



Product Documentation

ER/Studio® Data Architect

Quick Start Guide

2nd Edition

Version 9.0

Published December 2010

© 2010 Embarcadero Technologies, Inc. Embarcadero, the Embarcadero Technologies logos, and all other Embarcadero Technologies product or service names are trademarks or registered trademarks of Embarcadero Technologies, Inc. All other trademarks are property of their respective owners.

Embarcadero Technologies, Inc. is a leading provider of award-winning tools for application developers and database professionals so they can design systems right, build them faster and run them better, regardless of their platform or programming language. Ninety of the Fortune 100 and an active community of more than three million users worldwide rely on Embarcadero products to increase productivity, reduce costs, simplify change management and compliance and accelerate innovation. The company's flagship tools include: Embarcadero® Change Manager™, CodeGear™ RAD Studio, DBArtisan®, Delphi®, ER/Studio®, JBuilder® and Rapid SQL®. Founded in 1993, Embarcadero is headquartered in San Francisco, with offices located around the world. Embarcadero is online at www.embarcadero.com.

December 8, 2010

CORPORATE HEADQUARTERS

100 CALIFORNIA STREET
12TH FLOOR
SAN FRANCISCO, CALIFORNIA
94111 USA

EMEA HEADQUARTERS

YORK HOUSE
18 YORK ROAD
MAIDENHEAD, BERKSHIRE
SL6 1SF, UNITED KINGDOM

ASIA-PACIFIC HEADQUARTERS

L7. 313 LA TROBE STREET
MELBOURNE VIC 3000
AUSTRALIA

Contents

- Introducing ER/Studio Data Architect 5
 - Notice for Developer Edition Users 5
 - Application Interface 6
 - New Features 8
 - Additional Product Information 9
- Installation 11
 - ER/Studio Data Architect Installation Checklist 11
 - System Requirements 12
 - Installing ER/Studio Data Architect 12
 - Installing ER/Studio XE 14
 - Click-and-Run with Embarcadero® InstantOn™ 23
 - Licensing ER/Studio Data Architect 23
 - Registering a Trial or Workstation License at Application Startup 24
 - Registering a Workstation License after Application Startup 25
 - Registering by Phone 25
 - Registering a Networked License 26
- Tutorials 27
 - Getting Started with ER/Studio DA 28
 - Starting to Data Model with ER/Studio DA 29
 - Logical and Physical Modeling 30
 - Using Data Dictionary Domains to Populate New Entity 31
 - Establishing Relationships Between Entities 34
 - Creating and Working with Submodels in ER/Studio DA 35
 - Generating Physical Models from a Logical Model 38
 - Denormalizing the Physical Model 41
 - Finding out How an Entity Maps to the Physical Model 44
 - Documenting an Existing Database 46
 - Generating an HTML Intranet Dictionary Report 47
 - Documenting Data Lineage 57
 - Creating a Data Flow 58
 - Creating a Data Movement Rule 59
 - Defining External Source and Target Systems 61
 - Creating a Data Lineage and Transformation Visualization 63

CONTENTS

Diagram Navigation and Aesthetics.....66
 Navigating the Diagram..... 67
 Diagram Aesthetics 70

INTRODUCING ER/STUDIO DATA ARCHITECT

ER/Studio Data Architect (ER/Studio DA) is a visual modeling application used for platform-independent logical data architecture analysis and design, in addition to platform-specific physical database design and construction. Its powerful, multi-level design environment addresses the everyday needs of database administrators, developers, and data architects who build and maintain large, complex database applications and strive to consolidate, report, and re-use metadata across the enterprise.

ER/Studio DA's progressive interface and simplicity has been designed to effectively address the ease-of-use issues which have plagued data modeling and CASE tools for the past decade and more. The application equips you to create, understand, and manage the life-cycle of mission-critical database designs and business metadata within the enterprise.

ER/Studio DA is rich and customizable. ER/Studio DA offers:

Strong logical design capabilities

- The ability to spawn many physical designs from a corporate logical design
- Bi-directional model comparison and information synchronization
- Visual-Basic for Applications API for product customization
- Powerful DDL reverse engineering and generation
- Metadata import and export capabilities
- Data lineage documentation
- Sophisticated XML, HTML, and RTF-based documentation and reporting facilities.

NOTICE FOR DEVELOPER EDITION USERS

The Developer Edition of ER/Studio Data Architect provides the user with all the standard functionality of ER/Studio except for the following limitations:

- No access to the ER/Studio Repository and ER/Studio Enterprise Portal.
- Limited import and export functionality; the Import and Export Metadata Model Bridges are not supported.
- Limited database support.

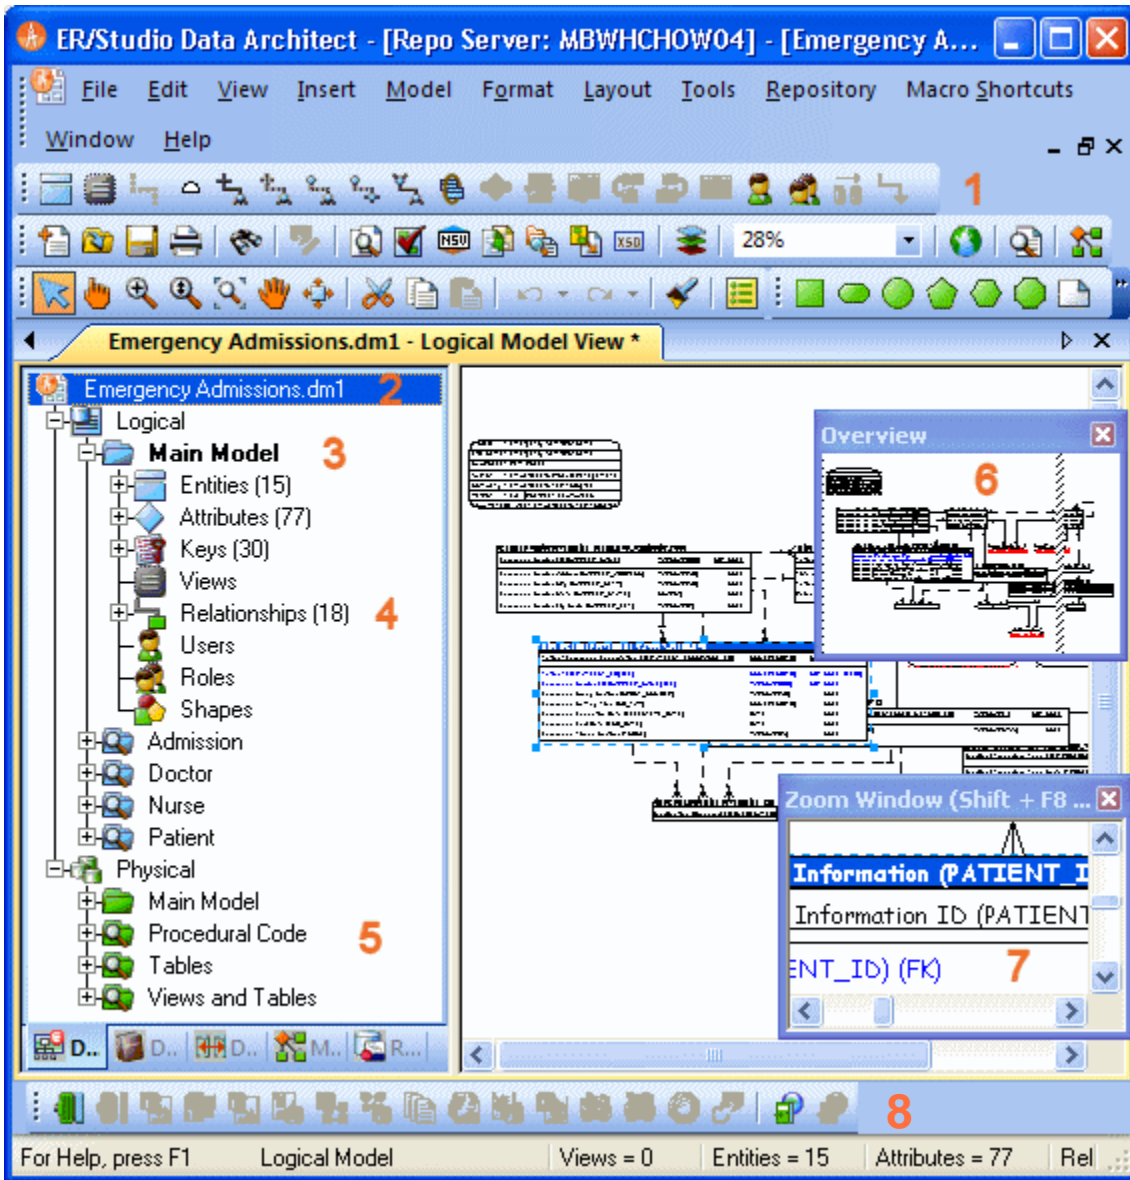
APPLICATION INTERFACE

The application interface is divided into several tabbed windows that let you navigate and customize your workspace. The application includes context-sensitive toolbars, menus, intuitive icons, keyboard commands, and other resources to help you work efficiently. ER/Studio DA's user interface has a standard Windows look and feel.

For a detailed description of the internal data structure of the ER/Studio DA application and the Repository, you can look at Metadata models (DM1 files), usually found in the following directories:

- **Windows XP:** C:\Documents and Settings\All Users\Application Data\Embarcadero\ERStudioDA_X.X\Sample Models
- **Windows Vista:** C:\ProgramData\Embarcadero\ERStudioDA_X.X\Sample Models

For information on enhanced UI features, click a number in the image below:



- 1 ER/Studio DA has an enhanced user interface that is similar to Windows XP with intuitive icons.
- 2 The name of the active model, with or without the full path to the file, can be displayed on the title bar of the main application window. This is an option you can set in **Tools > Options > Application > Application Defaults**.
- 3 ER/Studio DA displays the active model in bold type, so that you can easily see which model or submodel you are working on.
- 4 The **Data Model** tab of the Data Model Explorer displays indexes, relationships, and attributes as separate nodes.

- 5 The **Data Model** tab of the Data Model Explorer displays schema objects like packages, procedures, functions, materialized views, auxiliary tables, and triggers as separate nodes.
- 6 The **Overview Window** lets you navigate large Data Models.
- 7 The **Zoom Window** helps you focus on the details of a specific area of a large, reduced diagram.
- 8 You can dock toolbars anywhere in the ER/Studio DA user interface.

The Data Model Explorer is comprised of five tabs that offer easy access to important functionality and also includes an Explorer tree that lets you efficiently navigate the objects of your data models. The ER/Studio DA interface to the Data Model is divided into two main areas. On the left is the Data Model Explorer and on the right is the Data Model Window. This convention is carried out on other tabs of the user interface, such as the Repository and Data Lineage tabs, where you'll find an explorer on the left and a graphical representation on the right. The Data Model Window also provides a complete workspace for creating physical and logical data models and smaller, movable windows to help you navigate large data models.

NEW FEATURES

The following describes the new features of ER/Studio DA 9.0 and ER/Studio Repository 6.0.

Additional Platform Support

ER/Studio DA now fully supports Netezza 4.6 and 5.0.

User Interface Improvements

Since the denormalization feature has changed in 9.0, any documents containing denormalization mappings will be upgraded to use the new feature of ER/Studio DA 9.0. Denormalization mappings are now retained with the table data and can be viewed and edited on the Where Used tab of the Table Editor. Denormalization history, including Columns, Indexes, and Relationships of pre-map tables, is no longer retained, so the undo and make permanent functions specifically for denormalizations are no longer available. Denormalization mappings have been worked into the normal undo function of the product allowing them to be undone like other operations. However, the names of the pre-map tables are retained. Retaining only essential pre-map information means that the DM1 file with denormalizations is smaller and therefore opens faster.

For more information, see "New Features of ER/Studio DA 9.0 and ER/Studio Repository 6.0" and "Optimizing Query Performance on the Physical Model (Denormalizing the Physical Model)" in the *ER/Studio Data Architect User Guide* and online help.

ER/Studio Repository Performance Improvements

ER/Studio Repository software is now multi-threaded, enabling some processes to occur concurrently. The following describes some important rules that affect multi-threading on the Repository:

- If a Check In transaction is being processed on a given Diagram or Data Dictionary, all other operations on that Diagram or Data Dictionary must wait until the Check In completes.
- If an Add or Check In transaction is being processed, all other Add or Check In transactions must wait until the initial transaction completes. This same rule applies to Set Named Release, Branch and Branch/Merge Check In operations.
- Most other common tasks, including Log Ins, Log Outs, Gets, Check Outs and Security Center operations are processed concurrently by the server, but may be affected by database level-locking mechanisms.

ADDITIONAL PRODUCT INFORMATION

The Embarcadero Web site is an excellent source for additional product information, including white papers, articles, FAQs, discussion groups, and the Embarcadero Knowledge Base.

Go to www.embarcadero.com/support, or click any of the links below, to find:

- [Documentation](#)
- [Online Demos](#)
- [Technical Papers](#)
- [Discussion Forums](#)

INSTALLATION

This section will help you ensure your environment meets the technical requirements and helps you install ER/Studio Data Architect or ER/Studio Data Architect XE. It is comprised of the following topics:

- [ER/Studio Data Architect Installation Checklist](#) on page 11
- [System Requirements](#) on page 12
- [Installing ER/Studio Data Architect](#) on page 12
- [Installing ER/Studio XE](#) on page 14
- [Licensing ER/Studio Data Architect](#) on page 23

ER/STUDIO DATA ARCHITECT INSTALLATION CHECKLIST

Use the following checklist to ensure you correctly install and configure ER/Studio Data Architect.

Done	N/A	Requirements
		Download ER/Studio Data Architect You can download a 14-day trial version of ER/Studio Data Architect, from www.embarcadero.com/downloads . Embarcadero will send you a license registration code by email.
		Review ER/Studio Data Architect ReadMe For the most current installation and usage information, see the ReadMe at docs.embarcadero.com .
		Ensure System Compatibility For system requirements, see System Requirements on page 12.
		Install ER/Studio Data Architect For installation details, see Installing ER/Studio Data Architect on page 12.
		License ER/Studio Data Architect For ER/Studio Data Architect licensing details, see Licensing ER/Studio Data Architect on page 23.
		Run and Test ER/Studio Data Architect For step-by-step instructions on using ER/Studio Data Architect, see Tutorials on page 27.

SYSTEM REQUIREMENTS

Review the requirements before you install ER/Studio Data Architect. Adhering to these requirements optimizes ER/Studio Data Architect performance.

