messages can be routed to a specific application based on specific content values contained in the message. For example, an electronic funds transaction settlement application is designed to transmit bank transactions with a specific bank code to identify the destination bank system. When the EFT application publishes each message at runtime, the Oracle Application InterConnect runtime component determines the BankCode value based on objects stored in the repository, and routes the message to the appropriate recipient system.

Supports Major Messaging Paradigms

Applications InterConnect supports the three major messaging paradigms. These paradigms are defined by IT professionals using iStudio at design time. The definitions are used at runtime to route each message appropriately, to the widest possible pool of interconnected applications.

Connectivity with MQ-series Applications.

ednote: need content

Applications InterConnect's Value Add to your Integration

Applications InterConnect adds value at each stage of integration:

- building the integration through the product development channel
- customizing the integration through the consulting channel
- evolving the integration with the end user

In addition, it provides:

- auditing/persistence support for messages
- disconnected/asynchronous processing

Applications InterConnect Components

The four main components of Applications InterConnect are:

- iStudio, the GUI-based tool
- Adapters
- InterConnect Repository
- Oracle Message Broker (OMB)

iStudio

iStudio is designed to minimize (and ideally, eliminate) the need for source code development to enable application integration. iStudio is an easy to use, wizard-based tool used to specify and configure the seamless integration of applications using Applications InterConnect.

You use iStudio to model hierarchical data that represents the business objects you are integrating into multiple applications. You also use iStudio to specify data transformations, and implement Publish/Subscribe, Request/Reply, or Point-to-Point messaging paradigms.

Runtime Components

At application runtime, Oracle Applications InterConnect provides application level integration by using Adapters to communicate with the enterprise applications. An InterConnect Repository is used to store and deliver the integration specifications developed using iStudio.

Adapters

An adapter is the Applications InterConnect component that sits at the spoke with the application to make it InterConnect enabled. Internally, the adapter is written as two components for improved reuse of existing interfaces. These components are:

Bridge. This is the application specific piece of the adapter. The bridge communicates with the specific application interface to transfer data between the application and Applications InterConnect. For messages outbound from an application, the bridge is responsible for converting the data from the application's native system format to the agent's internal format (and conforming to the application view of data defined in iStudio). It then passing it on to the agent (described below) for further processing. For inbound messages, the bridge receives the message from the agent in the agent's internal format (and conforming to the application view of data defined in iStudio). It then converts the message back to the application's native format and pushes the data contained therein into the application. Each communication protocol requires a unique bridge.

Two products using the same protocol may use the same bridge code, though at runtime two separate processes are created. The bridge is also called the technology/protocol adapter.

2. Agent. The agent is a generic engine that carries out instructions for transformations and routing captured in repository metadata (populated by iStudio). The agent does not know how to talk to a particular application. For messages outbound from the application, the agent receives the message in it's internal format from the bridge. This internal format conforms to the application view of data (see iStudio description). The agent then queries the repository for metadata to transform this message to the common view and pushes the message to OMB.

For inbound messages, the agent receives a message from OMB that conforms to the common view defined in iStudio. The agent queries the repository for metadata and transforms the message from the common view to the application view. The message is then pushed to the bridge. The agent is also know as the integration adapter.

Interconnect Repository

The Interconnect Repository stores iStudio-generated metadata, encapsulating all the integration information.

Oracle Message Broker (OMB)

Oracle Message Broker is the hub component that controls all inter-application messaging at runtime. It uses a store-and-forward methodology based on Oracle Messaging Stack using Advanced Queuing (AQ).
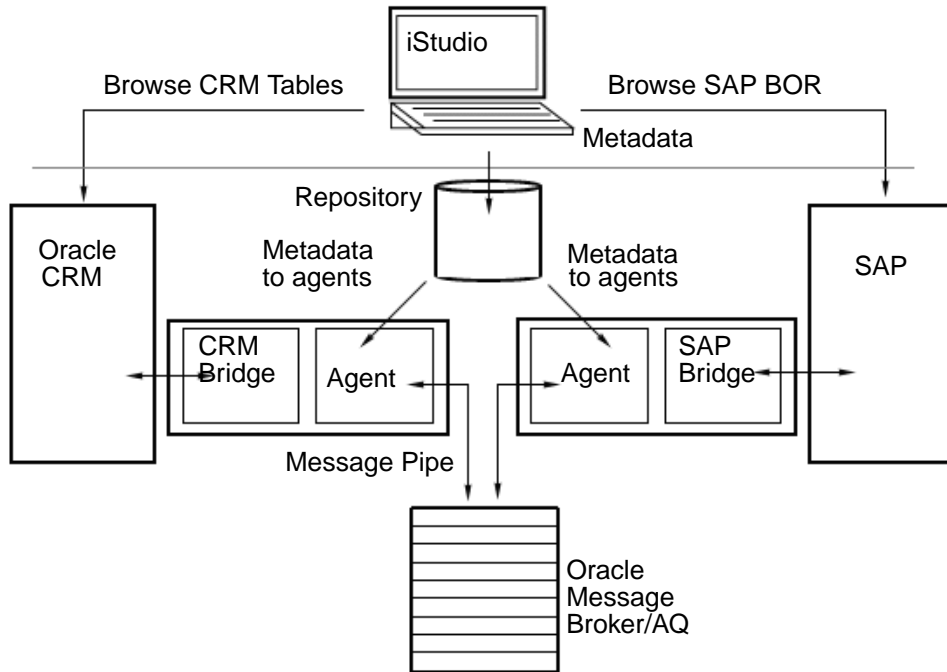
Applications InterConnect Architecture

Figure 1–1   An example of Applications InterConnect architecture—Oracle CRM 3i application integrated with SAP R/3 backend.

2

Design Time Concepts and iStudio

This chapter describes Applications InterConnect's design time concepts and iStudio, the GUI-based application for creating metadata that describes events, objects and other types of messages that interconnected applications use during runtime. It also explains how to perform important tasks including publishing and subscribing to an event using iStudio.

This chapter contains the following sections:

Modeling Paradigm

Applications InterConnect modeling is primarily based on the definition of a common format for data to be exchanged between interconnected applications. The model supports a hub-and-spoke architecture at the design level.

Each application possesses a specific format for each data structure used in the application. In iStudio, this "application view" of the data is mapped to a common view and stored as metadata in the Applications InterConnect Repository. At runtime, the metadata is retrieved from the repository by the adapter and used to map the application view of a specific data object to the common view and sent out as a message. Incoming messages likewise are converted from the common view to the application view and sent in to the application for processing.

This loose coupled paradigm ensures minimal modifications to integration when an application is modified or a new application is added.

## Hub and Spoke Methodology

Figure 2–1   Hub and spoke integration methodology used in Applications InterConnect

Applications InterConnect supports three messaging paradigms. These paradigms are defined in iStudio at design time. The definitions are used at runtime to route the messages appropriately.

The messaging paradigms are:

- Publish/Subscribe
- Request/Reply
- Point-to-Point

Publish/Subscribe

Publish/Subscribe messaging is used when an application sends—or publishes—a message to each subscribing application, and does not wait for a reply. This is often referred to as a non-blocking call. Using iStudio, you define which application is the publisher and the subscriber for a particular message.

Request/Reply

Request/reply messaging is used when an application sends a message and waits for a reply. In this paradigm, the application that sends the request blocks—or waits—for a reply. Even though many applications might be listening for this message, only one of them receives the message and sends a reply. The reply is sent using the point-to-point paradigm described below.

Point-to-Point

Point-to-Point messaging is used when the sending application specifies a destination and the message is delivered only to that destination.

Note: If this paradigm is used, a message can be delivered only to one destination.

Integration Process Overview

Application integration using Oracle Applications InterConnect involves two phases. They are:

- Design Time
- Runtime

Design Time

During the design phase, a developer uses iStudio to define the integration objects, applications which participate in the integration, and the specifications of the data exchanged between applications. All the specifications are stored as metadata in the Applications InterConnect Repository.

Runtime

For each application participating in a specific integration, Applications InterConnect attaches an adapter to it. At runtime, the adapter retrieves the metadata from the Repository to determine the format of messages, and perform

transformations between the various data formats, and route the messages to th3e appropriate queues under OMB.
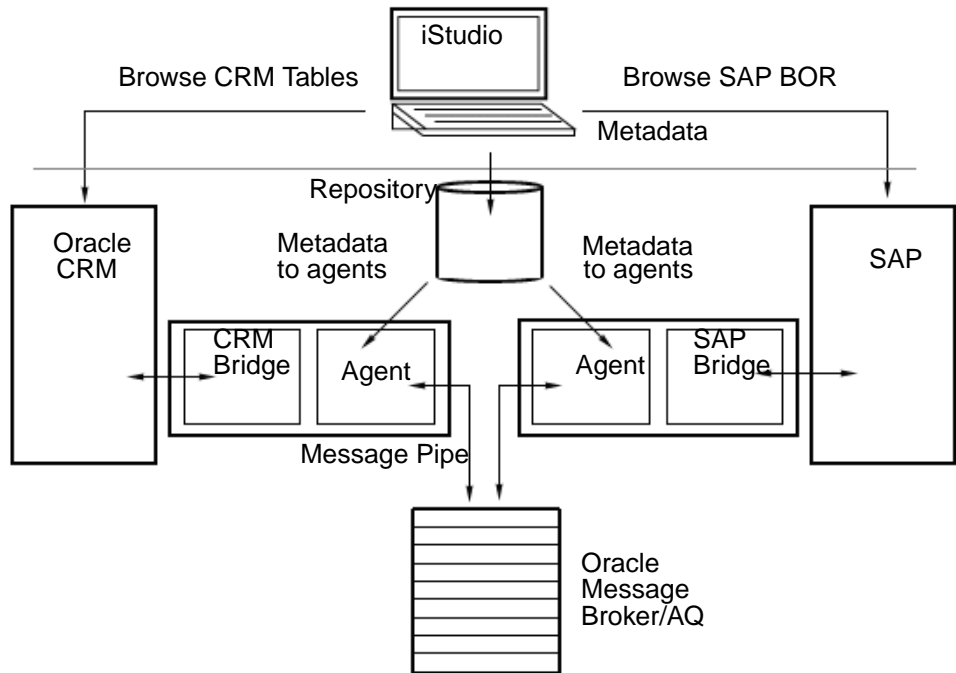
Figure 2–2   A graphical overview of design time and runtime phases in integration

iStudio Concepts

Applications InterConnect modeling is accomplished using iStudio, an easy-to-use graphical tool. These sections describe the concepts implemented by iStudio.

Projects

In iStudio, a project defines a Repository connection. To create a project, you have to specify the Repository connection information including the repository name and host. Several people may work on the same project.

Workspaces

Workspaces store user settings and preferences (for example, SAP and DB login credentials) in iStudio. In contrast, each iStudio project stores information about the current repository. When a project is opened in a workspace, it is automatically added to the list of recently opened projects. These are accessed by clicking **File** - **Reload**. To reload the repository, click on a project in the **reload** menu.

Applications

Each component that is integrated with Applications InterConnect is referred to as an application. Each application can specify the events and procedures it is interested in. The user can define the application views and transformations between these common objects and their application views.

Business Objects

For convenience and grouping, procedures and events are organized into business objects. A business object roughly corresponds to an object-oriented class. Procedures in the business object resemble methods in the class. For example:

Customer with a `getAddress()` procedure and a `newCustomer` event.

Events are used to model the publish/subscribe paradigm. Procedures are used to model the request/reply paradigm.

Common Objects

Applications InterConnect de-couples applications by introducing the concept of common objects. Procedures, events, and common data types (described below) are examples of common objects. The developer only needs to define interactions of an application with the common objects, and not with other applications directly. In other words, a point-to-point architecture is replaced by a hub and spoke architecture, with common objects serving as the representation of data in the hub.