ER/Studio Data Architect Requirements		
Hardware	Processor	Pentium III, 600 MHz or higher
	RAM	1 GB
	Disk Space	340 MB
		175 MB in the System TEMP directory for temporary file storage during installation.
	DVD Reader	
Operating System	Any of the following: <ul style="list-style-type: none"> • Windows 7 (32-bit and 64-bit modes) • Windows Server 2008 • Windows Vista • Windows XP Professional • Windows 2003 (32-bit mode) • Windows 2000 Pro Server • Advanced Server with Service Pack 2 	
Software	To use an Open Data Base Connectivity (ODBC) database as a data source, install the ODBC driver on your client computer. To connect to the DBMS using native connectivity, install the native client libraries available from the DBMS vendor.	
Privileges	Local administrator privileges (required for installation only).	

INSTALLING ER/STUDIO DATA ARCHITECT

NOTE: To ensure compatibility, the release dates of ER/Studio Repository and ER/Studio Data Architect must be the same.

- 1 Log on to Windows with local administrator privileges.
- 2 Start the ER/Studio Data Architect installation program.
- 3 Walk through the installation wizard following the onscreen prompts.

The following provides additional information for some options that are not self-explanatory:

- **Default Notation:** Choose the notation type you want ER/Studio Data Architect to use by default to display the model.
 - **IDEF1X:** Data modeling technique used by many branches of the United States Federal Government.
 - **IE (Martin/Finkelstein):** Information Engineering (IE) as developed by Martin and later revised by Finkelstein.
 - **IE (Crow's Feet):** Uses IE notation and represents relationships with connecting lines between entities, and symbols at the ends of those lines to represent the cardinality of the relationship. This is the notation used in Oracle texts and in other applications such as Visio. Foreign keys are displayed in entity boxes.
 - **IE (Hide Foreign Keys):** Uses standard IE notation, but foreign keys are not displayed in entity boxes.

NOTE: You can change the notation setting at a later time. To change the notation, choose Tools > Options from the ER/Studio Data Architect Main menu, and then click the Logical or Physical tab where you can choose the notation type.

- **Shared Application Data Location:** By default, ER/Studio Data Architect installs shared directories and application files such as the DatatypeMapping, Macros, and Sample Models directories in the following locations:

- **Windows XP:** C:\Documents and Settings\All Users\Application Data\Embarcadero\ERStudio_X.X

- **Windows Vista:** C:\ProgramData\Embarcadero\ERStudio_X.X

If you prefer to store these files in another directory, follow the procedure in [Changing the Location of ER/Studio Data Architect Shared Files](#) on page 13.

- For step-by-step instructions on how to use ER/Studio Data Architect, see the *ER/Studio Data Architect Evaluation Guide* at docs.embarcadero.com.

Changing the Location of ER/Studio Data Architect Shared Files

NOTE: You must have local administrator privileges to change the location where ER/Studio Data Architect stores and searches for shared application data.

- 1 Choose **Help > About ER/Studio Data Architect**, and then click **File Path**.
- 2 In the **File Path** dialog, browse to an existing directory where you want to store your shared files.

NOTE: The directory path you choose must be the location where you have created an Embarcadero\ERStudioDA_X.X directory or where you want ER/Studio Data Architect to create this directory.

ER/Studio Data Architect creates C:\Documents and Settings\All Users\Application Data\Embarcadero\ERStudioDA_X.X directory, if it does not already exist.

- 3 In the **File Path** dialog, click **OK**, and then in the **About ERStudio** dialog, click **OK** again.
ER/Studio Data Architect searches for the datatype mappings in this directory.
- 4 Choose **Tools > Options > Directories**, specify where you want to store your macros, and then click **OK**.
- 5 Copy the contents of the `DatatypeMapping`, `Macros`, and `Sample Models` directories to the new location.
- 6 To effect the changes, restart ER/Studio Data Architect.

INSTALLING ER/STUDIO XE

This guide tells you how to install Embarcadero ToolCloud, including the Embarcadero License Center (ELC), and get them up and running for use with ER/Studio XE. It assumes you are familiar with Embarcadero XE product descriptions.

For the latest update of this document, see http://docs.embarcadero.com/products/all_access/.

- [Step 1: Obtaining the Installer](#) on page 14
- [Step 2: Installing the ToolCloud and ELC](#) on page 15
- [Step 3: Logging into the ToolCloud Admin Console](#) on page 16
- [Step 4: Setting up Licenses in the ToolCloud](#) on page 17
- [Step 5: Named or Concurrent Licenses?](#) on page 17
- [Step 6: Setting up Named Network Users in the ToolCloud](#) on page 18
- [Step 7: Setting up Named Network Users in the ELC](#) on page 19
- [Step 8: Downloading XE Applications to the Product Repository](#) on page 20
- [Step 9: Distributing the ToolBox to Product Users](#) on page 20
- [Step 10: Configuring Product Catalog Updates](#) on page 22
- [Step 11: Changing the ToolCloud Administrator Password](#) on page 22

NOTE: For complete information on using the ToolCloud, see the *ToolCloud Administration Guide*, *Embarcadero License Center Administration Guide*, and *ToolBox User Guide*, available from the Start menu or your Embarcadero representative. For wiki FAQs, see <http://docwiki.embarcadero.com>.

Step 1: Obtaining the Installer

In most situations, you will receive an email containing instructions including how to download the installer.

Step 2: Installing the ToolCloud and ELC

- 1 Quit all programs.
- 2 Launch the ToolCloud installer and click **Next**.
- 3 Read the terms of the license agreement, and if you agree select **I Agree** and click **Next**.
- 4 If you do not have Embarcadero License Center (ELC) installed on your machine, click **Next**. Otherwise, clear the box and click **Next**.
- 5 The ELC installer opens. Click **Next**.
 - a Select **I accept the terms of the License Agreement** and click **Next**.
 - b Accept the default installation folder and click **Next**.
 - c Accept the default **Typical** installation and click **Next**.
 - d If there is no recent Java VM already installed, install the Java VM specifically for ELC. Otherwise, select **Choose a Java VM** already installed on this system, and click **Search**. Select a pre-existing VM from the list and click **Next**. *Note: The host name of the Java VM is the name used when creating .slip license files for Embarcadero products.*
 - e Accept the default Program Group for the icons and click **Next**.
 - f Review the Summary and click **Install**.
 - g When the installer prompts for configuration, click **Next**.
 - h Accept the check box to enable logging in ELC and click **Next**.
 - i Select the check box to run the ELC as a Windows service and click **Next**.
 - j The installer displays the Server Host Name and Server IP Address, which are later used by the license hosting site to create configuration files to manage licenses. Write down this information for later use and click **Next**.
 - k Select the check box to start the ELC and click **Next**.
 - l You want to configure the Reporting Engine, so click **Next**.
 - m The final ELC installation screen displays the directory where the application was installed. Note this for future reference. Click **Done**. The ToolCloud installer returns.
- 6 Accept the default installation directory and click **Next**.
- 7 Provide a new valid listening port for network communication. Write down the port number for later use.

NOTE: Do not use 8005 or 9001. The ToolCloud automatically uses ports 8005 and 9001 for Tomcat and other internal processes.
- 8 The ToolCloud distributes products through a shared repository created during installation. The Installer grants read-only access rights to everyone in order to provide a shared network directory from which products can be distributed. Normally you should accept the default and click **Next**. If there are specific restrictions for creating share directories, you can choose an existing shared directory.
- 9 The installer is ready to install. Click **Next**.

- 10 The installer prompts you to enable automatic updates. **Do not enable updates at this time.** You can configure this option later. Click **Next**.
- 11 Click **Next** to install the ToolCloud. Click **Finish** to complete.

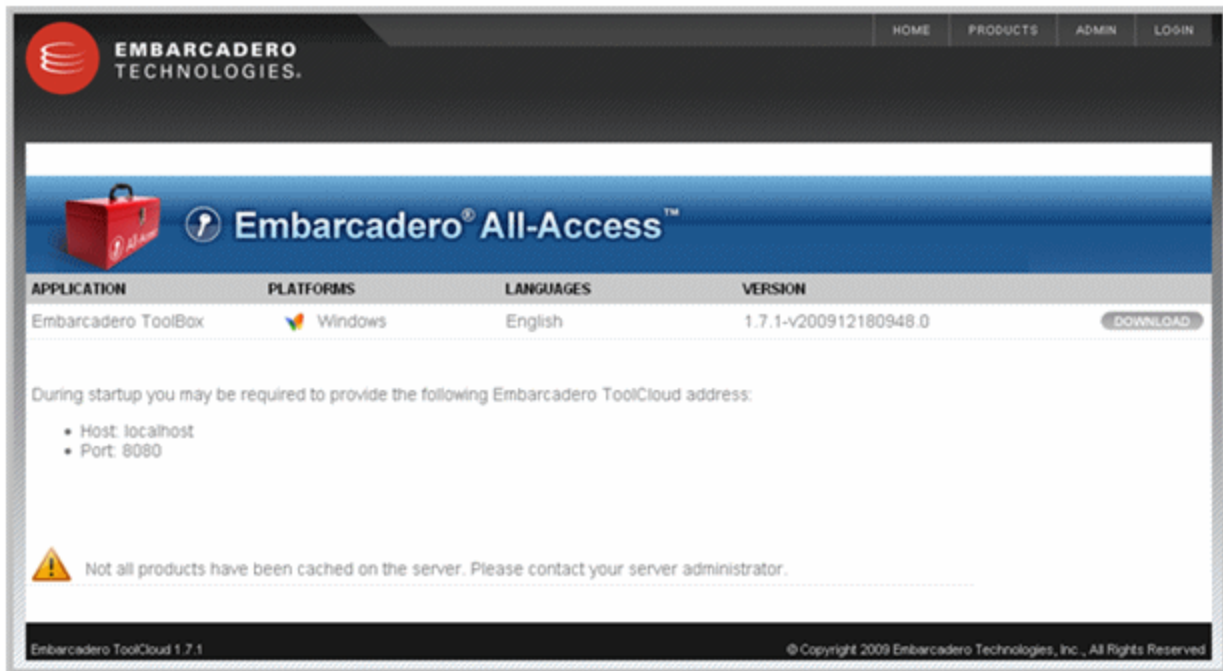
Step 3: Logging into the ToolCloud Admin Console

- 1 From the machine onto which you installed the ToolCloud, open a browser and go to the URL:

`http://localhost/aaserver`

If you gave it a port number other than 80 during installation, use:

`http://localhost:<port_number>/aaserver.`



- 2 At the top of the ToolCloud **HOME** page, click **ADMIN**.
- 3 Enter your login credentials. By default, the initial login is "admin" and the password is "admin".
- 4 Click **Login**. The **ADMIN** page opens.

Step 4: Setting up Licenses in the ToolCloud

- 1 In the **LICENSES** menu of the **ADMIN** page, click **Retrieve Licenses**.

Retrieve License

* Login Name

* Password

Certificate Number(s)

Leaving this field empty will retrieve all network licenses assigned to your account and associate them with this machine. If you have been assigned multiple license certificates and would like to enable only a subset on this server, then you may enter one or more comma-separated Certificate numbers, to retrieve just those licenses.

ELC Host Address

In case the server host name cannot be resolved to an IP address on your company network, enter either a fully qualified address (e.g. hostname.hostdomain.com) or an IP address (e.g. 123.123.12.21). This address will be used by clients to send licensing requests to the server.

- 2 Enter your login ID and password, as supplied to you by Embarcadero.
- 3 Enter the license certificate number, as supplied to you by Embarcadero.
- 4 Enter the ELC host address name or IP number, as you previously noted in Step 2.
- 5 Click **Retrieve**.

The license or licenses will be registered with the ToolCloud and the ELC.

Current Licenses

Default level Custom

License retrieval finished without any errors

License Title	Certificate	Server Name	File Location	Expire Date	
DBArtisan XE Professional with Analyst Pack	86971	SFWDELLO1	named_86971.slp	03/31/2010	

Step 5: Named or Concurrent Licenses?

If you have named network licenses, proceed to Steps 6 and 7. You must set up users in both the ToolCloud and the ELC.


If you have concurrent network licenses, you do not need to set up specific users. When a user connects to the ToolCloud from the ToolBox, the ToolCloud will provide access to its concurrent licenses. You can skip to Step 8.

Step 6: Setting up Named Network Users in the ToolCloud

For named network licenses, you add a user or users to the ToolCloud.

- 1 On the **ADMIN** page, click **Import Users**.
- 2 Select **Create user**.

- 3 Enter a network domain and user name.
- 4 If you are using All-Access, select an All-Access level to match the type of license you want to provide to the user. If you select the default, **None**, then the default access level specified in the **Current Licenses** page will apply.
- 5 Click **Create User**. The ToolCloud displays the user you have created.



Step 7: Setting up Named Network Users in the ELC

For named network licenses, you need to add the exact same user or users to the ELC as you added to the ToolCloud in Step 6.

- 1 On the **ADMIN** page, click **Embarcadero License Center**.
- 2 Enter your login credentials. By default, the initial login is "admin" and the password is "admin".
- 3 Select **Named Users**.
- 4 Click **Add entry**.


The screenshot shows the 'Named Users' management interface. A modal dialog titled '+ Add new entry' is open over a table. The table has columns: User Name, Host Name, Product Title, Product Id, Sku Id, Version, and Offline usage. The dialog fields are: User name: sqlninja, Host name: ANY, Product title: DBArtisan XE Professional (dropdown), Product id: 14002, Sku id: 64, Version: 0, Offline usage: 1 day (dropdown). At the bottom of the dialog are 'Add' and 'Cancel' buttons. Below the dialog, in the main window, are buttons for 'Add entry', 'Save to file', 'Undo changes', and 'Import from LDAP'.

- 5 Enter a user name that matches one in Step 5.
- 6 Click **Add**.
- 7 Click **Save to file**.

Step 8: Downloading XE Applications to the Product Repository

In the ToolCloud, the **ADMIN** page lists products along with their status.

- 1 Click the name of the product you want. The **Product:...** page opens.
- 2 Click **Cache Product version** or **Cache all versions**.



Step 9: Distributing the ToolBox to Product Users

Product users access software from the ToolCloud using the ToolBox. Have them follow these instructions, and substitute the correct values for the ToolCloud host name and port:

- 1 Connect to the ToolCloud at:

```
http://<serverhostname>:<portnumber>/aaserver
```
- 2 The ToolCloud **HOME** page lists the ToolBox. Click **Download**.
- 3 Save the ToolBox InstantOn to disk; the location depends on your browser setting.

- 4 After the copy completes, click **Yes** to launch the ToolBox. Subsequent launches will originate from the local machine.

Products that are not available are grayed out in the ToolBox interface.

Click the wrench icon to configure the InstantOn launch and installation options.