Procedures

Procedures are common objects that encapsulate functionality. Each procedure has a name and IN/OUT arguments. iStudio allows a developer to define procedures that different applications can invoke, or implement. For example:

```
getAddress(IN string Name, OUT string Address) where,
getAddress - name of the procedure
Name - IN argument of type string denoting person's name
Address - OUT argument of type string denoting person's address
```

If the case of two applications being integrated (for example, SAP and CRM), the CRM application can invoke the procedure getAddress. It does not need to know or specify which application implements the procedure. The procedure may be implemented by any ERP application such as the SAP application. Applications InterConnect does the necessary mappings at run-time and invokes the appropriate ERP function using the vendor-specific bridge.

There are two styles of the request/reply paradigm:

1. Asynchronous. The procedure invoking application does not block waiting for a reply. It makes an invocation and then continues normal processing. The OUT arguments are returned to the caller using a mechanism similar to callbacks.

2. Synchronous. The procedure invoking application blocks until it gets a reply. The OUT arguments are returned to the caller as part of the invocation. The caller then unblocks and continues with normal processing.

Events

Applications may be interested in events that originate from another application. For example, when a CRM application when adds a new customer, the ERP application must be notified of the new customer. iStudio allows developers to define events that encapsulate this kind of information. Applications may publish or subscribe to events.

Events are defined by a name and the data attributes it contains. For example:

```
EVENT newCustomer CONTAINS
{
String Name,
String StreetName,
String City
}
where,
newCustomer - Name of the Event, and
Name, StreetName, City - Event data attributes
```

NOTE: Events and procedures are defined as common objects. So, the data they contain is normally a superset of the data that being integrated across applications. (See Common Data Type for more information).

Application View

Each application that uses Applications InterConnect for integration defines its own definition or view of the common objects which is called the application view. This may be different or the same as that of the common object it mirrors.

For example, an application view that corresponds to the `getAddress()` procedure may be represented in this way:

```
getAddress(IN string LastName, IN string FirstName, OUT string Street, OUT
string City) where,
getAddress - name
LastName - IN argument of type string denoting person's last name
FirstName - IN argument of type string denoting person's first name
Street - OUT argument of type string denoting person's street address
City - OUT argument of type string denoting the city
```

Application views are mapped to or from common views depending on whether the application is invoking or implementing a function module.

Mapping and Transformations

Each application's view of data may be different than the common view. In these situations the developer should specify how fields in the application view map to fields in the common view. This mapping may also involve simple transformations on the fields being mapped.

For example, the `LastName` and `FirstName` fields of the application view may need to concatenated to the `Name` field in the common view. Applications InterConnect provides a set of standard transformations. However, the developer may define custom transformations in Java which can be imported into iStudio and plugged into the runtime system.

Data Types

Data types enable users to model complex hierarchical data that is exchanged between applications. For example, a purchase order contains a header object and one or more line item objects. Both the header and line item can be defined as data types that are then used to define a PurchaseOrder data type.

Data types are also useful when you want to reuse data. For example, you can create a common data type (described below) called `customer` if you have multiple events that utilize customer information. This is a superior alternative to importing or typing the data that is transmitted with these events repeatedly.

Application Data Types

Application data types are data types which each application defines to describe its application view of data.

Common Data Type

Data types defined in the common view are referred to as common data types.

iStudio Versioning

iStudio supports versioning for Common Data Types, Application Data Type, Events and Procedures.

Comprehensive versioning support is achieved by using the concept of an owner of these objects in addition to versions. An owner is the creator of the object. Only the creator of an object can modify the object. However, other users can create new versions or copy the original object under a new name. The owner is specified at the time of repository installation.

In the following example, the metadata is being created at Oracle, and at the time of repository installation, "Oracle" was specified as the owner of the metadata. The following functionality is available for versioning:

Automatic Versioning First, an event called "NewCustomerEvent" is created. When you create this object for the first time, the assigned owner is Oracle and the version is V1. This even is `NewCustomerEvent/Oracle/V1`.

Modify Object If you are not the owner, you can change the contents of the event (the data associated with it) by clicking **Modify**, but you cannot change the version number or the name of the event. The event is `NewCustomerEvent/Oracle/V1`.

Create New Version If instead, you want to keep the original `NewCustomerEvent` but want to create a new version of the information with modified data, click **Create New Version**. When you save this version, you now have two objects — `NewCustomerEvent/Oracle/V1` and `NewCustomerEvent/Oracle/V2`.

Load Version Not all versions of objects are loaded into iStudio. To work with a specific version of an object, use the **Load Version** capability. In the scenario above, when you created a new version it became your current version. Now to load `NewCustomerEvent/Oracle/V1`, you should use **Load Version**.

Copy Object To create a `NewBigCustomerEvent` which has a lot of common elements with `NewCustomerEvent/Oracle/V1`, first load

`NewCustomerEvent/Oracle/V1` and then click **Copy Object**. Using **Copy Object** allows you to not only modify the data, but also modify the name of the event. When you have modified the name of the even, `NewBigCustomerEvent/Oracle/V1` and `NewCustomerEvent/Oracle/V1` will both coexist in the repository.

NOTE: You cannot type in a name that already exists.

In our scenario above, all the metadata was built at Oracle. Now we can transmit this metadata to a customer, NewCorp. When NewCorp installs the repository and specifies the owner as `NewCorp`, the metadata is in a read-only state. Now, if they want to customize `NewBigCustomerEvent/Oracle/V1`, they cannot modify the existing version since the owners are different. They can however, use the other features described above.

To customize the metadata, they must create a new version, so that `NewBigCustomerEvent/Oracle/`V1 and `NewBigCustomerEvent/NewCorp/V2` coexist in the repository. The client can use both events in defining messages if required, and NewCorp can now modify the event it owns.

Cross Reference Tables and Domain Value Mapping

Cross Reference Tables and Domain Value Mappings are features supported by the Applications InterConnect platform to facilitate easy creation of mapping tables at design time, and their population, lookup, and deletion during runtime.

For example, consider two heterogeneous applications A and B being integrated through Applications InterConnect. To achieve integration, there is a requirement to maintain mapping tables that track the associations between the various identifiers (ids) created by the different systems. A *customer* created on application A has an id 100, while the same *customer* created on application B, as part of the integration, has an id 10000. The association between 100 and 10000 needs to maintained for subsequent information interchanges between the two applications. For example, if the address of the *customer* with id 100 is being updated by application A, and the change is propagated to application B, the id 100 needs to be automatically transformed to 10000 to ensure the correct record is updated at B. The Cross Reference Tables address this functionality.

In case of Domain Value Mappings the population of data occurs at design time. This feature is useful when the associations between the various ids (or names) is known at design time.

For example, the order status might be represented by *BOOKED* on application A, while its equivalent on B might be *ORDERED*. The association between *BOOKED* and *ORDERED* is maintained in a domain value map for order status.

Note: For information on content-based routing and load balancing, see Chapter 3, "Runtime Concepts and Components".

Message Type

Message Type specifies the mode of communication between an adapter and an application.

Using iStudio

iStudio is the GUI-based development tool that implements the concepts described in the previous section.

This section describes how to perform the various tasks required by Applications InterConnect during the design phase using iStudio. You can perform the tasks by referring to the steps described in the following sections.

- Start the iStudio tool. The following window is displayed: (Figure 2–3)

Figure 2–3 iStudio main menu

▶▶Toolbar

Figure 2–4 iStudio tool bar options

■ You can also select the tasks by clicking the icons in the toolbar.

Creating a New Project

To define a new project, click **File,** and select **New Project** in the iStudio main menu panel. iStudio displays the New Project Dialog box: (Figure 2–5)
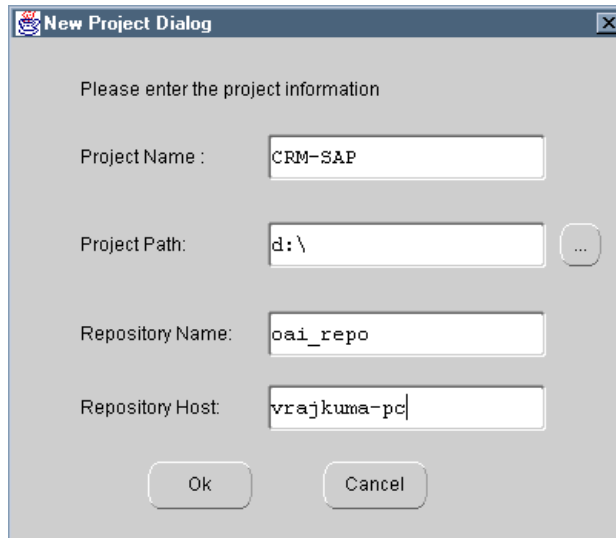

Figure 2–5 New Project dialog box

Enter the project details as follows:

■ **Project Name**—the project name for the integration.

- **Project Path**—the directory path where the project will be stored.

- **Repository Name** —the Repository where the project will be stored.

- **Repository Host**—the Repository host name.

Click **OK** to save the record, or **Cancel** to abort and exit.

Creating Workspaces

To create a new workspace, click **File** and select **New Workspace**. The New Workspace Dialog is displayed:
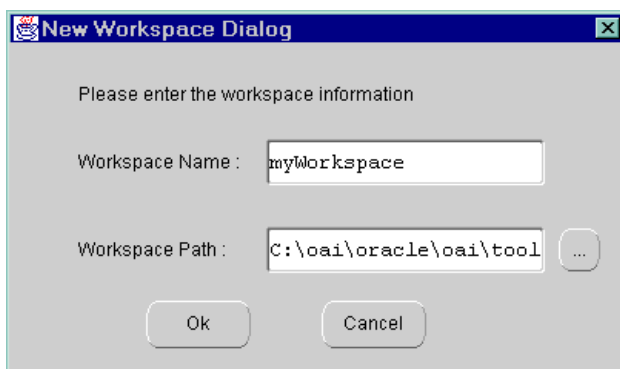

Figure 2–6   New Workspace Dialog

Enter the workspace details as follows:

- **Workspace Name**—the Workspace name for the project.

- **Workspace Path**—the Workspace directory location.

Click **OK** to save the record, or **Cancel** to exit without saving the workspace.

When you start iStudio, the default workspace **myWorkspace.iws** and the last opened project is automatically loaded. To save the SAP and DB login credits, check the box **Save settings as default** in the SAP and DB login dialogs. The user settings are automatically saved to the workspace.

Opening an Existing Workspace

To open a workspace that has already been created, click **File** and select **Open Workspace**. The file system dialog is displayed. Enter the workspace name and path to open the workspace.

Creating a Business Object

To define a new business object, click **File / New**, and select **Business Object**. The Create Business Object window is displayed: (Figure 2–7)


Figure 2–7   Create Business Object window

Enter the business object name:

■ **Business Object Name** —the name of the business object being created. (Only alphanumeric characters are allowed.)

Click **OK** to save the new business object or **Cancel** to exit.

Creating Common Data Types

To define a Common Data Type, click **File / New** and select **Common Data Type.** The Create Data Type window is displayed: (Figure 2–8)
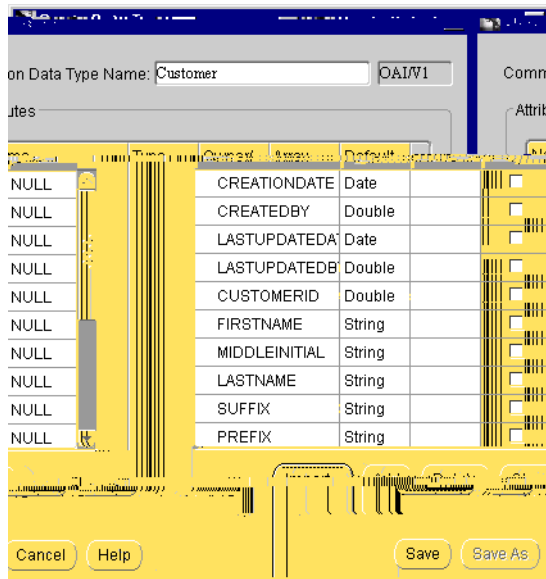
Figure 2–8   Create Data Type window

Enter common data type details as follows:

- **Common Data Type Name**—Data type name

- **OAI/V1**—the owner and version number of the Common Data Type. (It cannot be edited.)

NOTE: The first version is always V1. For subsequent versions, choose **Save As** to save it as different version number—V2, V3, V4. etc.

You can import attribute definitions from various sources. For example, you can import them from a pre-existing database table or an API Repository.

Refer to the section on Importing Attribute Information on page 15 to **Import**, **Add**, **Delete**, **Clear** the attributes.

Attributes

- **Name** —the attribute name.

- **Type** —may be integer, string, binary, float, double, or user-defined data type.

- **Array**—check this box if it is an array field. (Only user-defined data types can be of type array.)

- **Default** —the default value for the field is NULL.

Click **Save** to save and exit, **Save As** to save the Common Data Type as a different version, **Cancel** to exit without saving, or **Help** to display the help window.

▶▶Importing Attribute Information

Attributes can be imported from a database, SAP-BAPI, or SAP-IDOC. To import attributes from a database, perform these steps:

1. Click **Import** button to import the data types from a pre-existing database table.

2. Click **Database**. The database Login Info window is displayed: (**Figure 2–9**)
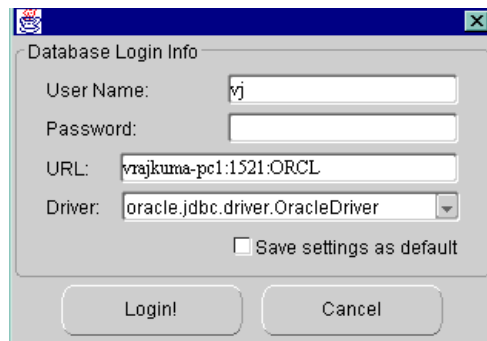

Figure 2–9   Database login info window for importing the common data types attributes

3. Enter the database log in information as follows:

- **User Name** —the log in name

- **Password** —the log in password

- **URL** —the machine name: port number: database SID

- **Driver—the** JDBC driver being used to connect to the database.

- **Save settings as default** —check this box to save the settings for the workspace.

4. After login, the database tables and arguments are displayed in the Database Browser Window. (Figure 2–10) Select **All** to import all the fields, or check the individual fields.

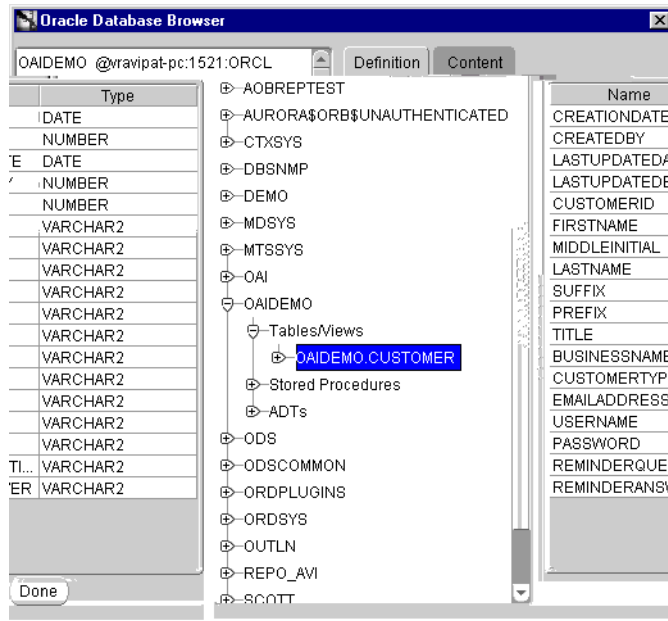5. Click **Done** to import the attributes into the common data type.

Figure 2–10   The Database Browser table to import the attributes for the common data type

6. If you select common view, all the arguments in the common view are copied over. After successful import, the arguments for the common data type is imported and populates the table. (Figure 2–11)
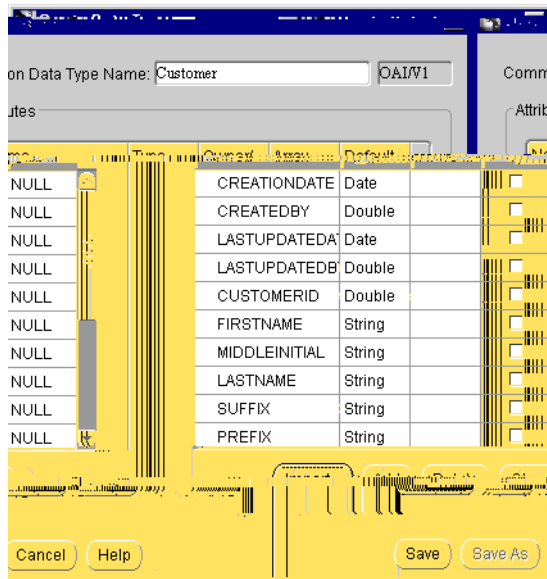
Figure 2–11   Common Data Type window with attributes

Click **Add** to add a new row of attributes, **Delete** to delete the selected attribute, or **Clear** to delete all attributes.

Creating Cross Reference Tables

To create a Cross Reference Table, click **File ⁄ New** in the pull-down menu, and select **Cross Reference Tables.** The Create X_Ref Table Dialog is displayed. (Figure 2–12)


Figure 2–12   Create X_Ref Table Dialog

1.   Enter the **Table Name**.

2. Click **OK** to save and exit, or **Cancel** to exit without saving.

Adding Applications to Cross Reference Tables

To add applications to the Cross Reference Table, select a Cross Reference Table in the project tree, right-click the Cross Reference Table, and click **Add App**. The Add Application to XRef Table window is displayed: **(**Figure 2–13**).**

Figure 2–13   Add Application to XRef Table

1. Choose an Application Name from the list.

2. Click **OK** to add the application and exit, or **Cancel** to exit without saving.

Removing Applications From Cross Reference Tables

To remove applications from the Cross Reference Table, select a Cross Reference Table in the project tree, right-click the Cross Reference Table, and click **Remove App**. The Remove Application from XRef Table window is displayed: **(**Figure 2–14**)**


Figure 2–14   Remove Application from XRef Table

1. Select the application name to remove from the list.

2. Click **OK** to remove the application and exit, or **Cancel** to exit without saving.

Creating Domain Value Mappings

To create a Domain Value Mappings table, click **File** in the iStudio main menu panel, click **New** in the pull-down menu, and select **Domain Value Mapping**. The Create DVM Table Dialog is displayed. (Figure 2–15)


Figure 2–15   Create DVM Table

Adding Applications to Domain Value Mappings

To add applications to Domain Value Mappings, select a Domain Value Mapping in the project tree, right-click the Domain Value Mapping, click **Add App**. The Add Application to DVM Table dialog is displayed. **(Figure 2–16)**

Figure 2–16   Add Application to DVM Table

1. Choose an Application Name from the list.

2. Click **OK** to add the application and exit, or **Cancel** to exit without saving.

Removing Applications From Domain Value Mappings

To remove applications from the Domain Value Mappings, select a Domain Value Mapping in the project tree, right-click **Domain Value Mapping**, and click **Remove App**. The Remove Application from Domain Value Mapping dialog is displayed.

In the Domain Value Mapping dialog, perform these two steps:

1. Choose the Application Name to remove from the list.

2. Click **OK** to remove the application and exit, or **Cancel** to exit without saving.

Editing Data in Domain Value Mappings

To add data to Domain Value Mappings, select a Domain Value Mapping in the project tree, right-click Domain Value Mapping, and click **Edit Mappings.** The Edit DVM Dialog is displayed: **(**Figure 2–17**)**
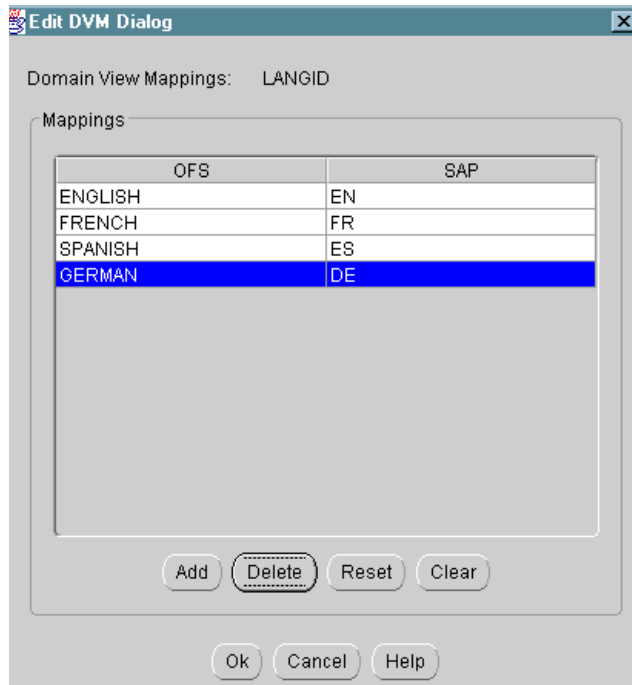

Figure 2–17   Edit DVM Dialog

Enter the mapping values and click **Add**.

To delete or clear the values in the mappings table, select a Domain Value Mapping, and click **Delete.**

To delete the DVM table, right-click the Domain Value Mapping, and click **Delete.** The DVM table is deleted**.**

To Create an Event

To create an event, click **Project (your project name)** in the iStudio main menu panel, click **New** in the pull-down menu, and select **Event**. The Create Event window is displayed: (Figure 2–18)
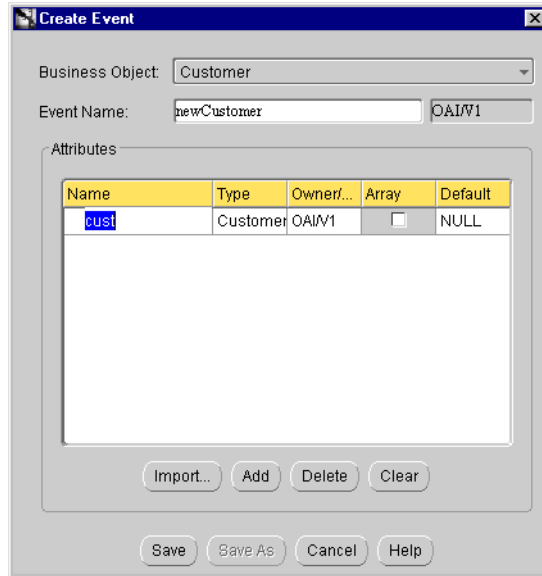

Figure 2–18   Create Event window and Common Data type selection for an attribute

Enter event details as follows:

- **Business Object Name**—the name of the category to which the event belongs to.

- **Event Name**—the event name. Only Alphanumeric characters are allowed.

- **OAI/V1**—the owner and version number of the Business Object.

Attributes

- **Name**—the attribute name.

- **Type** —Can be an integer, string, binary, float, double, data type. If you choose data type, you can select the common data type which have been defined already.

- **Array**—check this box if it is an array field. Only user-defined data types can be of type array.

- **Default** —default value for the field is NULL.

- Refer to the section on Importing Attribute Information on page 15 on how to **Import**, **Add**, **Delete**, **Clear** the attributes.

- Click **Save** to save and exit, **Save As** to save the Business object as a different version, or **Cancel** to exit without saving, or **Help** for help window.

To Create an Application

To create an application, click **File** in the iStudio main menu panel, click **New** in the pull-down menu, and select **Application**. The Create Application window is displayed: (Figure 2–19)

Figure 2–19   Create Application window

Enter event details as follows:

- **Application Name**—the name of the application.

- Click **OK** to save and exit, or **Cancel** to exit without saving.