Step 10: Configuring Product Catalog Updates

Set the ToolCloud behavior for obtaining product updates from Embarcadero.

Product catalog updates

Embarcadero regularly releases new versions of All-Access subscription products. Please choose how to update them.

Automatic (recommended)
 Notify me but don't automatically download updates
 Turn off Automatic updates

Do this at

- 1 In the **PRODUCTS** menu of the ToolCloud **ADMIN** page, click **Update Product Catalog**.
- 2 Select the option you want for getting updates.
- 3 If you select **Automatic** or **Notify...**, specify the frequency and time for notifications.
- 4 Click **Save**. If you selected **Automatic** or **Notify...**, the **NOTIFICATIONS** button will display the number of new messages in red text. Click **NOTIFICATIONS** to see the messages.

Step 11: Changing the ToolCloud Administrator Password

Most commonly, you will have internal policies about software passwords. Using your company guidelines, change the Administrator password.

- 1 At the top of the ToolCloud **HOME** page, click **CHANGE PASSWORD**.
- 2 Enter the new password.
- 3 Click **Save**.

NOTE: Do not forget the Administrator password, as the system cannot recover it for you. Write it down and store it in a secure location.

CLICK-AND-RUN WITH EMBARCADERO® INSTANTON™

In addition to the standard installation executables for ToolCloud products, Embarcadero provides InstantOn versions of each product. InstantOn installers provide click-and-run functionality for a product, to simplify deployment and enable side-by-side versioning. InstantOn executables are also ideal in locked-down desktop environments because they do not affect system files or system registry settings.

InstantOn files use a technique called application virtualization to run within their own space without sharing DLLs or system-wide settings that could conflict with other applications. Application data, including configuration and license information, is stored locally on the hard disk, separate from the executable itself. InstantOn executables run on the local machine and are launched from either the local hard drive or the network using the ToolBox.

In a ToolCloud-centered environment, also called network mode, InstantOn applications are provisioned and licensed by the ToolCloud (leveraging the Embarcadero License Center) and then accessed using the ToolBox. In this scenario, the ToolCloud can obtain software updates from Embarcadero servers and make them available to the ToolBox.

In a ToolBox-centered environment, also called workstation mode, the ToolBox downloads products directly from Embarcadero. Users license the products using the Embarcadero Developer Network (EDN). See the *All-Access ToolBox User Guide* for more information.

Using InstantOn executables on a shared ToolCloud avoids the need to keep multiple machines updated with new versions. By keeping the ToolCloud version up to date, users will automatically gain access to the latest versions of the products.

InstantOn executables launched from the ToolCloud benefit from application streaming, because the application will begin to open before the product has been fully downloaded. However, a user who needs frequent access to an application should run the file locally, which is typically faster than running it from the ToolCloud.

LICENSING ER/STUDIO DATA ARCHITECT

ER/Studio Data Architect requires one or more licenses in order to run. A baseline license provides basic feature support for the product. In addition, incremental licenses may be required to support specific DBMS platforms, product add-ons or other functions.

ER/Studio Data Architect supports three types of licenses:

- **Trial licenses:** A license for a 14-day, full-featured trial version of the product. The trial license must be registered before you can use the product. For more information, see [Registering a Trial or Workstation License at Application Startup](#) on page 24.
- **Workstation licenses:** A license or set of licenses that is tied to a particular workstation. The product can only be used on that workstation. For more information, see [Registering a Workstation License after Application Startup](#) on page 25.

- **Networked licenses:** Networked licenses are administered and distributed by a central License Server (Embarcadero License Center or Acresto FLEXnet Publisher). There are two types of networked licenses:
 - **Concurrent licensing:** Users on different machines take turns using licenses from a shared pool.
 - **Networked Named User licensing:** Licenses are pre-assigned to specific users that are setup on the license server's user list. Those users are guaranteed to have licenses available any time.

For more information, see [Registering a Networked License](#) on page 26.

If you have problems registering, see [Registering by Phone](#) on page 25.

REGISTERING A TRIAL OR WORKSTATION LICENSE AT APPLICATION STARTUP

Shortly after downloading a trial version of an Embarcadero product, you should receive an e-mail from Embarcadero with a serial number, which must be registered as follows. These instructions also apply if you have purchased a product and received a serial number prior to running the application for the first time, or if the trial period has expired.

- 1 Start the application.

The Embarcadero License Registration dialog appears.

- 2 Copy the serial number from the e-mail, and then paste it into the **Serial Number** field.
- 3 Enter your Embarcadero Developer Network (EDN) account credentials in the **Login or Email** and **Password** fields.

If you have not previously created an EDN account, or have forgotten your password, click **I need to create ...** or **I've lost my password**.

- 4 Click **Register**.

The activation file should be downloaded and installed automatically. If this does not happen, click the **Trouble Connecting? Try Web Registration** link in the license registration dialog and follow the prompts. If you still have problems, see [Registering by Phone](#) on page 25.

REGISTERING A WORKSTATION LICENSE AFTER APPLICATION STARTUP

The following instructions assume that you have received a workstation license by e-mail and that you currently have a valid trial license. If you did not install a trial version or if the trial period has expired, follow the instructions in [Registering a Trial or Workstation License at Application Startup](#) on page 24 instead.

- 1 Choose **Help > Launch License Manager**.
- 2 In the **License Manager** dialog, select the version of ER/Studio Data Architect the license applies to: **Standard**, **Developer**, or **Enterprise**.
- 3 Choose **Serial > Add**.
- 4 Copy the serial number from the e-mail, paste it into the **Add Serial Number** dialog, and then click **OK**.
- 5 Expand **Unregistered Serial Numbers**, right-click the serial number just added, and then click **Register**.

A registration dialog appears.

NOTE: The Registration Code box shows a machine-specific identifier required with other registration methods.

- 6 Ensure that **Register using Online Registration** is selected.
- 7 Provide Developer Network credentials in the **Login name or Email** and **Password** fields.
- 8 Click **Register**.

The activation file should download and install automatically. If this does not happen, click the **Trouble Connecting? Try Web Registration** link and follow the prompts. If you still have problems, see [Registering by Phone](#) on page 25.

REGISTERING BY PHONE

If you have problems with either of the above procedures, you can register licenses by phone. You must provide the following information:

- Developer Network credentials
- The registration code displayed in the Embarcadero License Registration dialog that appears when you start an unlicensed application
- The product base license serial number
- The license serial numbers for any additional features you have purchased.

For North America, Latin America, and Asia Pacific: call (415) 834-3131 option 2, and then follow the prompts. The hours are Monday to Friday, 6:00 A.M. to 6:00 P.M. Pacific time.

For Europe, Africa, and the Middle East: Call +44 (0)1628-684 494. The hours are Monday to Friday, 9 A.M. to 5:30 P.M. U.K. time.

Shortly after phoning in, you will receive an e-mail containing an activation file. Then do the following:

- 1 Save the activation file to the desktop or to a scratch directory such as `c:\temp`.
- 2 From within the application, choose **Help > Launch License Manager**.
- 3 Choose **License > Import**.
- 4 Navigate to the directory where you saved the activation file, and then double-click the activation file.
- 5 Click the **Import** button to import the activation file, and then click **Finish**.

REGISTERING A NETWORKED LICENSE

If your organization uses networked licensing, an administrator, department head, or someone providing a similar function, will provide you with activation files: `license.dat` and `license.txt`. In the ER/Studio Data Architect About dialog, ensure that the File Path field contains the correct location of these files.

Once you receive the files, save them to the license subdirectory of the main installation directory (typically `C:\Program Files\Embarcadero\<product><version>\license\`), then restart the application.

No additional steps are necessary.

TUTORIALS

The tutorials are intended to help you get started using Embarcadero's data modeling and database design solution, ER/Studio Data Architect (ER/Studio DA).

After completing these tutorials, you'll have the foundation you need to explore the many features and benefits of ER/Studio DA. You will have learned how to do the following:


- Create a new data model.
- Work with logical and physical diagrams.
- Leverage productivity-focused features such as powerful reporting engines.
- Use common tasks and commands to make you more productive.

The tutorials are divided into nine sessions. Do them all at once or complete them individually as your time permits.

- [Getting Started with ER/Studio DA](#) on page 28
- [Logical and Physical Modeling](#) on page 30
- [Documenting an Existing Database](#) on page 46
- [Documenting Data Lineage](#) on page 57
- [Diagram Navigation and Aesthetics](#) on page 66

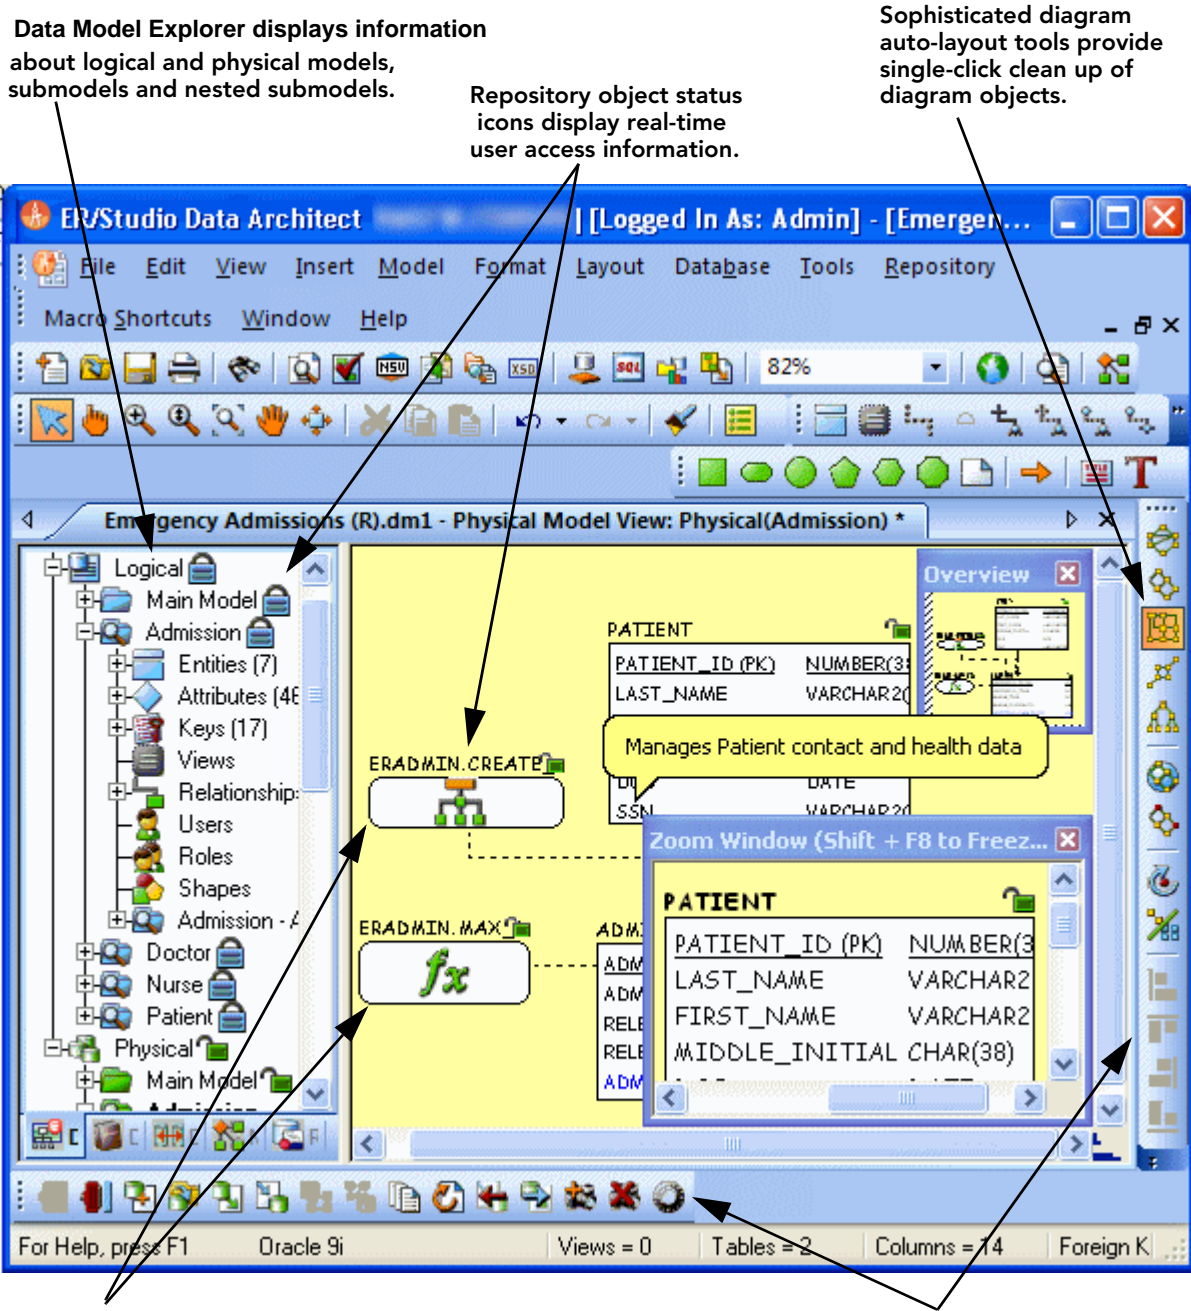
You can use this basic tutorial as a road map of product highlights, but also to help you find your own path in exploring ER/Studio DA.

Once you have started, from the Main menu you can click Help to find many additional resources that complement and build on many of the activities shown in this brief guide.



GETTING STARTED WITH ER/STUDIO DA

The graphic below names and describes the functionality of some key elements of the ER/Studio DA user interface.



Data Model Explorer displays information about logical and physical models, submodels and nested submodels.

Repository object status icons display real-time user access information.


Sophisticated diagram auto-layout tools provide single-click clean up of diagram objects.

Complex schema objects like functions can be displayed to illustrate dependencies.

Toolbars are dockable anywhere in the ER/Studio DA application window.

STARTING TO DATA MODEL WITH ER/STUDIO DA

- 1 On the Windows **Start > Programs** menu, choose **Embarcadero > ERStudio Data Architect**.
- 2 Choose **File > New > Draw a new data model**.




As you can see in the **Create a New Model** dialog, there are a number of ways to begin modeling with ER/Studio DA:

- Build a new design from the ground up by drawing a new data model.
- Build a data model from an existing database through live reverse engineering.
- Import designs from other modeling products such as Computer Associate's ERwin or SQL files.

TIP: You can select an initial layout style for your model before the SQL import takes place.

- 3 Ensure **Relational** is selected for the new model type and then click **OK**.

After selecting **Draw a new data model** and clicking **OK**, ER/Studio DA will resemble the image below:



LOGICAL AND PHYSICAL MODELING


ER/Studio DA supports both logical (non-DBMS or technology-specific) modeling and physical (DBMS-specific) modeling. ER/Studio DA is designed to allow organizations the flexibility to analyze and design a business problem or application logically and generate as many different physical interpretations from the logical model as required. Multiple physical models can be generated from the logical model for the same DBMS (for example, Oracle) or other DBMSs (such as Oracle, SQL Server and DB2). Generating logical and physical models will be discussed in detail in the following sessions.

- [Using Data Dictionary Domains to Populate New Entity](#) on page 31
- [Establishing Relationships Between Entities](#) on page 34
- [Creating and Working with Submodels in ER/Studio DA](#) on page 35
- [Generating Physical Models from a Logical Model](#) on page 38
- [Denormalizing the Physical Model](#) on page 41
- [Finding out How an Entity Maps to the Physical Model](#) on page 44

USING DATA DICTIONARY DOMAINS TO POPULATE NEW ENTITY

As instructed in [Getting Started with ER/Studio DA](#) on page 28, you have chosen to draw a new data model to begin a logical model from the ground up. Before we begin to add entities, we will populate ER/Studio DA with some sample domains (Domains are re-usable attributes).

- 1 Choose **File > Import Data Dictionary**.



- 2 Next to the **File Location** box, click the ellipsis and browse to the `Sample Models` folder, which is located at:

- **For Windows XP:**

`C:\Documents and Settings\All Users\Application Data\Embarcadero\ERStudioDA_X.X\Sample Models`

- **For Windows Vista:**


`C:\ProgramData\Embarcadero\ERStudioDA_X.X\Sample Models`

- 3 Double-click the `Orders.dm1` sample model and then click **OK**.


This model contains a pre-populated, sample data dictionary.

NOTE: Under Resolve Imported Objects with Duplicate Names, you can choose between a couple of options to determine how the dictionary objects are imported. This is important when importing into a diagram that already has dictionary objects in it.

Once opened, you will see that the ER/Studio DA Model Explorer has automatically switched to the **Data Dictionary** tab to allow immediate drag and drop access to domains.




- 4 Now, to add an entity to the Data Model Window, click the **Entity** tool on the **Modeling Toolbar**,




and then click in the Data Model Window to drop the entity.

TIP: The cursor will change to an entity symbol once the **Entity** tool is clicked, so you can click to drop as many entities on the Data Model Window as you want.



- 5 Right-click to return your mouse to the selection cursor.

- 6 In the entity name field, type `Customer`, replacing the default entity name, `Entity1`.



- 7 In the **Domains** folder of the **Data Dictionary** tab, locate the **ID** domain in the **Numeric Domains** folder.
- 8 Click the **ID** domain (do not release your mouse), drag it onto the **Customer** entity, and then release it just below the entity's name, which is the entity's Primary Key field.

TIP: You can edit or rename an entity and insert, edit or rename attributes by holding down the **Shift** key and then clicking the name or attribute. Pressing the **Tab** key cycles between the entity's name, and primary key and non-primary key fields. After naming the field, press **Return** to insert a new field.


TIP: If you need to zoom in on the entity to read the entity name and attributes, press **F8** to view the **Zoom Window**, then use the scroll bars of the **Zoom Window** to center the entity on the Data Model Window. Press **Shift** while rolling the mouse wheel forward to increase the view magnification. You can use the **Pan** tool to reposition the view to better see the entity.

- 9 In the entity, click **ID**, the name of the attribute we just created from the ID domain, and change its name to `CustomerID`, as seen in the illustration above.
- 10 Repeat the process in step 7 and step 8 to populate the Customer entity with the following domains:
- **Name** and **Phone** from the **Property Domains** folder
 - **Address, City, State, Zip Code** from the **Address Domains** folder.
- 11 Drop another entity on the **Data Model Window** and call it `Order`.
- 12 Drag the ID domain onto the Order entity's Primary Key field, change the ID domain name to `OrderID`.
- 13 Right-click the Order entity and select **Comments**.
- 14 Enter some text in the **Comments** editor, click **Add** and then click **OK**. You can add comments to any database object in the same manner. This can be very useful when you want to provide feedback on a model to other users.

15 Save your data model. We will use it in the next session of this tutorial.

ESTABLISHING RELATIONSHIPS BETWEEN ENTITIES

- 1 On the **Modeling toolbar**, click the **Non-Identifying, Mandatory Relationship** tool.



NOTE: The screen shots of this tutorial were taken using a model where the notation was set to IE (Crow's Feet). Depending upon the notation your model is currently set to, the icons for each relationship will be slightly different.

TIP: You can change the model notation by choosing **Model > Model Options**, and then selecting another notation option in the Notation area.

- 2 To establish a relationship between Customer and Order, click the parent entity, **Customer** and then click the child entity, **Order**.

NOTE: ER/Studio DA supports sound design practices by automatically propagating the primary key, from parent to child entities. If there are candidate alternate keys that should also be propagated to the child, in the **Relationship Editor** you can choose all available parent entity keys in the Parent Key list. Deleting a relationship will remove a non-native propagated attribute. However, if you want to keep the child columns of the relationship or foreign constraint, when you delete the relationship you can check the Make Foreign Keys Native option. In this case, if the relationship between Customer and Order is deleted, the CustomerID will be left in the Order entity.

What is a Domain? A Domain is a valuable tool in establishing standard, re-usable Attributes/Columns. They allow data modelers to create a data element once (such as an ID field you require all of your entities to leverage as its primary key) which has the same data type, definition, rule, and constraint no matter where the data element is distributed and bound. Read more about Domains in ER/Studio DA Help.

CREATING AND WORKING WITH SUBMODELS IN ER/STUDIO DA

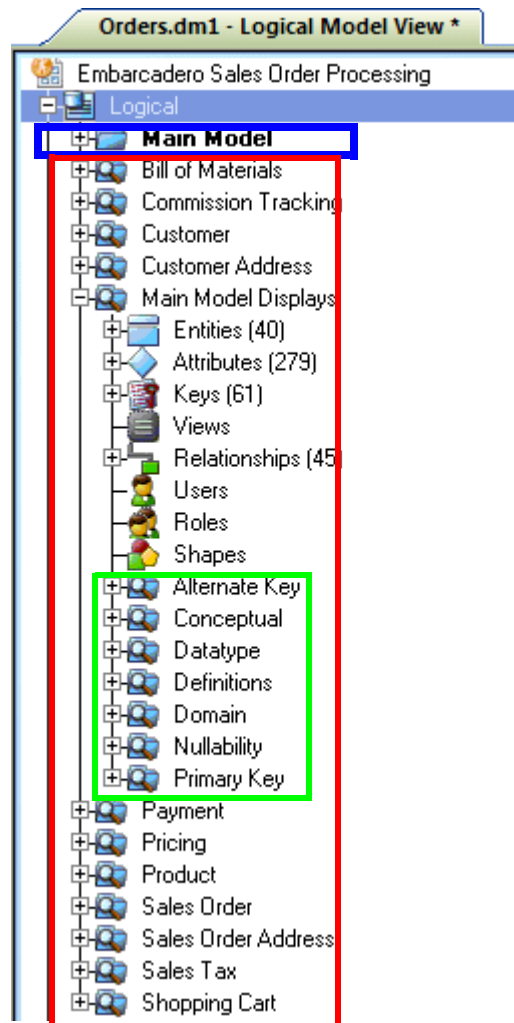
Now that you have a general understanding of how to build logical models from the ground up in ER/Studio DA, it is important to understand how to work with an important navigation feature of ER/Studio DA, called Submodels. Submodels and nested submodels are designed to break down large, complicated views of a data model in order to focus on a specific area. An important aspect of Submodels to understand is that any changes made in the submodel...*other than layout, color, display settings, notation, etc. which can be unique to the submodel...*will occur automatically in the Main Model view. In other words, change or add an attribute to an object in a Submodel and the change is automatically propagated to its Main Model counterpart.

Close the current sample model, and open a more mature model. Use this exercise to learn more about submodeling.

- 1 Choose **File > Open**.
- 2 Select **Orders.dm1** and then click **Open**.
- 3 To preserve this sample model for future use, choose **File > Save As** and then save the `Orders.dm1` file with a new name.

In this exercise, we will be modifying this model.

- 4 Collapse the folders in the **Data Model** tab of the Data Model Explorer to look like the image below:




In the `orders.DM1` sample model, there are no physical models. This model includes several submodel folders that help to describe the logical model:

- **Main Model** – This is the entire collection of all logical objects in the `Orders.DM1` file. Note the absence of the magnifying glass on the folder icon which designates it as the main model.
- **Bill of Materials** through **Shopping Cart** – These are submodels, which are smaller collections of entities derived from the Main Model that help to describe specific areas of the Main Model free from other entities.
- **Alternate Key** through **Primary Key** – These are nested submodels, which can go 'n' levels deep and are literally submodels of submodels.


Feel free to explore. Click the plus sign (+) to expand these folders.

Let's create a submodel with all the objects related to the `orders.DM1` Address components.

- 5 To make a new submodel, navigate to Logical Main model and with the **CTRL** key depressed, click the objects in the Data Model Explorer, as seen in the image below.




NOTE: Any objects selected in the Data Model Explorer will also be selected on the Data Model Window as well (also seen here). You can also select entities by lassoing them on the Data Model Window. With the entities selected, choose **Model > Create Submodel**.



- 6 Below **Create Submodel**, enter `Address Components` as the name for the submodel.
- 7 Click **OK**.

ER/Studio DA creates the **Address Components** submodel.



What do the results look like and how do I navigate to the submodel? Once created, you will see the new submodel listed in the Data Model Explorer, denoted as a submodel by the magnifying glass over its folder, as in the case with Bill of Materials and the other submodels.

GENERATING PHYSICAL MODELS FROM A LOGICAL MODEL

ER/Studio DA can generate as many physical models from a single logical model as desired. There are many ways to leverage multiple physical models in ER/Studio DA to help the design process. Examples of how multiple physical models are used are:

- Managing change in an existing application: Maintain independent development, test, and production physical model diagrams that represent specific databases.
- Migrating database applications: Use ER/Studio DA as an analysis and design hub for migrating database applications. Manage a physical model of the legacy source database application in addition to its new target physical model, which can be for an entirely new DBMS than originally maintained in the legacy database.


Let's generate a new physical model from a logical model in order to build a database. We will use the `Orders.DM1` sample model.

- 1 Open your version of the `Orders.DM1` sample model.

TIP: Use the steps shown in the last session to do so.

- 2 Select the **Main Model** and then click **Model > Generate Physical Model**.

ER/Studio DA invokes a step-by-step wizard to walk you through the process of generating a DBMS-specific physical model.



- 3 Name the new physical model, `DB2 Physical Model` and then select **DB2 UDB for OS/390 8.x** as the target DBMS to generate.
- 4 Continue through the **Generate Physical Model** wizard, which prompts very clear and concise questions about how you want your physical model to be generated.

NOTE: The wizard prompts you to customize items such as individual object selection, index assignment, default storage parameters, resolution of many-to-many relationships that may be in the logical model, and naming conventions. A DBMS-specific validation check is also provided in this wizard.

TIP: The **Quick Launch** can store common settings so that an operation can be reused on this model or on any other models. You can reuse the settings on another model by choosing the **Use File-Based Quick Launch Setting** option when saving the Quick Launch information on the last page of the wizard.

- To generate the new Physical Model, on the last page of the wizard, click **Finish**.

Now that a physical model has been generated from the logical model, feel free to navigate to specific objects via the Data Model Explorer, such as the CUSTMR table selected here. Double-click and view the physical details of the object such as DDL, Indexes, Partitions, and Storage.

The screenshot shows the ER/Studio Data Architect interface. The main window displays the Data Model Explorer for the 'Orders.dm1 - Physical Model View: DB2 Physical Model'. The tree view shows the following structure:

- Embarcadero Sales Order Processing
 - Logical
 - DB2 Physical Model
 - Main Model
 - Tables (40)
 - Addr_Cmpnnt
 - Addr_Dtl
 - Addr_Role
 - Addr_Typ
 - Addr_Typ_Cmpnnt
 - App_Usr
 - Authrzd_Prdct_Disct
 - Bill_Mtrls
 - Cmmsn_Credt
 - Custmr** (selected)
 - Custmr_Addr
 - Custmr_Phn

The Table Editor on the right shows the columns for the selected 'Custmr' table:

Column ID	Column Name
1	Custmr_ID
2	Frst_Nme
3	Lst_Nme
4	Cmpny_Nme
5	Eml_Addr
6	Onyx_ID
7	Is_Exstng_Custmr
8	Crted_By


Annotations on the left side of the image:

- New Physical Model in the Data Model Explorer**: Points to the 'Main Model' folder in the tree view.
- Double-click CUSTMR, to invoke the Table Editor and view design details, such as the DDL script.**: Points to the 'Custmr' table in the tree view.

DENORMALIZING THE PHYSICAL MODEL

ER/Studio DA comes equipped with denormalization wizards to help you optimize the physical design once the physical model is generated. The wizards help automate the process and keep the ties between the physical tables and the logical entities.


The active, denormalization wizards available depend on which tables are selected in the physical model when you select Denormalization Mapping. For example, if two tables that are related to each other are selected, the valid operations would be Rollups or Rolldowns.



When only one table is selected, the option to split the tables becomes available. The Table Merge option is available when two unrelated tables are selected.

Let's walk through an example of a denormalization operation using the generated physical model in a previous session of this tutorial. We may want to reduce the overhead on the `Custmr` table by splitting it into two physical tables, `Custmr_East` and `Custmr_West`. Splitting the table can reduce the query time and provide opportunities to store the tables at different locations which could further reduce lookup time.