Creating a Procedure

To create a procedure, click **File** in the iStudio main menu panel, click **New** in the pull-down menu, and select **Procedure**. The Create Procedure window is displayed: (Figure 2–20)
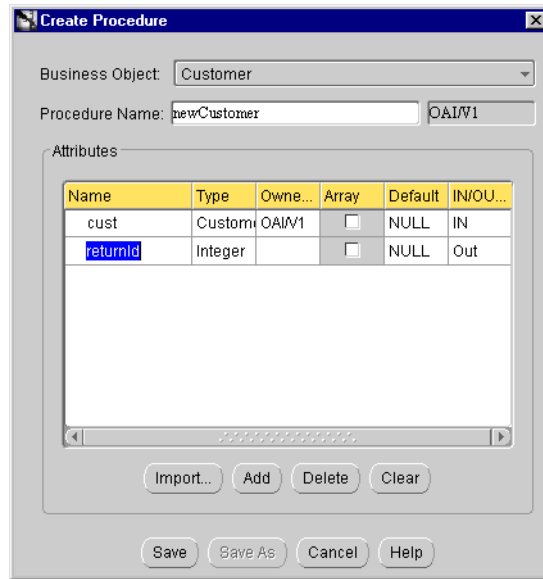
Figure 2–20   Create Procedure window with the arguments fields imported

Enter the procedure details as follows:

- **Business Object Name**—the name of the category to which the procedure belongs to.

- **Procedure Name**—the procedure name. Only Alphanumeric characters are allowed.

- **OAI/V1**—the owner and version number of the procedure.

Arguments

- **Name**—the argument name.

- **Type** —integer, string, binary, float, or double data type. If you choose data type, you can select the common data type which has been defined already.

- **Array**—check this box if it is an array field. Only user-defined data types can be of type array.

- **Default** —default value for the field is NULL.

- **IN/OUT/INOUT**—IN - input parameter, OUT- output parameter, INOUT - input and return parameter.

- Refer to the section on Importing Attribute Information on page 15 on how to **Import**, **Add**, **Delete**, **Clear** the attributes.

- Click **Save** to save and exit, **Save As** to save the procedure as a different version, **Cancel** to exit without saving, or **Help** for help window.

Creating a Publish Event

To create a publish event in an application, click **Event** in the iStudio main menu panel, then click **Publish** in the pull-down menu. The Publish Wizard window is displayed: (Figure 2–21)

The Event Publish Wizard navigates you through these four steps:

1. Select an event.

2. Specify the application view.

3. Map and transform the common objects.

4. Finish the event publish.

The steps are described in detail in the following sections.

Select an Event

Figure 2–21   window 1 Publish Wizard window - Select an Event

Enter the event details as follows:

- **Application**—the name of the application which is publishing the event.

- **Message Type**—this specifies the mode of communication between the adapter and the application.

  You can select from the following Message Types:

  Database—the adapter picks the message data from the database.

  SAP-BAPI—the adapter communicates with the application using BAPI.

  IBP—the adapter communicates with the application using IBP.

  SAP-IDOC—the adapter communicates with SAP using IDOC.

  XML—the adapter communicates with the application using XML.

- Select the event name.

- Click **Next** in the Publish Wizard. The **Define Application View** window is displayed: (Figure 2–22)

Specify the Application View

After selecting the event to publish, you define the application view. The application view window is initially an empty table. You may define the attributes using the **add**

- Click **Next** in the Publish Wizard. The **Mapping and Transformation** window is displayed: (Figure 2–23)

Map and Transform

Mapping can either involve copying the individual fields, or simple shape-change transformations.
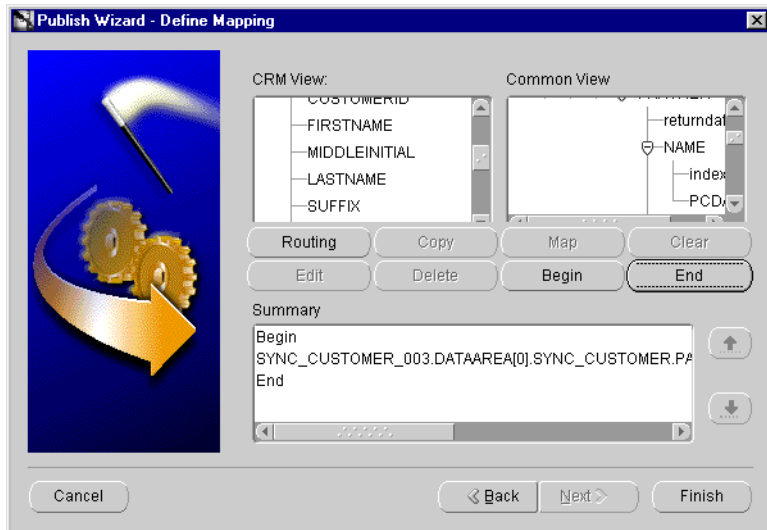

Figure 2–23   window 3 Publish Wizard - Define Mapping of Application View to Common View

- Click **Begin** to define mappings.
- Select the appropriate fields in the application view and map it to the fields in the common view. See Mapping Attributes on page 2-28 to map and copy attributes.
- Click **End** when you are done defining mappings.
- **Summary**—displays the results of the mapping selection.
- Click **Finish**.

The publish event is now created.

◗◗ Mapping Attributes

You may use the **map** button to specify transformations. For example, to map fields `FirstName` and `LastName` in the application view to `Name` in the common view, use the `concat` transform.

The following steps illustrate this example:

1. Select fields to map in both the application view and the common view.

2. **^left-click** the mouse to select multiple fields in a view.

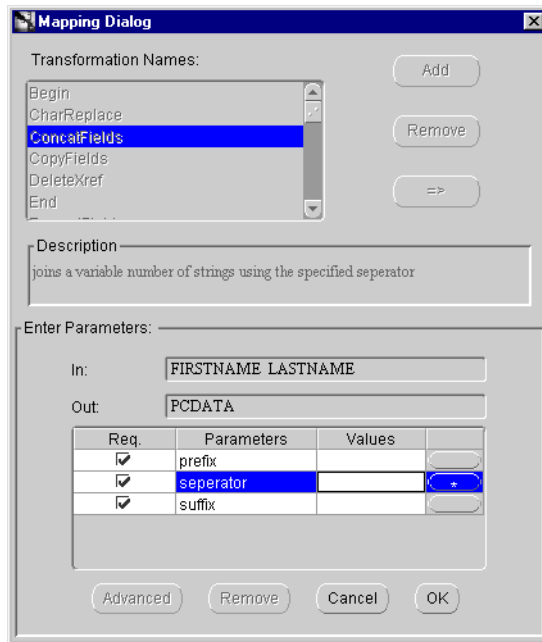3. Click **Map** button to display the Mapping dialog box: (Figure 2–24)

Figure 2–24   Mapping dialog box

4. Enter the following information:

   - **Transformation**—select the transformation to perform for the fields. (For example, ConcatFields).

   - Click => button

- **Parameters**—Enter values for the transformation parameters. (For example, (blank) value for separator parameter.)

5. Click **OK** to confirm selection or **Cancel** to exit.

▶▶Copying Attributes

1. You can copy individual fields or whole objects (which have the same attributes). While copying, no transformation occurs on the fields.

2. Select **Copy** to duplicate all the fields of the application view in the common view.

▶▶Delete Attribute Mappings

1. Click **Delete** to delete a mapping, or **Clear** to clear the existing attributes.

▶▶Edit Attribute Mapping

1. Select a mapping.

2. Click **Edit** to change the mapping.

Creating a Subscribe Event

To create a subscribe event in an application, click **Event** in the iStudio main menu panel, then click **Subscribe** in the pull-down menu. The Subscribe Wizard is displayed: (Figure 2–25)

The Event Subscribe Wizard navigates you through four steps:

1. Select an event.

2. Specify the application view.

3. Map and transform the common objects.

4. Finish the event subscription.

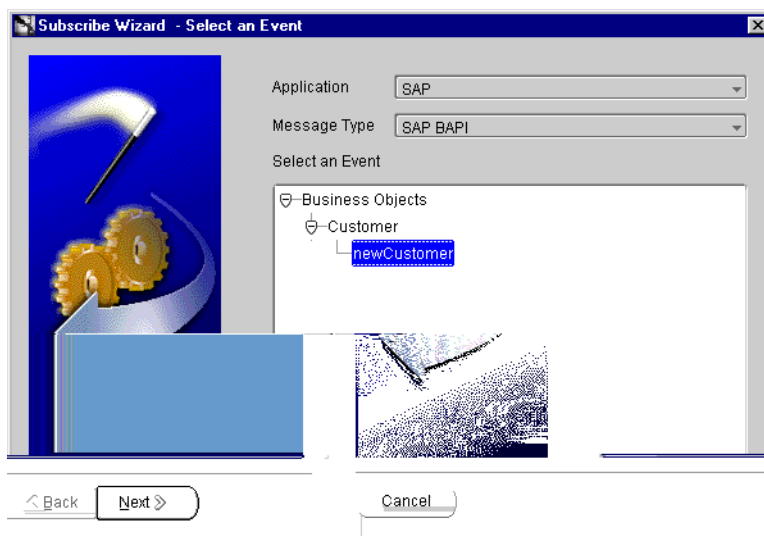The steps are described in detail in the following sections.

Select an Event

Figure 2–25   window 1 Subscribe Wizard - Select an Event

Enter the event details as follows:

- **Application (e.g.SAP)**—the application name which is subscribing to the event.

- **Message Type**—this specifies the mode of communication between the adapter and the application.

  You can select from the following Message Types:

  Database—the adapter picks the message data from the database.

  SAP-BAPI—the adapter communicates with the application using BAPI.

  IBP—the adapter communicates with the application using IBP.

  SAP-IDOC—the adapter communicates with SAP using IDOC.

  XML—the adapter communicates with the application using XML. (See Publish Message, figure 2-28).

- Select the event to subscribe.

- Click **Next** in the Subscribe Wizard to display the Define Application View window: (Figure 2–26)

Specify the Application View

After selecting the event to publish, you define the application view. The application view window is initially an empty table. You may define the attributes by using the **add** button, or importing the definitions from a database or an API Repository.
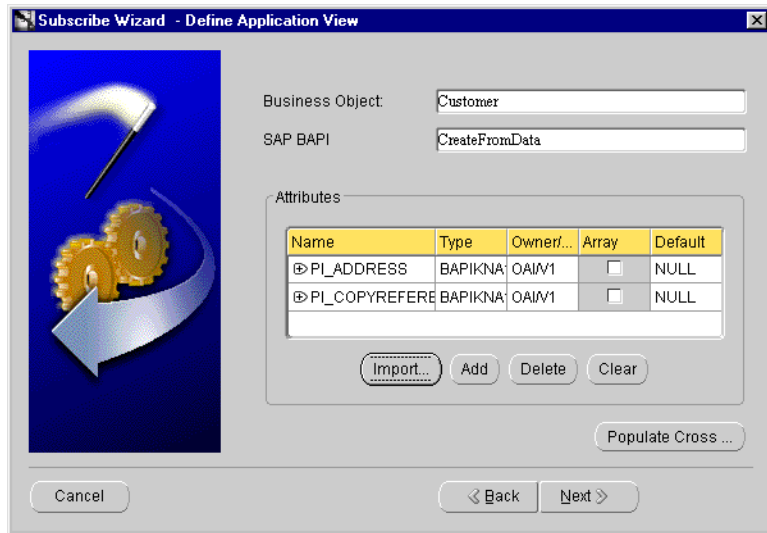

Figure 2–26   window 2 of Subscribe Wizard Event - Define Application View

Enter the application view details:

- **Business Object**—name of the business object.

- **SAP-BAPI**—the SAP-BAPI name which is imported.

Attributes

- **Name**—the attribute name.

- **Type**—may be an integer, string, binary, float, or double data type.

- **Owner/Version**—the owner and version number of this application view.

- **Array**—check this box if it is an array field. Only user-defined data types can be of type array.

- **Default**—default value for the field is NULL.

Refer to the section Importing Attribute Information on page 15 to **Import**, **Add**, **Delete**, **Clear** the attributes.

Click **Populate XRefT**, and a dialog is displayed: Enter the name of the Cross Reference Table. To populate and look up Cross Reference Table, refer to the section Populating Cross Reference Tables on page 2-33.