Before the operation, the `Custmr` table should look like:



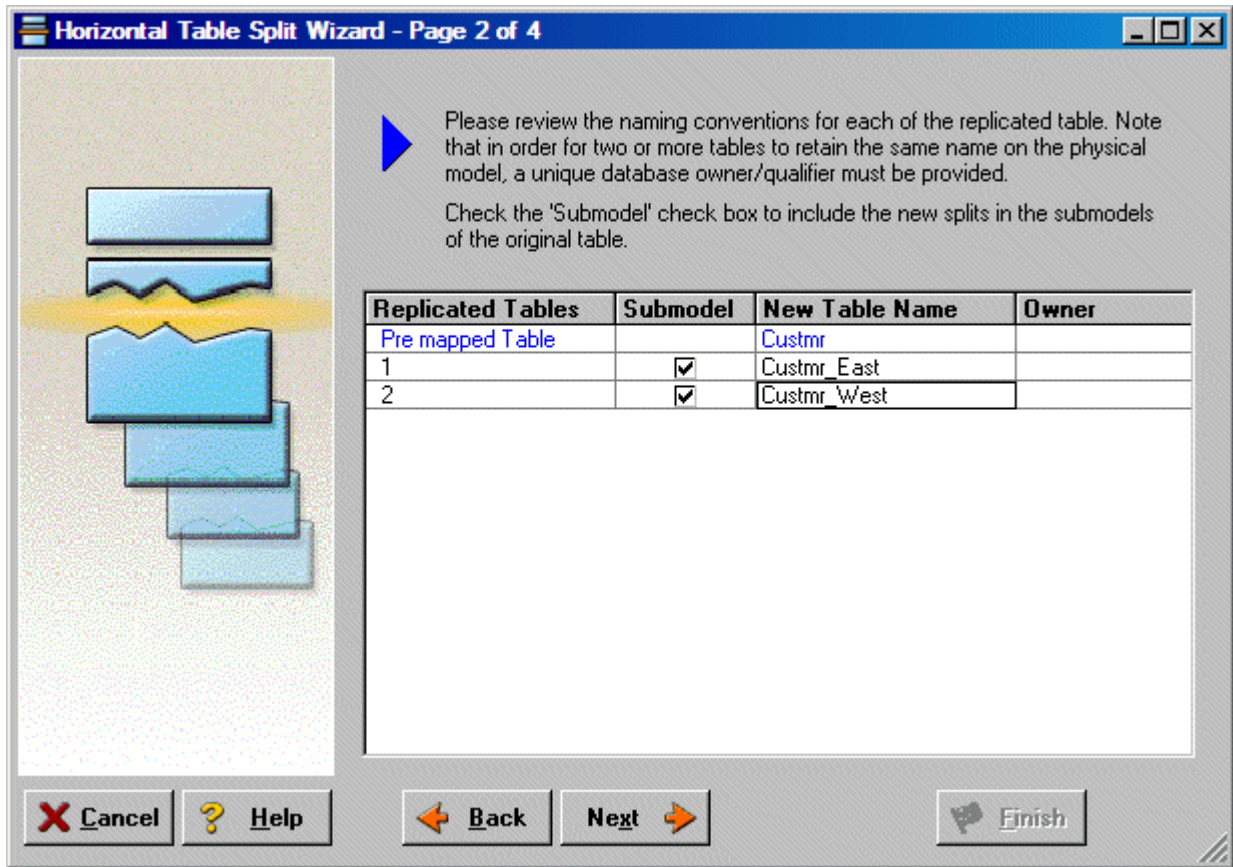
- 1 Open the `orders1.dml` model you modified and saved in the last session.
- 2 In the Data Model Explorer, right-click the **Custmr** table in the Physical Model.
- 3 Choose **Denormalization Mapping > Horizontal Splits**.

Notice that since only `Custmr` is selected, the only possible mappings are vertical and horizontal splits.

The **Horizontal Table Split Wizard** launches.

- 4 On Page 1, type 2 for the number of splits.

- 5 On Page 2, rename splits1 and 2 to `Custmr_East` and `Custmr_West` respectively.

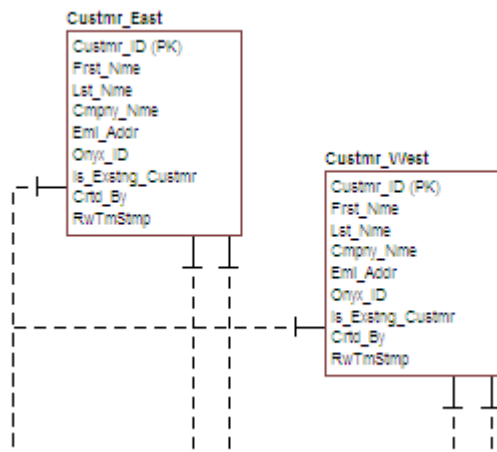


- 6 On Page 3, click **Next**.

We will keep all the relationships.

- 7 On Page 4, type a name and definition for the denormalization operation, and then click **Finish**.

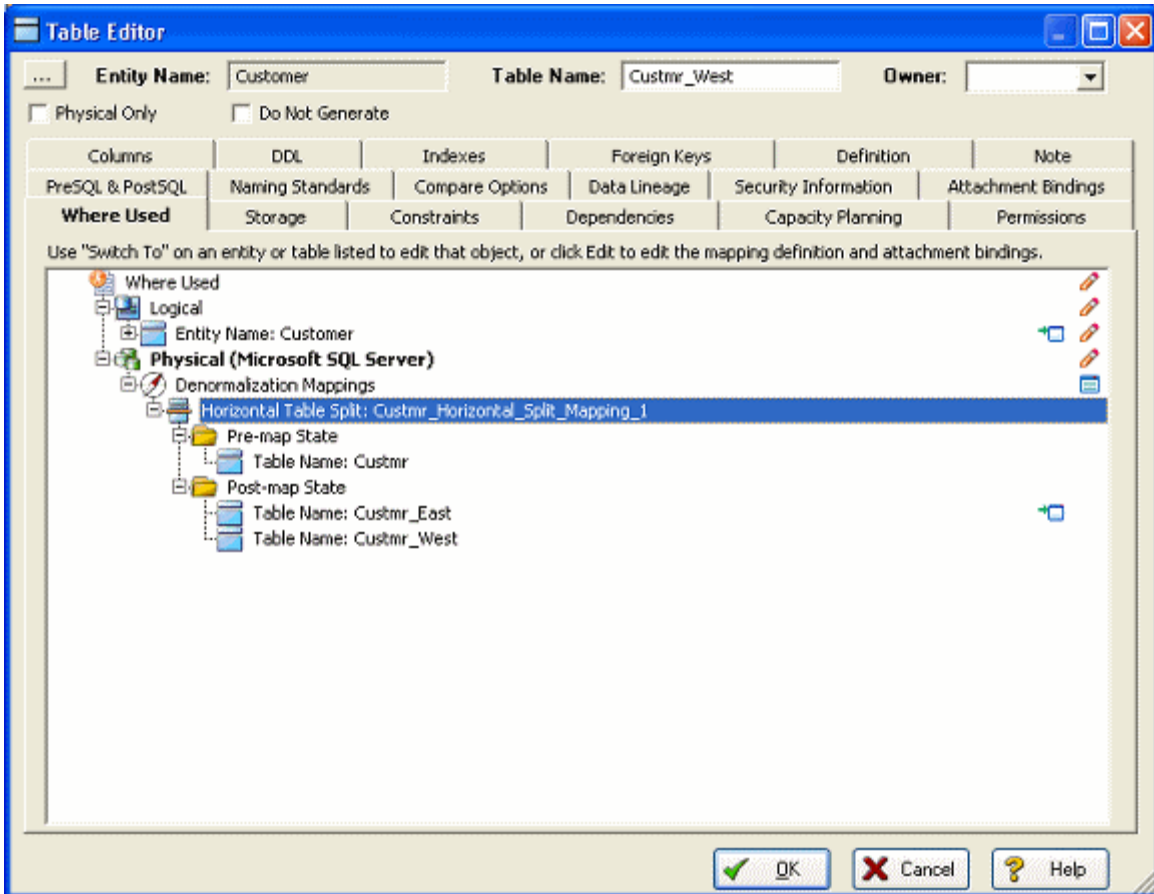
Finished! After the split the `Custmr` table will be two physical tables that look like this:



The two tables are identical except for the name.

You can selectively choose which attributes are included in the resultant tables by using a vertical split.

The denormalization mapping is stored with each of the tables.



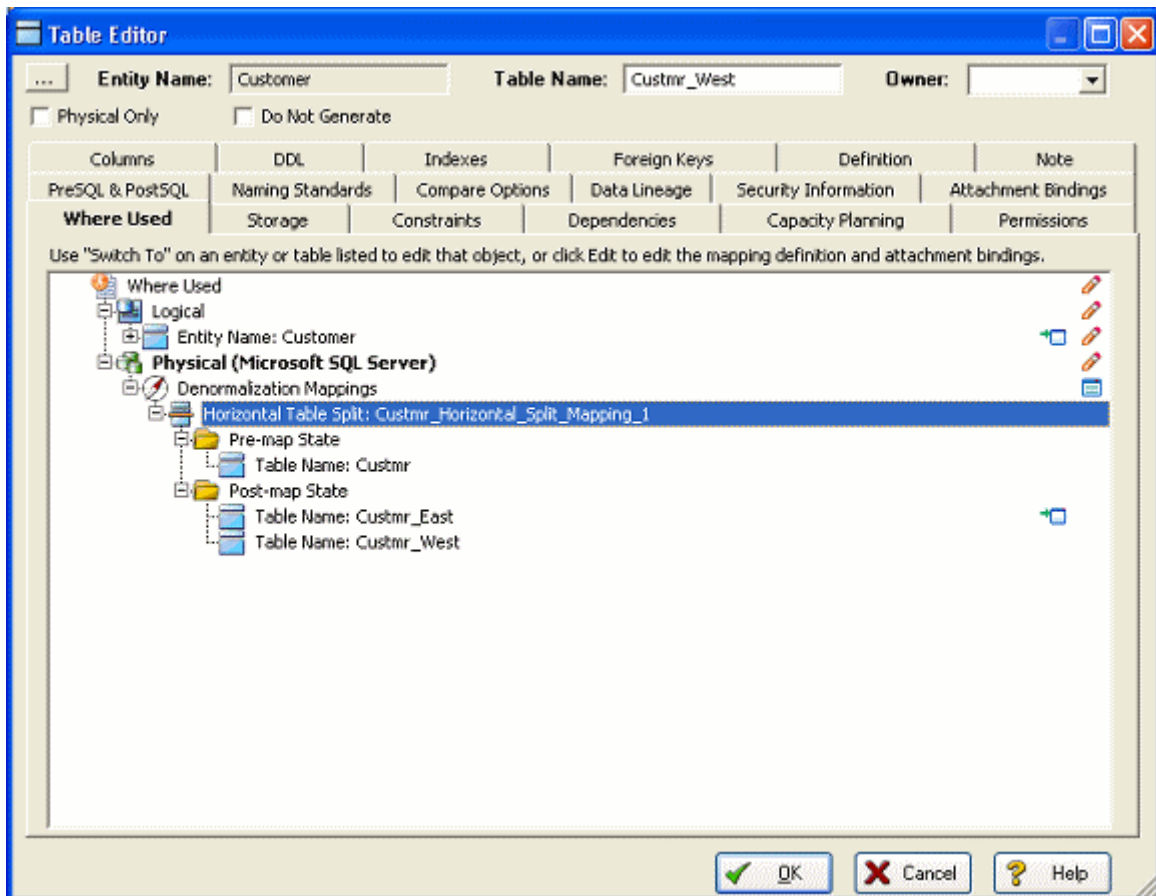
You can use the denormalization information to undo the operation or see the history of what happened. ER/Studio DA tracks the before and after states of these operations. This comes in handy in the next section where we discuss the **Where Used** analysis that can be performed between the logical and physical models.

FINDING OUT HOW AN ENTITY MAPS TO THE PHYSICAL MODEL

Now that we have performed a denormalization operation, the logical entity, Customer, essentially has become two physical tables, Custmr_East and Custmr_West. The ties between the logical and physical models are not lost. ER/Studio DA allows you to see what Customer in the logical model maps to in the DB2 physical model.

Let's take a look at the Customer entity in the logical model.

- 1 In the Data Model Explorer, navigate back to the **Customer** entity in the Logical model.
- 2 To start the Entity Editor, double-click the **Customer** entity.
- 3 Click the **Where Used** tab.



Once the tree is expanded, you can see the lineage of what has happened to the object. Notice that `Custmr_East` and `Custmr_West` are listed as physical implementations of the `Customer` entity. The denormalization mapping object shows how the end result was achieved.

The **Where Used** tab also displays the submodel usage of a particular entity within the logical or physical model. This allows you to see which business areas the entity belongs to.

NOTE: **Where Used** information is also available for attributes and columns.

Conclusion

In this session, you have seen how incredibly quick and easy it is to:

- Build a logical data model from scratch.

- Create a new submodel view to understand how to model on specific parts of a larger Main Model.
- Generate a physical model from a logical database in preparation for building a new database.
- Denormalize objects in the physical model.
- View the mappings between the logical and physical models using the Where Used tab.

For more information, refer to the ER/Studio DA Help and the review the topics “Developing the Logical Model” and Developing the Physical Model,” “Synchronizing Physical and Logical Models,” and “Generating a Script File or Database.”

DOCUMENTING AN EXISTING DATABASE

One of ER/Studio DA’s most powerful applications is that of a documentation generator to communicate complex databases and associated metadata to the Enterprise. ER/Studio DA is equipped with extensive report generation capabilities:


- **HTML Report Generation:** Instantaneous generation of an HTML-based Web site designed to provide simple navigability through data models and model metadata using standard browsers such as Microsoft’s Internet Explorer or Mozilla Firefox.
- **RTF Report Generation:** Instantaneous documentation generation compatible with applications like Microsoft Word.

In the exercise below, we will reverse-engineer an existing database and generate an HTML report for distribution and navigation to those who depend upon the information about the data model, but who may not be permitted to connect to the database for security or organizational reasons.

GENERATING AN HTML INTRANET DICTIONARY REPORT

PRE-REQUISITE: This exercise assumes that you can connect to an existing database in order to document it. Please refer to “Reverse-Engineering an Existing Database” in ER/Studio DA Help for explicit setup details if needed. If you cannot connect to an existing database, you can still generate documentation from the installed sample models; skip steps 1 through step 7 below which relate to in reverse-engineering and begin at step 8 after opening a sample model included with ER/Studio DA.

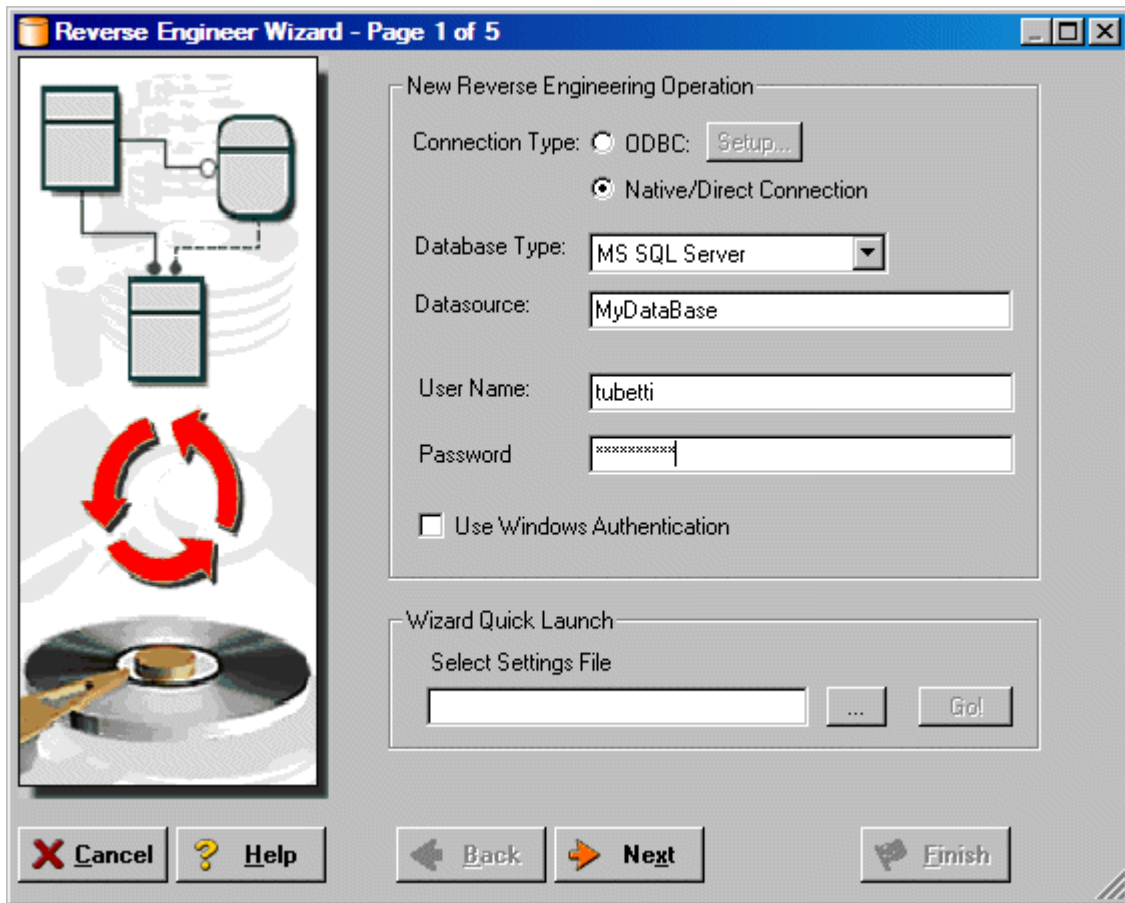
- 1 Choose **File > New**.



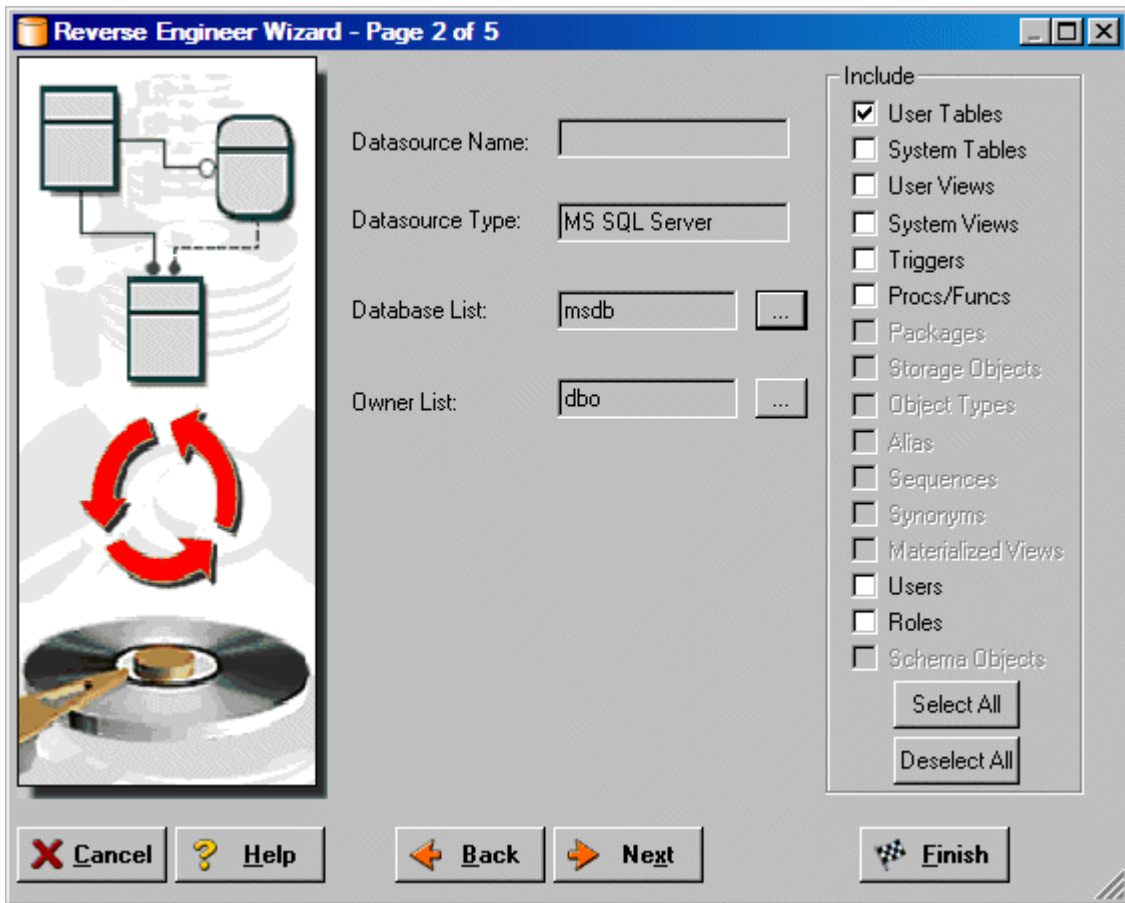
- 2 Select **Reverse-engineer an existing database**.
- 3 Click **Login**.

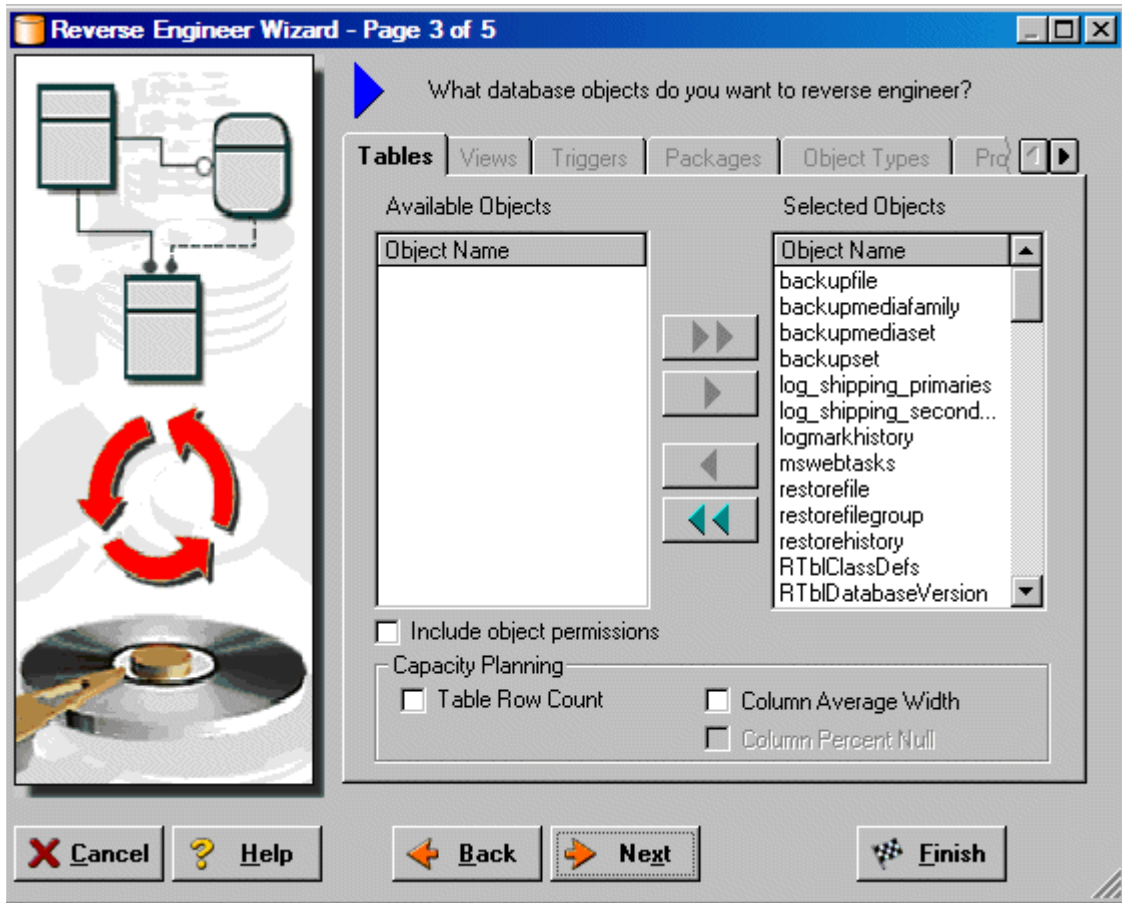
You can reverse engineer the database from either an ODBC datasource or via Native RDBMS client connectivity. In this example, Native Connectivity to Microsoft SQL Server will be demonstrated.

- 4 Type the relevant connectivity information such as the data source name, and user name and password and then click **Next**.





- 5 Walk through the **Reverse Engineer Wizard** selecting the objects, options, and layout preferences for the model.





- 6 Continue through the wizard to select layout styles and other preferences.



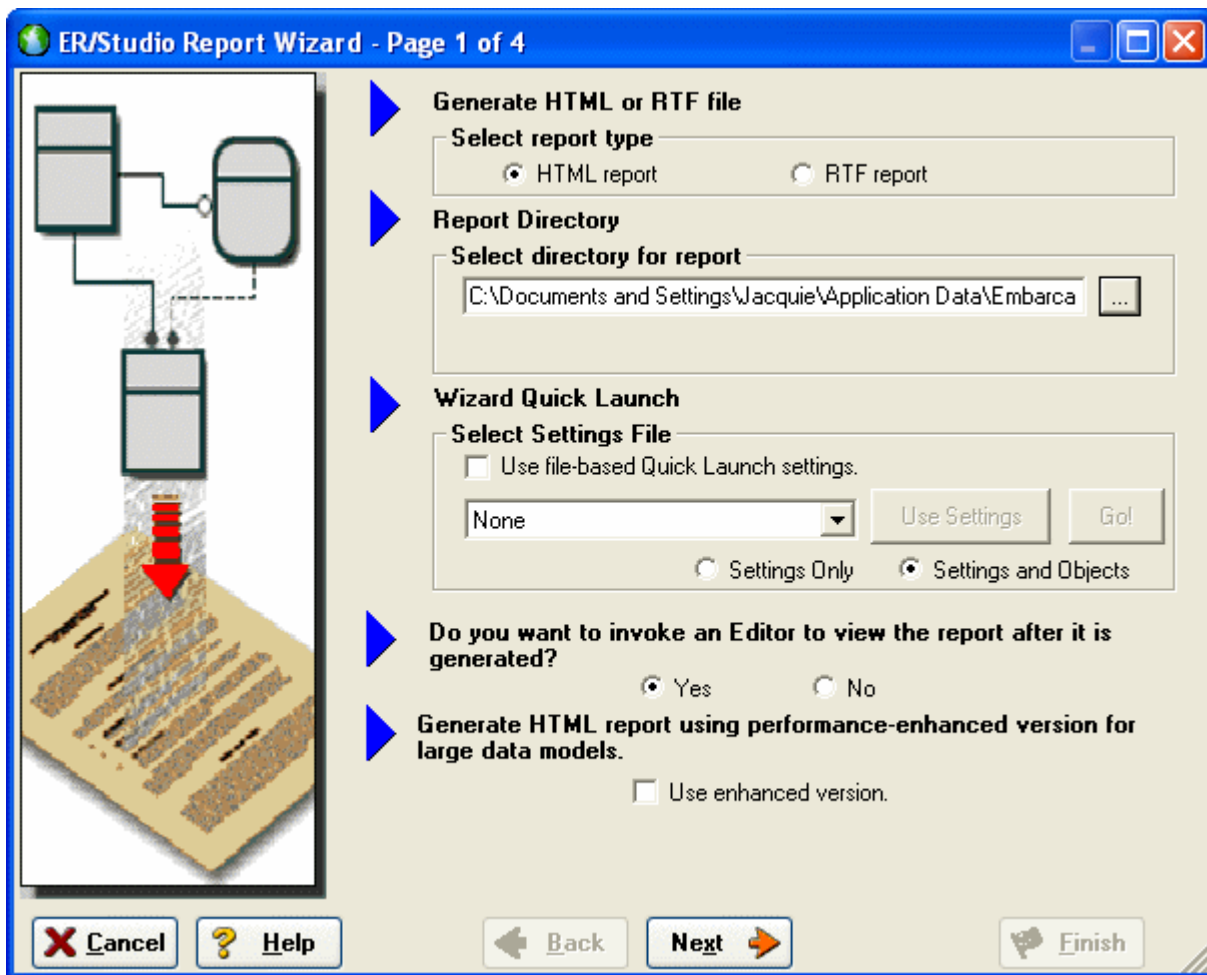


- 7 Click **Finish** and ER/Studio DA reverse engineers your database!

Once reverse engineering of your database is complete, we will generate a complete HTML report of the database for others in your organization to review.