Click **Next** in the Subscribe Wizard to display the Mapping and Transformation window: (Figure 2–27)

Map and Transform

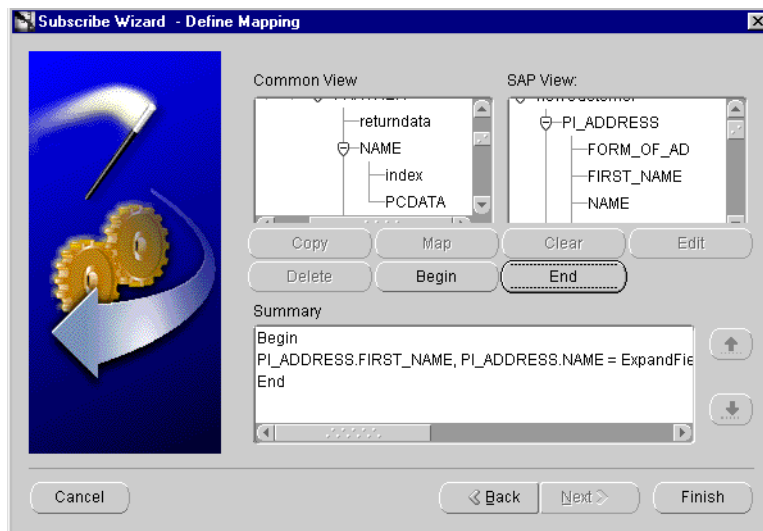Mapping can either involve copying the individual fields or simple shape change transformations.

Figure 2–27   window 3 Subscribe Wizard - Define Mapping of Common View to Application View

- Select the fields in the common view and map it to the fields in the application view. See Mapping Attributes (Common to Application) on page 2-34 to map and copy attributes.

- **Summary**—displays the results of the mapping selection.

Click **Finish**. The subscribe event has now been created.

◗◗ Populating Cross Reference Tables

To populate the Cross Reference Tables, you define a Returned Application Object which is the value returned by subscribe/implement code for populating the Cross Reference Table. The **Returned Application Object** is defined in the wizard.
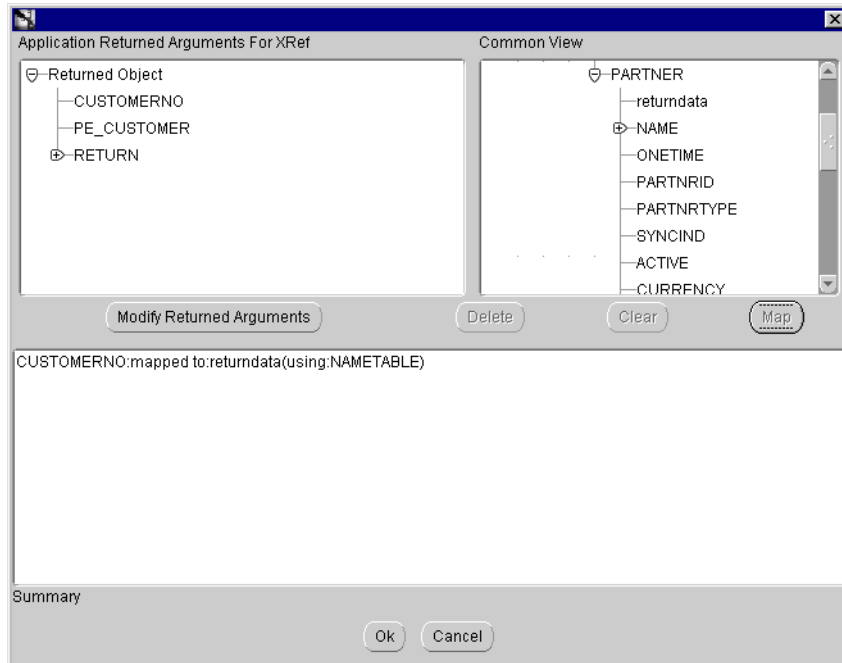

Figure 2–28

- The Returned Application Object window displays the structure of the Returned Application Object in tree format. The Returned Application Object is initially defined to be the same as the application view.

- You can modify the Returned Application Object only in Event Subscribe by clicking **Modify**. You cannot modify the Returned Application Object in Implement Procedure.

- Select the Common View of the message.

- Select the corresponding values in the two windows and click **Map**.

- Specify the Cross Reference Table name to be populated using these values.

To view the lookup mapping in the Mapping option, refer to

◗◗ Mapping Attributes (Common to Application)

You may use the **map** button to specify transformations. For example, to map fields FirstName and LastName in the common view to Name in the application view you can use the concat transform.

The following steps illustrate this example:

1. Select the fields to map in both the common view and the application view.

2. **^left**-**click** the mouse button to select multiple fields in a view.

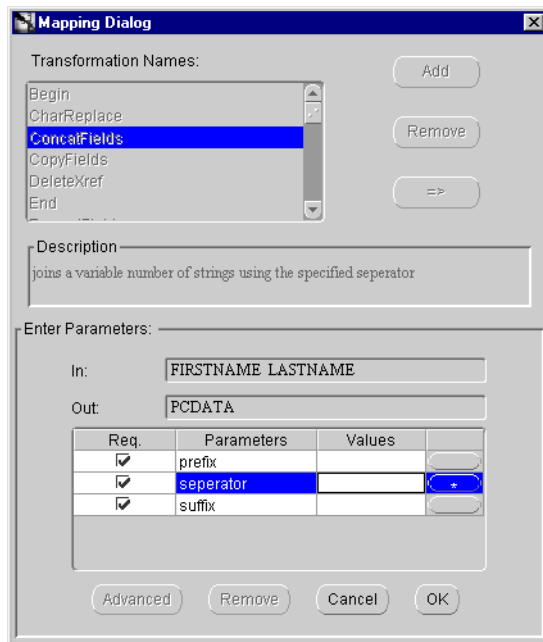3. Click **Map** to display the Mapping dialog box: (Figure 2–29).

Figure 2–29   Mapping dialog box

4. Enter the following information:

   - **Transformation**—select the transformation criteria for the fields. e.g. ConcatFields

■ **Parameters**—enter the parameters for transforming the data. (e.g., comma)

5. Click **OK** to confirm selection or **Cancel** to exit.

▶▶Copying Attributes

1. Select **Copy** to duplicate all the fields of the common view in the application view.

2. Click **Delete** to delete a mapping or **Clear** to clear all the existing attributes.

▶▶Delete Attribute Mappings

1. Click **Delete** to delete a mapping or **Clear** to clear all the existing attributes.

▶▶Lookup or Delete Cross Reference Mappings

1. Select fields to map in both the application view and the common view.

2. ^**right**-**click** the mouse to select multiple fields in a view.

3. Click **Map** button to display the Mapping dialog box: (Figure 2–24)

4. Select **LookupXRef** or **DeleteXRef** transformation depending on the functionality.

5. Enter the table name as an argument for the cross reference.

6. Click **OK**.

Creating an Invoking Procedure

To create an invoking procedure in an application, click **Procedure** in the iStudio main menu panel, then click **Invoke** in the pull-down menu. The Invoke Wizard is displayed: (Figure 2–30)

The Invoke Procedure Wizard navigates you through these steps:

1. Select a procedure.

2. Specify the application view.

3. Map and transform the application view and the common view.

4. Map and transform the common view and the application view.

5. Specify the stored procedure for database type messages.

6. Finish the procedure invocation.

The steps are described in detail in the following sections.
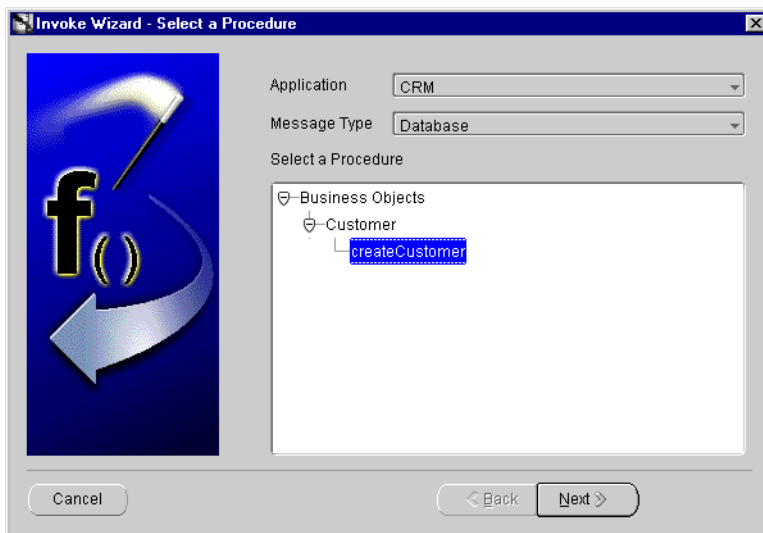
Select a Procedure

Figure 2–30   window 1 Invoke Wizard - Select a Procedure

Enter these procedure details:

- **Application (e.g.iStore)**—application name that is invoking the procedure.

- **Message Type**—specifies the mode of communication between the adapter and the application.

You can select from the following Message Types:

Database—the adapter picks the message data from the database.

SAP-BAPI—the adapter communicates with the application using BAPI.

IBP—the adapter communicates with the application using IBP.

SAP-IDOC—the adapter communicates with SAP using IDOC.

XML—the adapter communicates with the application using XML.

> **Note:** In this release of Applications InterConnect, to invoke a procedure, the message type can be a database only.

- Select the procedure to invoke.
- Click **Next** in the Invoke Wizard to display the Define Application View window: (Figure 2–31)

Specifying the Application View

After selecting the procedure to invoke, you define the application view. The application view window is initially an empty table. You may define the attributes by using the add button, or importing the definitions from a database or an API Repository.
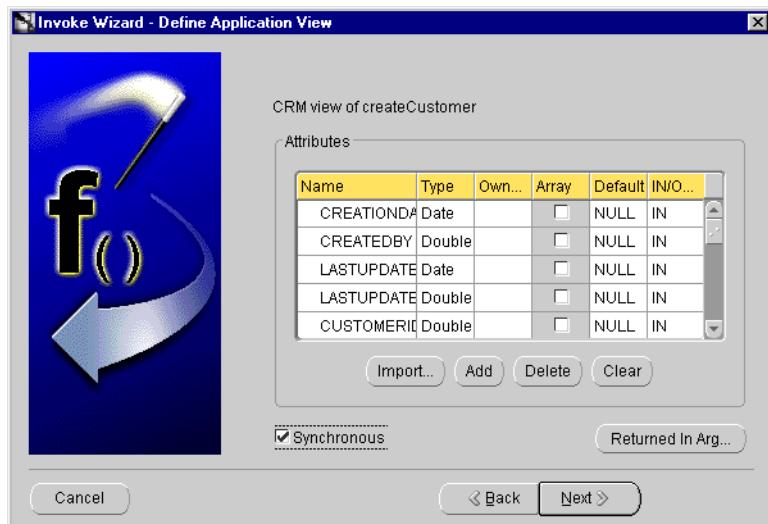

Figure 2–31   window 2 Invoke Wizard - Define Application View

Enter these application view details:

Attributes

- **Name**—the attribute name.

- **Type**—integer, string, binary, float, or double data type.

- **Owner/Version**—the owner and version of this application view.

- **Array**—check this box if it is an array field. Only user-defined data types can be of type array.

- **Default**—default value for the field is NULL.

Refer to the section Importing Attribute Information on page 15 to **Import**, **Add**, **Delete**, **Clear** the attributes.

- If you need to specify the IN arguments which are to be returned, click **ReturnedInArgs**. The **Add Returned In Args** window is displayed: (Figure 2–32)
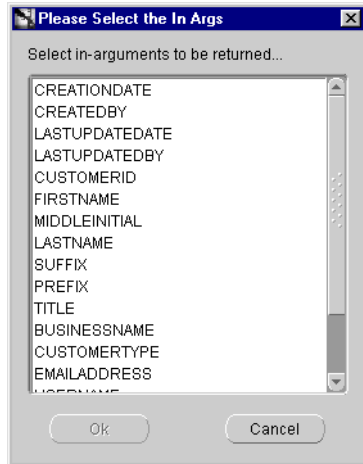

Figure 2–32   Add Returned IN arguments

- Select the input arguments to be returned. Only non user-defined input attributes are shown for selection. Click **OK** or **Cancel** to exit.