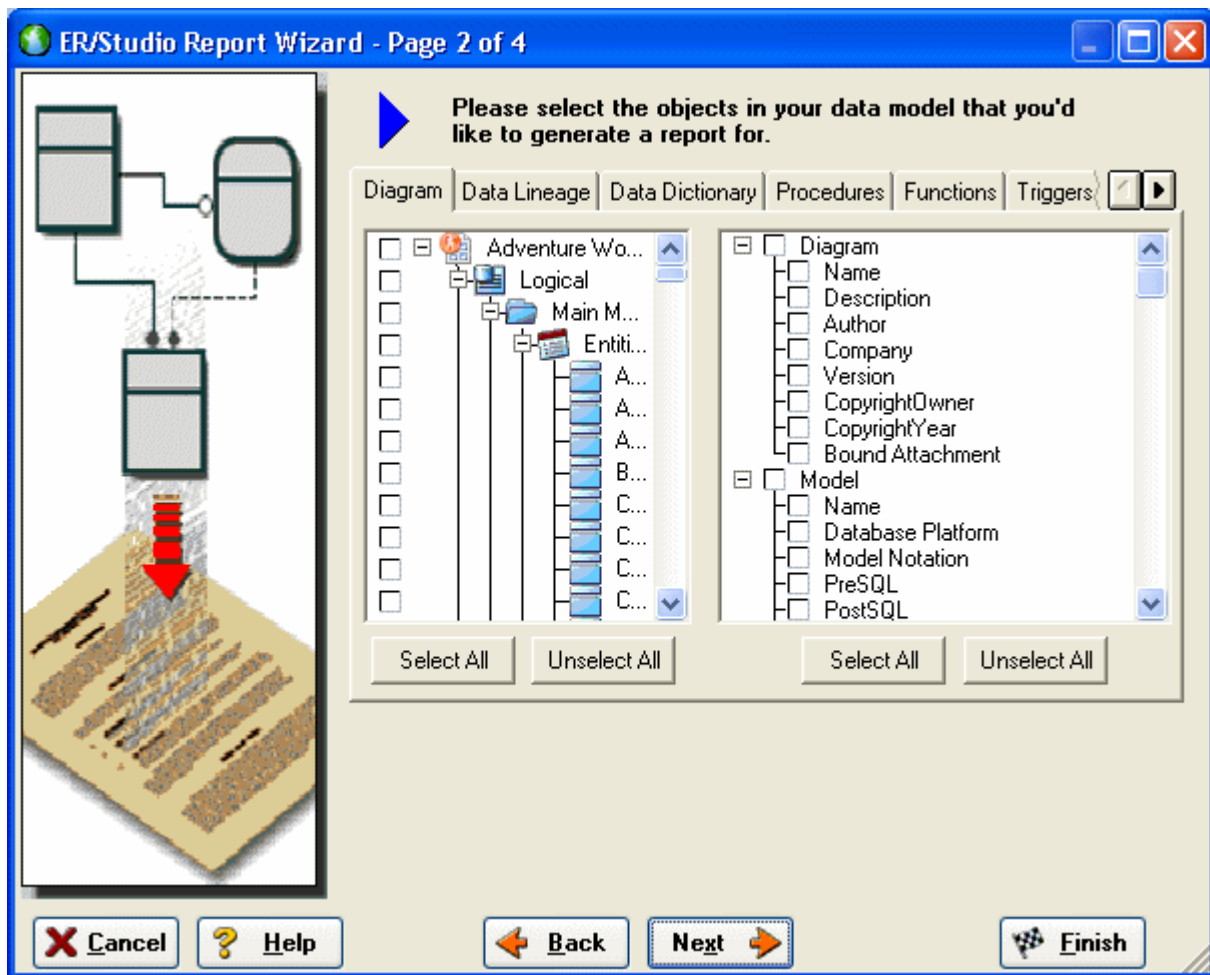
- 8 In the Data Model Explorer, select the **Physical Main Model**.

9 Choose **Tools > Generate Reports**.



10 On the first page of the wizard, for the report type, select **HTML**.

- 11 On page 2, click **Select All** in both areas of the **Diagram** tab.

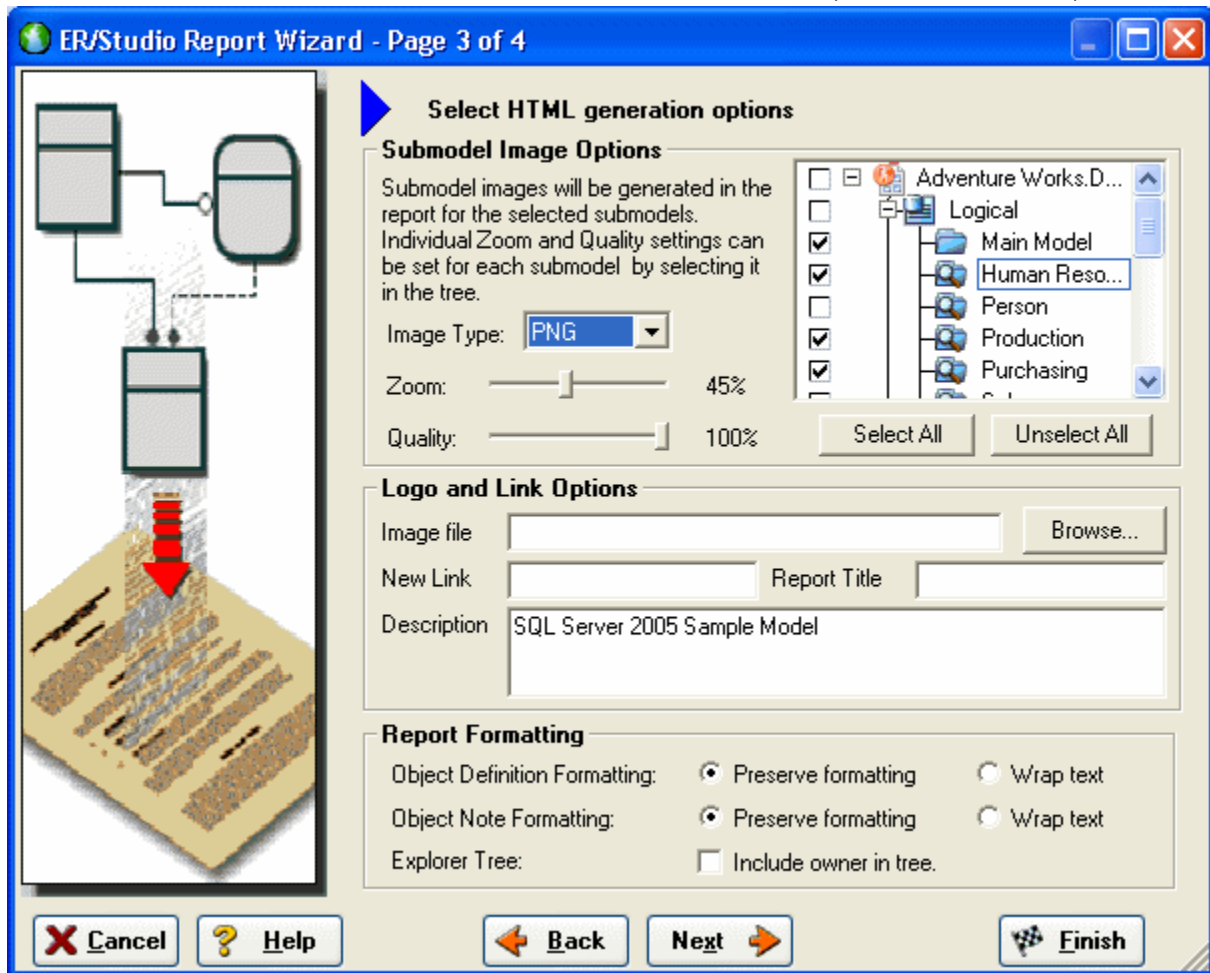


- 12 Click **Select All** in both areas of the **Data Dictionary** and the **Procedures** tabs, and then click **Next**.

NOTE: The tabs available on page 2 depend on what objects are supported by the database platform of the selected model. Some databases support more database objects than Microsoft SQL Server 2005, such as Oracle11g and IBM DB2 LUW 9.x for which there are also tabs on this page for procedures, functions, triggers, packages, and tablespaces. If the model was previously denormalized, a **Denormalization Mappings** tab would also appear.

- 13 On page 3, in the **Submodel Image Options** area, click **Select All**.

TIP: In the **Logo and Link Options**, you can choose to replace the ER/Studio DA default Embarcadero Technologies logo in favor of your own corporate logo (and Hyperlink).



Because HTML formatting can be included in object definitions, you can also choose to preserve the formatting specified on the **Definitions** tab of the various object editors.

- 14 Click **Next** to advance to Page 4 of 4, and then click **Finish**.

ER/Studio DA then begins the report publication process and launches the default browser so you can review the report.

15 Finished!

Start navigating the report via your browser. Navigation will perform exactly as it does when you are using ER/Studio DA! Expand the tree to find Model Image and click on it (see below). You will see a read-only version of your data model (as seen below). Use the Explorer to navigate to any metadata you want or select the entities and relationships in the model image to jump to their information.



Conclusion

In this session, you learned how to:

- Connect to and reverse-engineer an existing database with ER/Studio DA.
- Document a database in seconds by using ER/Studio DA's automatic HTML documentation publication facility.

For more information on Reporting, review the "Generating RTF and HTML Model Reports" in the *ER/Studio Data Architect User Guide*.

DOCUMENTING DATA LINEAGE

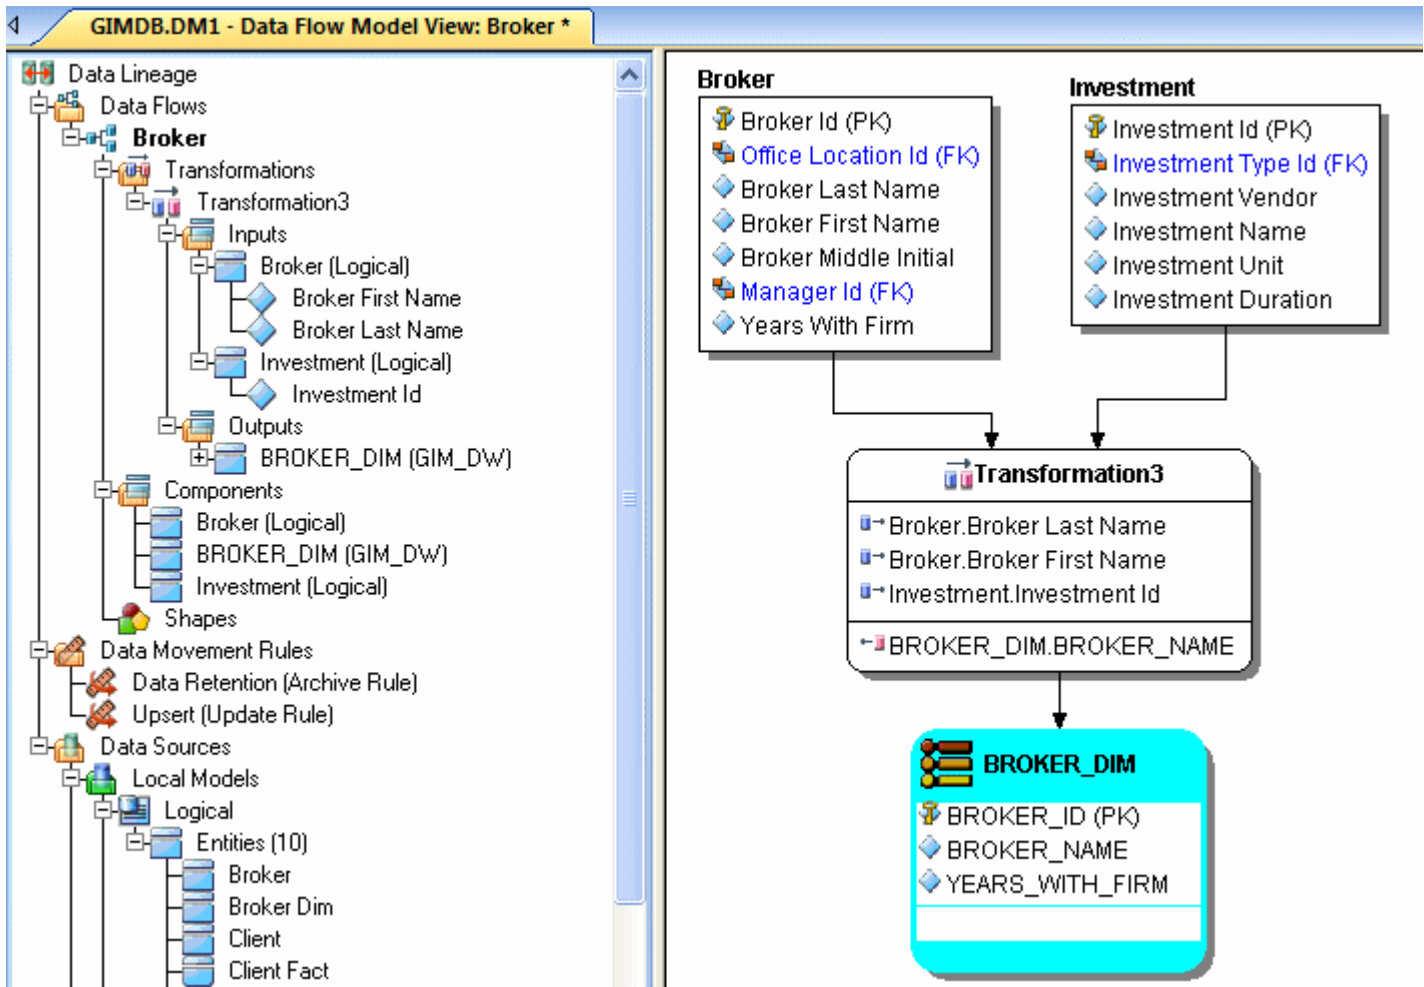
The Data Lineage feature of ER/Studio DA enables you to document the movement of data from point A to point B (and any intermediate steps in between). This movement is sometimes referred to as Extraction, Transformation and Load (ETL). Points A and B can be anything from flat files, high-end databases such as Oracle and DB2, XML, Access databases, and Excel worksheets. This is sometimes referred to as source and target mapping. A model produced in ERStudio can represent any point along the way. Data Architects need the ability to specify the *source* or *target* of data down to the column-level. Along with the metadata that defines the source and target mapping are rules for how the data is manipulated along the way.

This section will help you document the data lineage of your systems. It is comprised of the following tasks which correspond to the general ETL workflow:

- [Creating a Data Flow](#) on page 58
- [Creating a Data Movement Rule](#) on page 59
- [Defining External Source and Target Systems](#) on page 61
- [Creating a Data Lineage and Transformation Visualization](#) on page 63

CREATING A DATA FLOW


The Data Flow organizes and encapsulates one data transformation and the source tables and columns used in the transformation to produce the target data. Multi-tiered mappings are possible and there can be multiple transformations involving different columns between two tables as illustrated below.



Create a Data Lineage Data Flow

- 1 Choose **File > Open** and select the GIMB.DM1 diagram in the Sample Models directory.
- 2 Click the **Data Lineage** tab at the bottom of the application window.


You are prompted to create a **Data Lineage Data Flow**.




- 3 Click **Yes**.

*If this is not the first time you click the Data Lineage tab after opening a diagram, from the **Data Lineage Explorer**, right-click the **Data Flows** node and then click **Create Data Flow**.*

- 4 Enter a data lineage data flow name and then click **OK**.



NOTE: The name that appears in the diagram title tab at the top of the application window is appended with: data flow name, when you click a task in the **Data Lineage** explorer, such as GIMDB.DM1 - Data Flow Model View: Broker*.




The Data Flow has been created.