- Click **Next** in the Invoke Wizard to display the Mapping and Transformation window: (Figure 2–33)

Map and Transform (Application View to Common View)

Mapping can either involve copying the individual fields, or simple shape-change transformations.
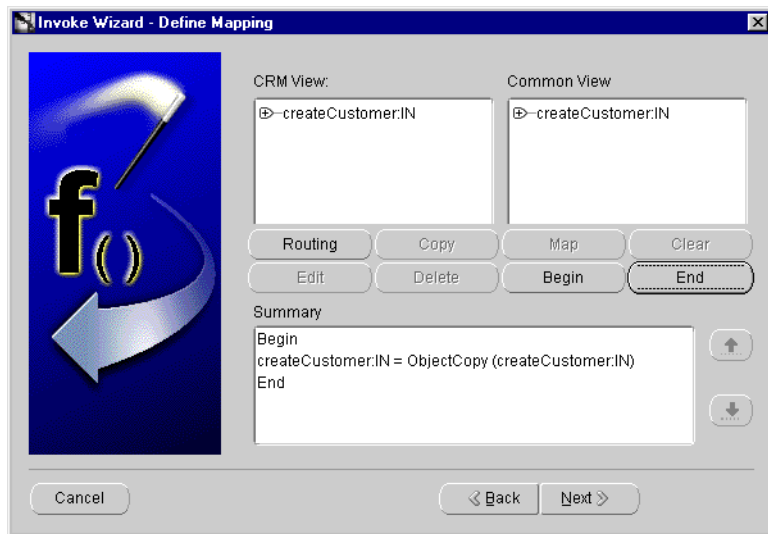
***Figure 2–33   window 3 Invoke Wizard - Define Mapping of Application View to
Common View***

- Enter the mapping details. Refer to the section on Mapping Attributes on
  page 28 and Copying Attributes on page 2-29 to map and copy attributes.

- **Summary**—displays the results of the mapping selection.

- Click **Next** in the Invoke Wizard to display the Mapping and Transformation
  (Common View to Application View) window: (Figure 2–34)

Map and Transform (Common View to Application View)

Mapping can either involve copying the individual fields, or simple shape-change
transformations. You map the common view return arguments to the application
view return arguments in this step.

Figure 2–34   window 4 Invoke Wizard - Define Mapping of Common View to Application View

- Enter the mapping details. Refer to the section on Mapping Attributes on page 28 and Copying Attributespage w
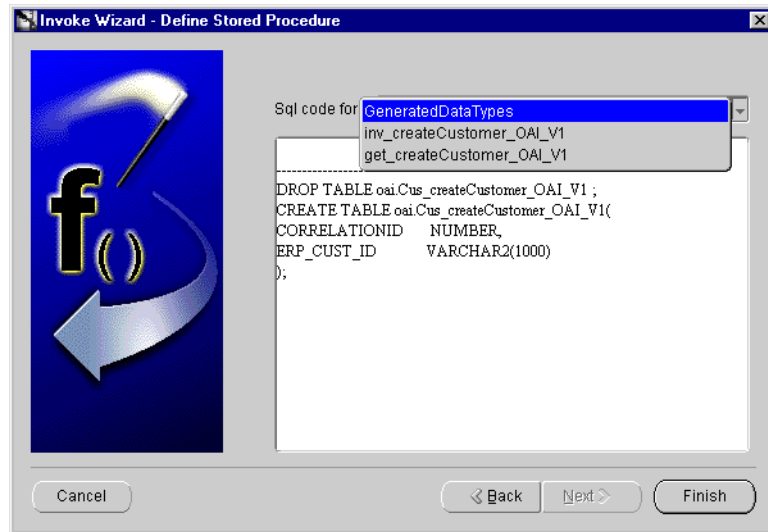
Figure 2–35   window 5 Invoke Wizard - Define Stored Procedure

Enter the stored procedure details as follows:

- **SQL Code**—select from a list of system generated procedures.

- Add new procedure code and click **Finish**.

The procedure invoke message is created.

To Create a Procedure Implementation

To implement a procedure, select **Procedure** in the iStudio main menu panel, and click **Implement** in the pull-down menu. The Implement Wizard is displayed: (Figure 2–36)

The Implement Procedure Wizard navigates you through these steps:

1. Select a procedure.

2. Specify the application view.

3. Map and transform the application view and the common view.

4. Map and transform the common view and the application view.

5. Finish the procedure implementation.

The steps are described in detail in the following sections.
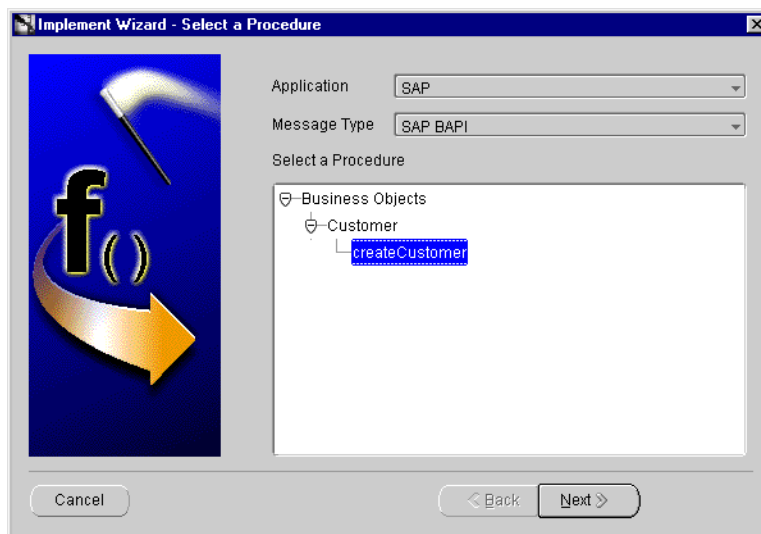
Select a Procedure

Figure 2–36   window 1 Implement Wizard - Select a Procedure

Enter these procedure details:

- **Application (e.g.SAP)**—the application name which is implementing the procedure. The example shown here is SAP.

- **Message Type**—specifies the mode of communication between the adapter and the application.

  You can select from the following Message Types:

  Database—the adapter retrieves the message data from the database.

  SAP-**BAPI**—the adapter communicates with the application using BAPI.

  IBP—the adapter communicates with the application using IBP.

  SAP-**IDOC**—the adapter communicates with SAP using IDOC.

  XML—the adapter communicates with the application using XML.

- Select the procedure to invoke.

- Click **Next** in the Implement Wizard to display the Define Application View window: (Figure 2–37)

Specifying the Application View

After selecting the procedure to implement, you define the application view. The window is initially an empty table. You may define the attributes by using the add button, or importing the definitions from a database or an API Repository.
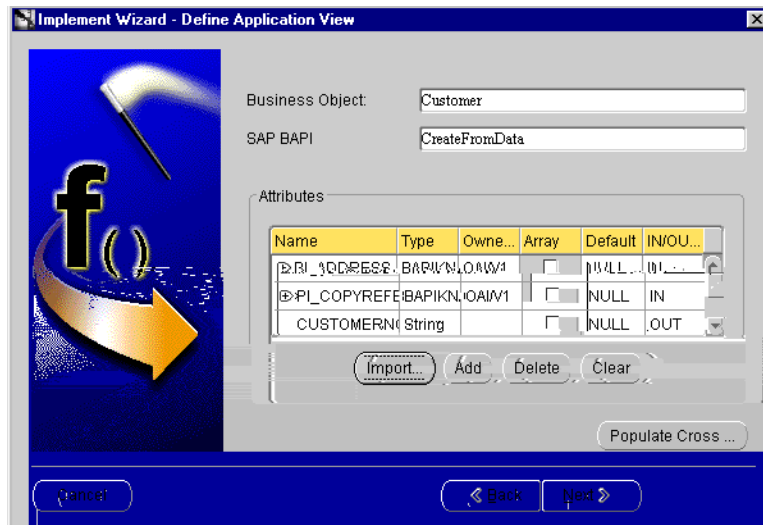

Figure 2–37   window 2 Implement Wizard - Define Application View

Enter these application view details:

- **Business Object**—name of the business object.

- **SAP-BAPI**—the SAP-BAPI name which is imported.

Attributes

- **Name**—the attribute name.

- **Type**—integer, string, binary, float, or double data type.

- **Owner/Version**—the owner and version of this application view.

- **Array**—*c*heck this box if it is an array field. Only user-defined data types can be of type array.

- **Default**—default value for the field is NULL.

Refer to the section Importing Attribute Information on page 15 to **Import**, **Add**, **Delete**, **Clear** the attributes. If you are importing BAPI information the SAP log in window is displayed: (Figure 2–38)

Click **Populate XRefT**, and a dialog is displayed: Enter the name of the Cross Reference Table. To populate and look up Cross Reference tables, refer to the section Populating Cross Reference Tables on page 2-33.
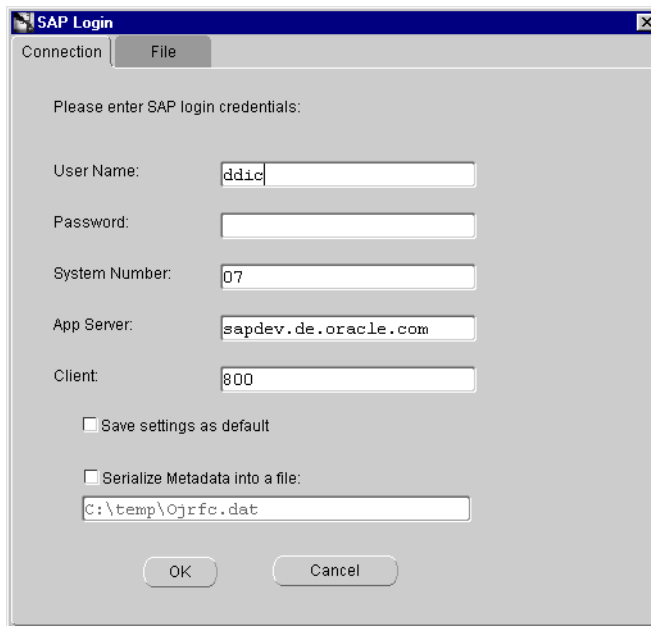

Figure 2–38   SAP login window

Enter the SAP log in information as follows:

- **User Name**—the SAP log in name

- **Password**—the password for this username

- **System Number**—the system you are logging in to

- **Application Server**—the server name

- **Client**—the client host name

- **Save settings as default**—check to save workspace settings.

- **Serialize Metadata**—check to save SAP metadata in a file.

- **Filename**—the name of the metadata file, as a fully-qualified path.

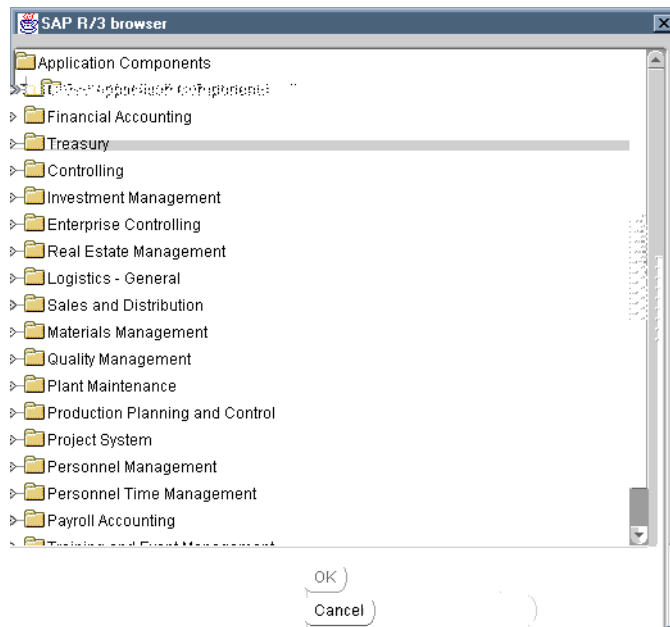- Click **OK**. The SAP application browser window is displayed: (Figure 2–39)

Figure 2–39   SAP R/3 browser window

- Select the SAP-BAPI and click **OK.** The SAP-BAPI attributes will be imported into the application view.

- Click **Next** to display the Mapping and Transformation window: (Figure 2–40)

Map and Transform

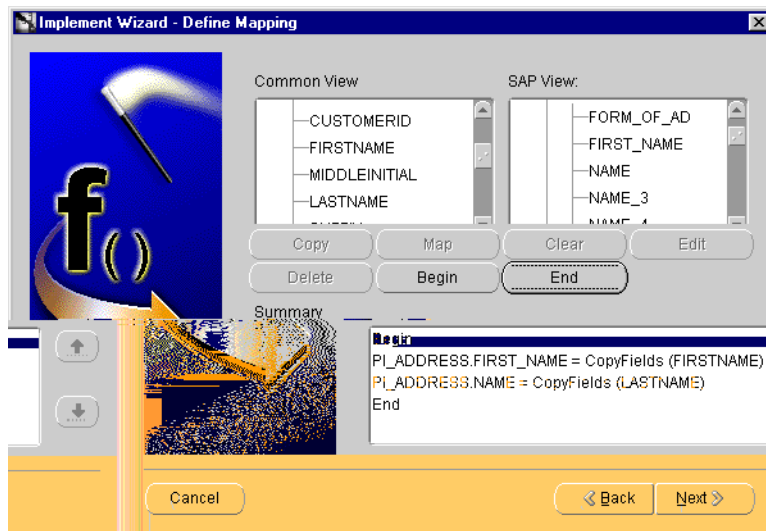Mapping may involve copying individual fields, or simple shape-change transformations.

Figure 2–40   window 3 Implement Wizard - Define Mapping of Common View to Application View

- Enter the mapping details.

Refer to Mapping Attributes on page 2-28 to map and Copying Attributes on page 2-29 to copy attributes.

- **Summary**—displays the results of the mapping selection.

- Click **Next** in the Implement Wizard to display the second Mapping and Transformation window: (Figure 2–41)
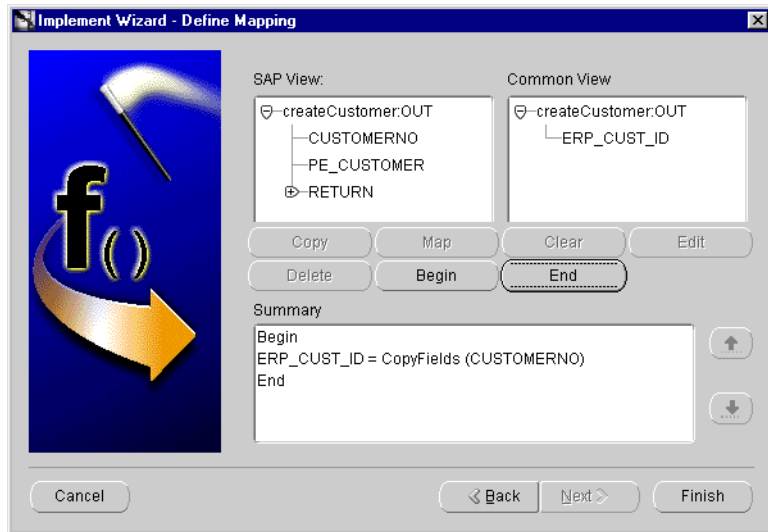
Figure 2–41   window 4 Implement Wizard - Define Mapping of Application View to Common View

- Enter the mapping details.

Refer to Mapping Attributes on page 2-28 to map and Copying Attributes on page 2-29 to copy attributes.

- Click **Finish**.

The procedure implement message is created.

Exporting Stored Procedures

iStudio generates stored-procedure stubs to enable an application to interface with the Applications InterConnect run-time easily. These stubs are exported to a file using the export functionality.

To export stored-procedures, select **File** in the main menu, then click **Export**. The Export Application dialog box is displayed: Figure 2–42
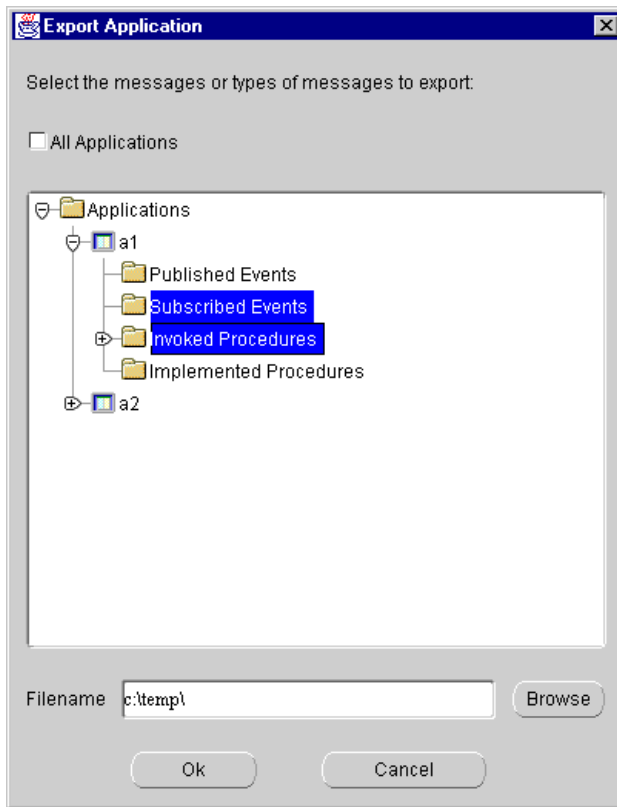
Figure 2–42   Export Application dialog

Filtering Messages

In the Export Application dialog box, you can choose which messages to export stored procedures. Messages may be filtered in any of the following ways:

- to export all messages select **Applications**.

- to export all messages of a certain type (i.e. Published Events) for all applications, check the **All Applications** check box. Then select one or more types of messages to export.

- to export all messages for a specific application, select the **Application Name**.

- to export all messages of a certain type (i.e. Published Events) for a specific application, select that **type** under the **application name**.

- to export specific messages, select the messages by **name**.

To select more than one message or class of messages **^click** the application.

- **File Name**—choose the name of the file to contain the exported stored procedures. The name you give generates multiple files. Click **Browse** to view the directory path.

- Click **OK** to export the stored-procedure, or **Cancel** to exit without storing.

The stored-procedure is now exported.

3

Runtime Concepts and Components

This chapter describes the runtime concepts, components and processes of Applications InterConnect.

This chapter contains the following sections:

Introduction to the Runtime Component

The runtime component consists of an event-based distributed messaging system. An event is any action that triggers a message. The messaging system can be distributed with different components of the system communicating over a WAN.

Features

The Applications InterConnect runtime features are categorized as follows:

- Integration architecture
- Message delivery
- Messaging paradigms
- Persistence
- Content-based routing
- Load balancing through message partitions

- Logging and tracing
- Fault tolerance

Integration Architecture

The runtime enables inter-application communication through hub and spoke integration. This methodology keeps the applications decoupled from each other by integrating them to a central hub only. The applications are at the spokes of this arrangement and are unaware of the applications they are integrating with. To them, the target is the hub. Since each application integrates with the hub, translation of data between the application and hub (in either direction) is sufficient to integrate two or more applications.

Refer to Chapter 2, "Design Time Concepts and iStudio", on page 2-1 for a graphical representation of the hub and spoke architecture.

Message Delivery

Guaranteed delivery

All messages when handed over to runtime, to be sent to another application have guaranteed delivery.

Exactly once delivery

The messages are neither lost nor duplicated. The destination application will receive each sent message exactly once.

In order delivery

The messages are delivered in the exact same order as they were sent.

Messaging Paradigms

Refer to Chapter 2, "Design Time Concepts and iStudio", on page 2-1 for an explanation of the messaging paradigms supported in Applications InterConnect.

Persistence

Messages remain in the runtime system until they are delivered. The message is deleted when each application that is scheduled to receive a specific message has done so. For auditing purposes, you can configure the system to retain all messages even after they have been delivered successfully to each application.

Content-based Routing

Content-based routing increases message delivery efficiency by routing each message directly to the intended application by examining specific elements in the data object. The diagram below illustrates content-based routing in action.
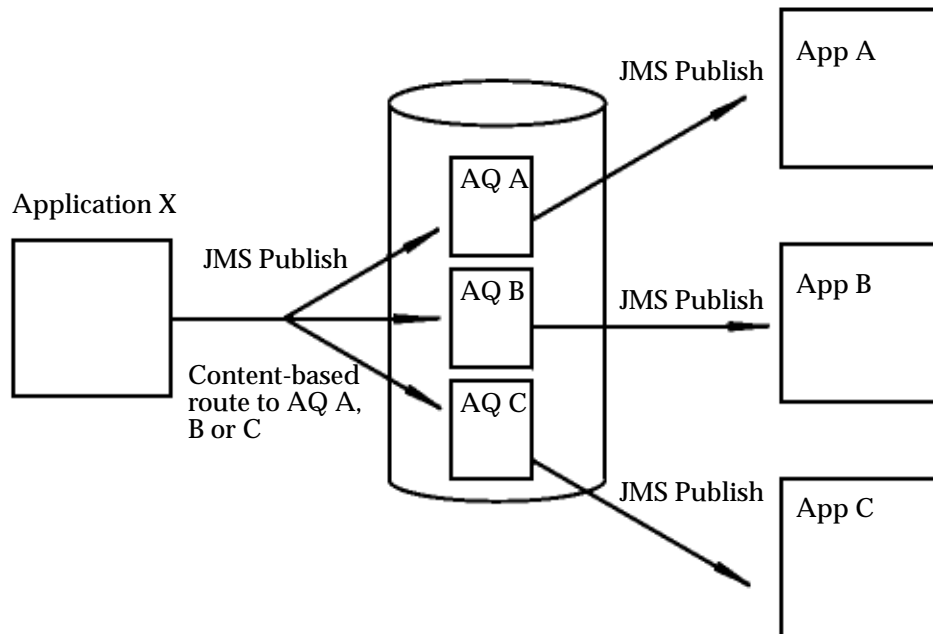
The database in the middle is the one that resides in the hub under the Oracle Message Broker (OMB). OMB has been purposely omitted from the diagram to better illustrate the functionality. Also omitted are adapters (agents and bridges) that are attached to each application.

Messages can be routed to a specific application based on specific content values contained in the message. For example, an electronic funds transaction settlement application is designed to transmit bank transactions with a specific bank code to identify the destination bank system. When the EFT application publishes each message at runtime, the Oracle Application InterConnect runtime component determines the BankCode value based on objects stored in the repository, and routes the message to the appropriate recipient system.

To implement content-based routing in this scenario, the condition set captured in iStudio may be coded in this fashion:

```
if (BankCode EQUAL A) then route message to AQ A
if (BankCode EQUAL B) then route message to AQ B
if (BankCode EQUAL C) then route message to AQ C
```

iStudio allows you to specify SQL-like routing conditions based on the message content from the publishing application. This information is stored as metadata in the repository. At runtime, the publishing agent (connected to EFT application in the example above), utilizes this information to route the message to the specific recipient application.

Default Routing Support

On a per application basis, for every published message, you can specify a hub queue in which the message should be stored. Conversely, for each subscribing application, you can specify a hub queue from which the message should be retrieved. This pairing of publish and subscribe queues constitutes the Message Capability Matrix for each application. Using this matrix, the integrator can determine which queues need to be created in the hub. This matrix is stored in the repository as metadata and is used by the agents (see runtime section) to route messages to the appropriate queue on behalf of publishing applications, and to listen for messages on the appropriate queues on behalf of subscribing applications.

Load Balancing Through Message Partitions

At runtime, for performance reasons, you may need more than one adapter attached to a specific application. For example, Application A publishes three different kinds of events—EventA, EventB, and EventC. Three potential scenarios should be examined to determine whether (and how) one or more adapters should be attached to the application to meet performance objectives:

Scenario 1 The order in which the messages are sent by Application A must be adhered to strictly for the life of all messages. For example, if Application A publishes messages in a specific order, they must be received by the subscribing applications in the exact same order (even if they correspond to different event types). In this case, you cannot add more than one adapter to Application A for load balancing.

Scenario 2 The order in which messages are sent by Application A must be adhered to but not across different event types. For example, Application A publishes the following messages in order: M1_EventA, M2_EventB, M3_EventA. M1 and M3 must be ordered with respect to each other because they correspond to the same event. However, M2 has no ordering restrictions with respect to M1 and M3.

In addition, EventA messages are transformation/size/computation heavy and EventB and EventC messages are very light. In this case, you can create message partitions from the Message Capability Matrix. Partition1 can process EventA messages, and Partition2 can process EventB and EventC messages. When you install the adapters, you specify not only the application it is attach to but also the partition it uses. These message partitions can be used to effectively load balance integrated applications.

Scenario 3 There is no message order dependency, even within the same event type. Since there are no ordering restrictions, two approaches for load balancing can be employed:

A. No message partitions are created. One or more adapters are added utilizing the entire Message Capability Matrix. This means that at runtime any one of the adapters would be available to receive any message, though only one of them would actually receive the message.

B. Message Partitions can be created based on projections of the number of messages for a particular event type. For example, if there will be three times as many EventA messages than EventB or EventC messages, you could create two partitions—one for handling EventA messages, and the other for handling the other two event types.

Logging and Tracing
The runtime provides different levels of tracing to capture all information needed for troubleshooting or monitoring. The trace information is logged to local trace files that can be read using the Runtime Management Console. For more information, refer to Chapter 4, "Runtime Management Console".

Fault Tolerance
If at any time one or more of the runtime components or applications fail, none of the messages will be lost.

Load Balancing
Adapters (see components) offer multi-threaded support for load balancing. There can be multiple adapter instantiations attached to each application. The broker may be used as part of Oracle Application Server for load balancing purposes.

Components

There are four major components in the runtime system:

- Adapter (composed of an Agent and a Bridge)
- Oracle Message Broker
- Repository

Adapters

An adapter is the Applications InterConnect component that sits at the spoke with the application to make it InterConnect enabled. Internally, the adapter is written as two components for improved reuse of existing interfaces. These components are:

Bridge. This is the application specific piece of the adapter. The bridge communicates with the specific application interface to transfer data between the application and Applications InterConnect. For messages outbound from an application, the bridge is responsible for converting the data from the application's native system format to the agent's internal format (and conforming to the application view of data defined in iStudio). It then passing it on to the agent (described below) for further processing. For inbound messages, the bridge receives the message from the agent in the agent's internal format (and conforming to the application view of data defined in iStudio). It then converts the message back to the application's native format and pushes the data contained therein into the application. Each communication protocol requires a unique bridge.

Two products using the same protocol may use the same bridge code, though at runtime two separate processes are created. The bridge is also called the technology/protocol adapter.

2. Agent. The agent is a generic engine that carries out instructions for transformations and routing captured in repository metadata (populated by iStudio). The agent does not know how to talk to a particular application. For messages outbound from the application, the agent receives the message in it's internal format from the bridge. This internal format conforms to the application view of data (see iStudio description). The agent then queries the repository for metadata to transform this message to the common view and pushes the message to OMB.

For inbound messages, the agent receives a message from OMB that conforms to the common view defined in iStudio. The agent queries the repository for metadata and transforms the message from the common view to the application view. The message is then pushed to the bridge. The agent is also know as the integration adapter.

Oracle Message Broker

Oracle Message Broker (OMB) is the message store and forward component of the runtime. Adapters send messages to OMB which stores them in an Oracle 8i database using Advanced Queuing (AQ). OMB then delivers the messages to adapters who have subscribed to them. Messages are deleted from the persistent store after each recipients who expects the messages have received them. OMB conforms to the Java Messaging System (JMS) specification for a messaging server. JMS communication between each adapter and the broker is built on CORBA.

Repository

The Repository communicates with adapters at runtime using CORBA to provide translation information for messages. This translation information is called metadata. The Repository is populated with metadata during the design phase using iStudio. Metadata customizes a generic adapter to tend to a specific application's integration reqirements.

Example
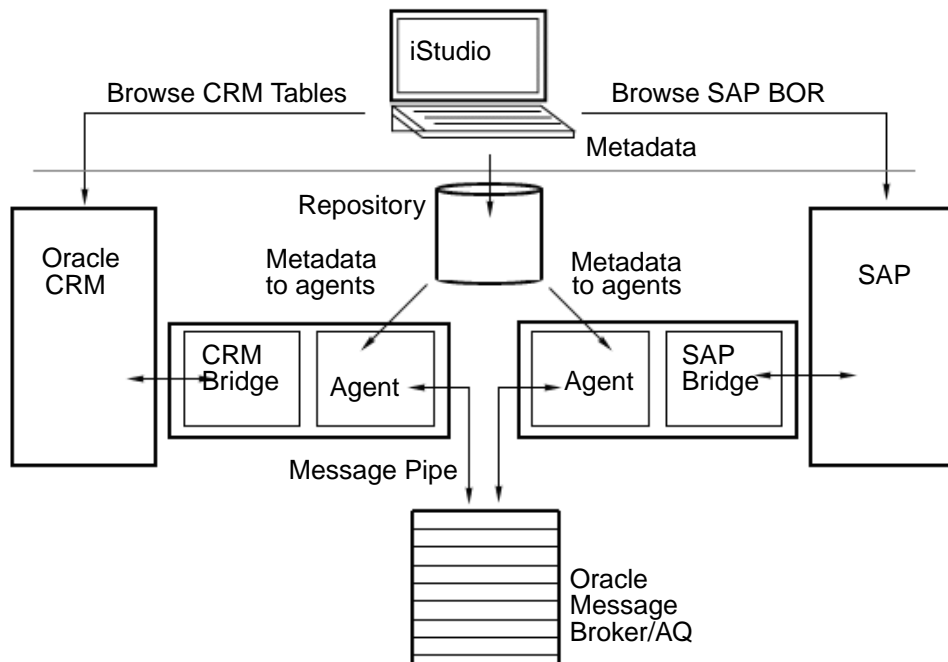
Figure 3–1   CRM Application and SAP communication - A graphical interpretation

The CRM application and SAP communication interpretation from the above diagram is as follows:

1. The CRM Adapter at startup, queries the Repository for message translation information. It is aware that it's plugged into an Oracle CRM Application, so it queries the Repository for all Oracle CRM Application-related message translation information. The metadata containing this information is cached in the CRM Adapter.

2. The SAP Adapter at startup, expresses an interest in all (or some) messages that the CRM application is publishing. It also caches all SAP-related metadata from the Repository. Note that the Repository was populated with the metadata at design time using iStudio.

3. An event occurs in the Oracle CRM application.

4. As a result of the event, the application transfers all event-related information to the bridge through bridge APIs called by the application infrastructure.

> **Note:**   This bridge was custom-developed as an extension to the agent to tailor it to support an Oracle CRM application.

5. The bridge creates a message with the event information using published agent APIs. It then transfers the message to the agent also using published agent APIs.

6. On receiving the message in question, the CRM agent looks up the metadata information in the cache and performs the necessary translations. These translations are from the Oracle CRM Application's view to the hub or common view as described above in the hub and spoke architecture section.

7. The common view message is now shipped to the broker which stores it in an Oracle 8i database.

8. The SAP Adapter (using JMS APIs) receives this message into the adapter space.

9. On receiving the message, the SAP agent reviews at its cached metadata information and translates this message from the hub view to the SAP application view.

10. It then calls into the SAP bridge APIs to transfer the message into SAP space.

> **Note:** This bridge was custom developed as an extension to the agent to tailor it to handle an SAP application.

11. The bridge calls into SAP to deliver the message contents to the application infrastructure.

4

Runtime Management Console

This chapter describes the Oracle Applications InterConnect Runtime Management Console.

This chapter contains the following sections:

-

Introduction to Runtime Management Console

Oracle Applications InterConnect console is a centralized tool to administer the Applications InterConnect runtime components, the Repository and the adapters.

Using the console, the administrator can start, stop, monitor and troubleshoot the Applications InterConnect runtime. The console can be launched from the tools menu of Oracle Enterprise Manager console, or run as a standalone application.

The Oracle Applications InterConnect runtime includes two OAI components—the adapter and Repository. The console locates each instance of Oracle Applications InterConnect components in the system using a directory service. The console also monitors the functioning of each component, and provides various statistics about the runtime. The administrator can troubleshoot the system, in case of component failures or malfunction, using the console.

Console Monitors

The Oracle Applications InterConnect Console has three monitors:

- Adapter monitor
- Repository monitor

- Error monitor

Adapter Monitor

The Adapter monitor enables you to perform the following operations on an Applications InterConnect Adapter running in a remote location:

- Stop

  Stop causes the runtime to disable the adapter from sending and receiving messages.

- Start

  Start causes the runtime to enable the adapter to send and receive messages.

- Suspend

  Suspend causes the runtime to stop the adapter from processing the messages being sent and received.

- Resume

  Resume causes the runtime to restart the adapter, and processing of messages that have been sent and received.

Repository Monitor

The Repository monitor enables you to perform the following operations on an Applications InterConnect Repository running in a remote location:

- Stop
- Start

1. Stop will cause the runtime to:
   - stops all Repository services with the exception of administration.
   - release all locks and free sessions.

2. Start will cause the runtime to:
   - start all Repository services with the exception of administration.

Error Monitor

The error monitor enables you to manage errors during runtime. All the errors in the various runtime components are written into a table in the database. Using the error monitor, the users can view each error and take necessary action.

The user can change the status of the error, add a comment, or delete the error. The error monitor screen in the console consists of a table with **NextSet** and **PrevSet** buttons.

The database table has the following columns:

- **Timestamp**—the time the error occurred.
- **Source**—the component name where the error occurred (indicated by the application name and host name).
- **Severity**—the priority of the error (Warning, Critical, Fatal).
- **Status**—the current status of the error.
- **Description**—the description of the error

The table has a fixed number of errors listed at any time. To display the next or previous set of errors, use the `NextSet` and `PrevSet` buttons respectively.

Glossary

Advanced Business Application Programming (ABAP)

Advanced Business Application Programming (ABAP) is a programming language developed by SAP for application development purposes.

Business Application Programming Interface (BAPI)

Standardized programming interface that enables external applications to access the business processes and data of the R/3 System. Business Application Programming Interfaces (BAPIs) are defined in the Business Object Repository (BOR) as methods applied to SAP business objects, in order to perform specific business tasks.

Business Object

Represents a central business object in the real world. R/3 business objects describe complete business processes. This type of encapsulation reduces complexity because the inner structure of the business object remains concealed. By invoking methods known as BAPIs (Business APIs), external applications can access and manipulate the business objects via the Internet, DCOM or CORBA.

BAPI Browser

BAPI-oriented view of the Business Object Repository (BOR). The BAPI Browser displays all the business objects, for which BAPIs have been implemented.

IDoc

A specific instance of an IDoc type.

IDoc type

The IDoc type indicates the SAP format that is to be used to transfer the data for a business transaction. An IDoc is a real business process in the form of an IDoc type. An IDoc type is described using the following components:

•a control record is the format of the control record which is identical for all IDoc types.

•one or more data records is a data record consists of a fixed administration part and a data part (segment). The number and format of the segments can be different for each IDoc type.

•status records is the status record describing the processing stages which an IDoc can pass through, and have an identical format for each IDoc type.

Example: Purchase order no.4711 was sent to a vendor as IDoc no.0815. IDoc no.0815 is formatted in IDoc type ORDERS01 and has the status records "created" and "sent".

IDoc interface

Definition of IDoc types and data interchange methods (port definition) between SAP systems and partner systems. Partner systems can be:

- EDI subsystem

- other R/3 System

- other R/2 System

- third-party software

Index

views
    application view,   2-7

W

workspaces,   2-5
    new,   2-5
    old,   2-5