CREATING A DATA MOVEMENT RULE

Data Movement rules describe how source and target tables and entities are related. You can relate source data to one or more tables and entities in the same model, the active diagram, or to tables imported from external systems. The rules defined here are used at the table level on the Data Lineage tab of the entity and table editors.

Create a data movement rule

- 1 On the **Data Lineage** tab, right-click **Data Movement Rules** and choose **New Data Movement Rule**.



- 2 Complete the **Data Movement Rule** editor as required and then click **OK** to exit the editor.

TIP: Once created, you can edit the Data Movement rule by double-clicking it to launch the **Data Movement Rule** editor.

The following describes options that require additional explanation:

Rule Information tab

- **Rule Name:** Enter a name that indicates the operation and objects acted on, depending on the specifics of your binding definition.
- **Rule Type:** Select a generic movement rule type that best describes the data movement.
- **Rule Text:** Document your data movement plan here, perhaps adding instructions or contingency plans.

Binding Information tab

Select the object classes and/or specific objects to which you want to bind this attachment. You can override this setting using the **Data Lineage** tab of the entity or table editor.


DEFINING EXTERNAL SOURCE AND TARGET SYSTEMS

Data sources can originate from models in the active diagram (local models) or from external sources that are either imported into the active diagram or created on the Data Lineage tab. A data source can be imported from *.dm1 files, *.at1 files, database or from SQL files, flat files, and other common application files. The following describes how to import metadata from an external source.

NOTE: Source data imported through the Data Lineage tab only includes information such as table and column name, datatype, nullability, primary key, and column definitions. To obtain more details, reverse engineer the database or import it into ER/Studio DA using the Metadata Wizard.

Import external source or target data

- 1 From the **Data Lineage** tab, expand the **Data Sources** node.
- 2 Right-click **Other Sources** and choose **Import New Source**.



- 3 Complete the **Import Source** wizard as required and then click **Finish** to import the source.

The new source will appear under the **Other Sources** node.

The following describe options that require additional explanation:

Page 1 - Please select where you would like to import the source metadata from

- **From a Repository based DM1 file:** Lets you obtain source from data models and Named Releases managed within the ER/Studio Repository. When you select this option, ER/Studio DA opens the Repository Operation Status dialog box and the **Get From Repository** dialog box. This process connects to the current Repository Server defined in the Repository settings. The **Import Source** wizard automatically gets the diagram.
- **From an SQL file** ER/Studio DA imports the SQL file.
- **From a live database:** If you select this option, a page appears where you can select the database and connection type. The connection type can be either ODBC or Native/Direct Connection. For information about connecting to databases, including troubleshooting information, see "Connecting to Database Sources and Targets" in the *ER/Studio Data Architect User Guide*.
- **Comparison Quick Launch:** The Compare Quick Launch data is saved as an *.rvo file. For information on using the **Quick Launch** option in the wizard, see "Saving and Using Quick Launch Settings" in the *ER/Studio Data Architect User Guide*.


Page 5 - Results

- **Current and Target Model Display Grid:** Between the Source and Target models is a Resolution column. The default merge decision is Merge the data into the new source file. You can click on any item in the **Resolution** column to enable the decision list. If you want to change the decision, click the list and then click the new resolution. When you change the default resolution of an object, the decisions of their dependent properties and objects are automatically updated. You can also click the category folders, like the **Tables Resolution** column to change all the decisions for all the underlying objects in that object category. And, you can use the **CTRL** key to select multiple items, and then right click to enable the decision list.
- **SQL Difference:** To enable the SQL Difference utility, select any difference that is a long text field, such as a Definition, Note, or DDL, and then click SQL Difference to view the differences between the SQL of the models. This utility only allows you to view the differences; difference resolutions are performed on the **Results** page of the **Compare and Merge Utility**.
- **Filter Report on Results:** Create a report of the source content and your chosen resolutions. You can choose to create an HTML or an RTF report.

TIP: You can modify the default display using the options at the bottom of the page.

CREATING A DATA LINEAGE AND TRANSFORMATION VISUALIZATION


- 1 To create the data source or transformation input tables, expand **Data Sources > Local Models > Logical > Entities** and then drag and drop the **Broker** and **Investment** tables onto the Data Lineage window.
- 2 To create the data target or transformation output tables, navigate to **Data Sources > Local Models > GIM_DW** and then drag and drop the **Broker** table onto the Data Lineage window.



- 3 To obtain the Transformation Insertion tool, right-click an empty space in the **Data Lineage** window and then click **Insert Transformation**.
- 4 To insert the transformation, click in the **Data Lineage** window between the source and target data sources and then right-click to drop the **Transformation Insertion** tool.
- 5 Reposition and resize the transformation object to suit your needs.
- 6 Right-click an empty space of the **Data Lineage** window and then click **Insert Data Stream**.

TIP: Transformation and Data Flow tools are also available on the toolbar. Hover the mouse over the tools to find the tool you need.

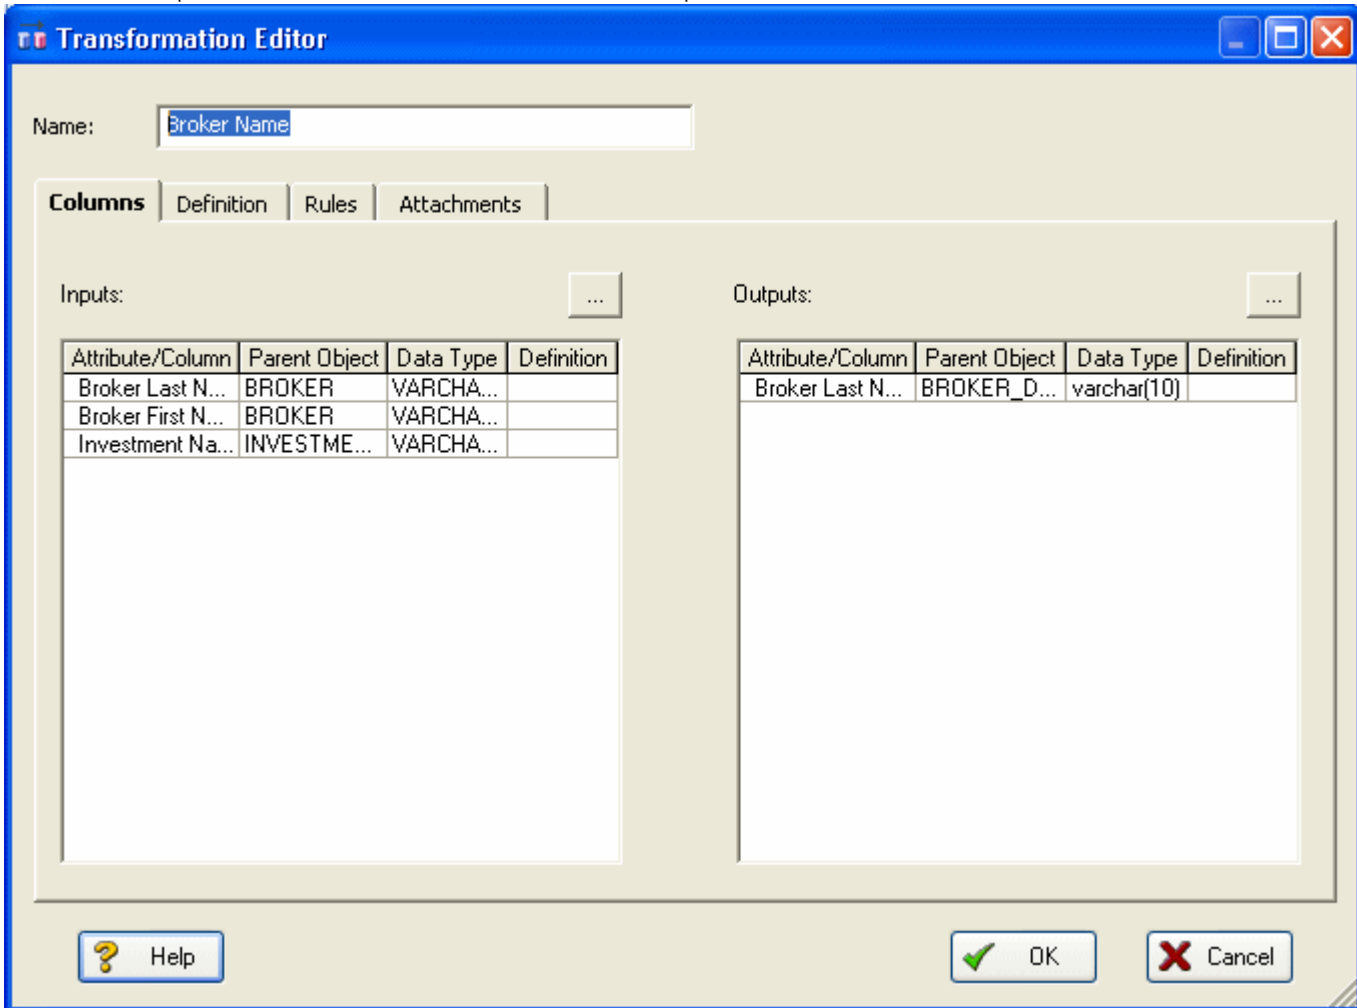
- 7 Click an input and then click the transformation object. Repeat as many times as necessary to link all the inputs to the transformation object.
- 8 Click the transformation object and then click an output.



TIP: If the Inputs and Outputs do not display in the diagram as they do in the illustration above, choose **View > Diagram and Object Display Options > Transformation** and then click **Inputs and Outputs Columns**.

- 9 To define which columns should be used in the transformation and any transformation rules, double-click the new transformation to open the **Transformation Editor**.

10 Complete the **Transformation Editor** as required and then click **OK** to exit the editor.



You are done! Now you can more easily share your ideas with your colleagues!

TIP: Once the Data Flow is created, you can double-click it to change its name, or double-click a transformation or component to change its properties.

The following describes options in the Transformation Editor that require additional explanation:

Columns tab

- **Inputs:** Click the ellipsis (...) button to choose the inputs to be transformed in this task.
- **Outputs:** Click the ellipsis (...) button to choose the outputs resulting from the transformation.

Definition tab

- **Business:** Describe the transformation for your audience.
- **Code:** Enter the code that will perform the transformation, such as a SELECT statement, or a VBasic or Java Script function or procedure.

Data Movement Rules tab

These are the rules from the Data Movement Rules node of the Data Lineage Explorer.

NOTE: You can delete or edit an input or output column by double-clicking the transformation in the **Data Lineage** window, clicking the ellipsis in the **Transformation Editor** and then deselecting the column you want to remove.

Attachments tab

Bind an external piece of information or attachment to the transformation. You can also remove an attachment from an object, override the default value of an attachment, or change the position of a bound attachment. To override the value of the attachment you have moved to the **Selected Attachments** grid, double-click the **Value** field of the target attachment. ER/Studio DA opens the **Value Override Editor** or a list, depending on the attachment datatype. Attachments are created in the **Attachments** folder of the **Data Dictionary** and must be applied to the default before they will display on this tab.


DIAGRAM NAVIGATION AND AESTHETICS

To assist with the creation of presentation-quality diagrams that are easy to navigate and are aesthetically pleasing, ER/Studio DA offers progressive diagram Auto Layout and Navigation utilities that also help you to clean up complex diagrams. Modelers should spend time solving complex database or business data model problems, not forcing boxes and lines to look a certain way.

NAVIGATING THE DIAGRAM

To demonstrate some of ER/Studio DA's layout and navigation utilities, we will import a sample SQL script provided with ER/Studio DA.

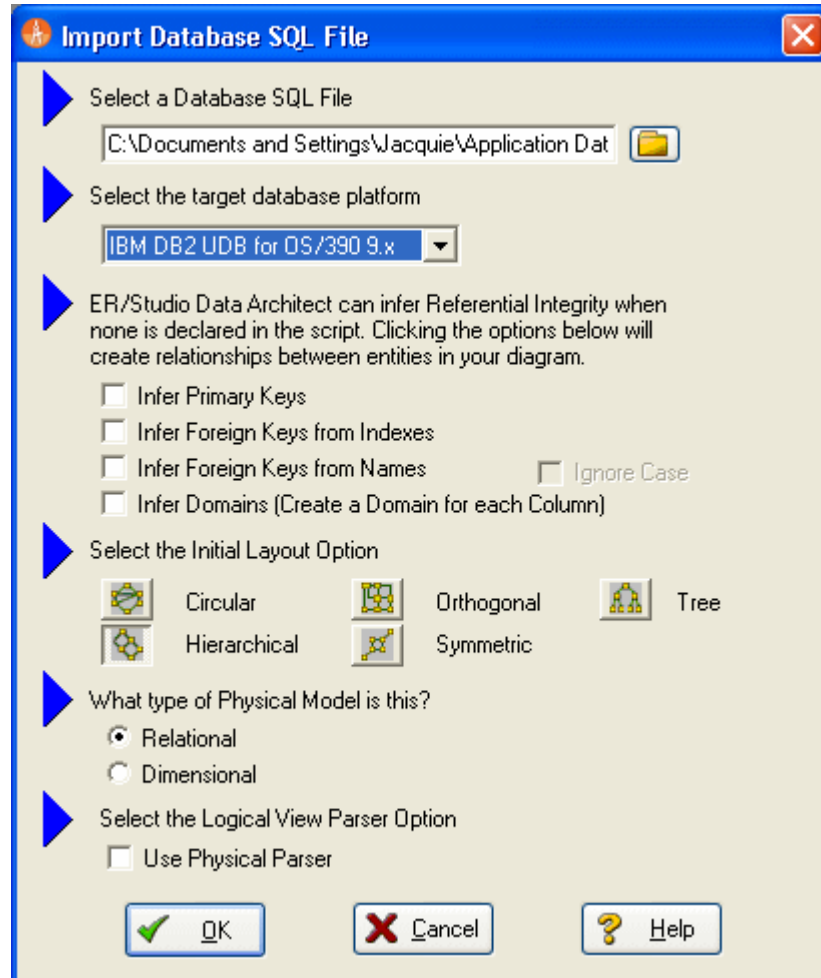
- 1 Close any files you have open.
- 2 Choose **File > New**.



- 3 Select **Import Model From:** and then in the import list, click **SQL File**.

NOTE: The ERX File choice enables you to import Computer Associates ERwin 3.5.2 ERX files. Selecting **From External Metadata** launches the MetaWizard to import from alternative sources.

The **Import Database SQL File** dialog appears:



- 4 To the right of **Select a Database SQL File** click the folder icon, click **IBM DB2 OS390.SQL**, and then click **Open**.

The full path to this file is:

For Windows XP:

C:\Documents and Settings*<user>*\Application
Data\Embarcadero\ERStudio\SQLCode

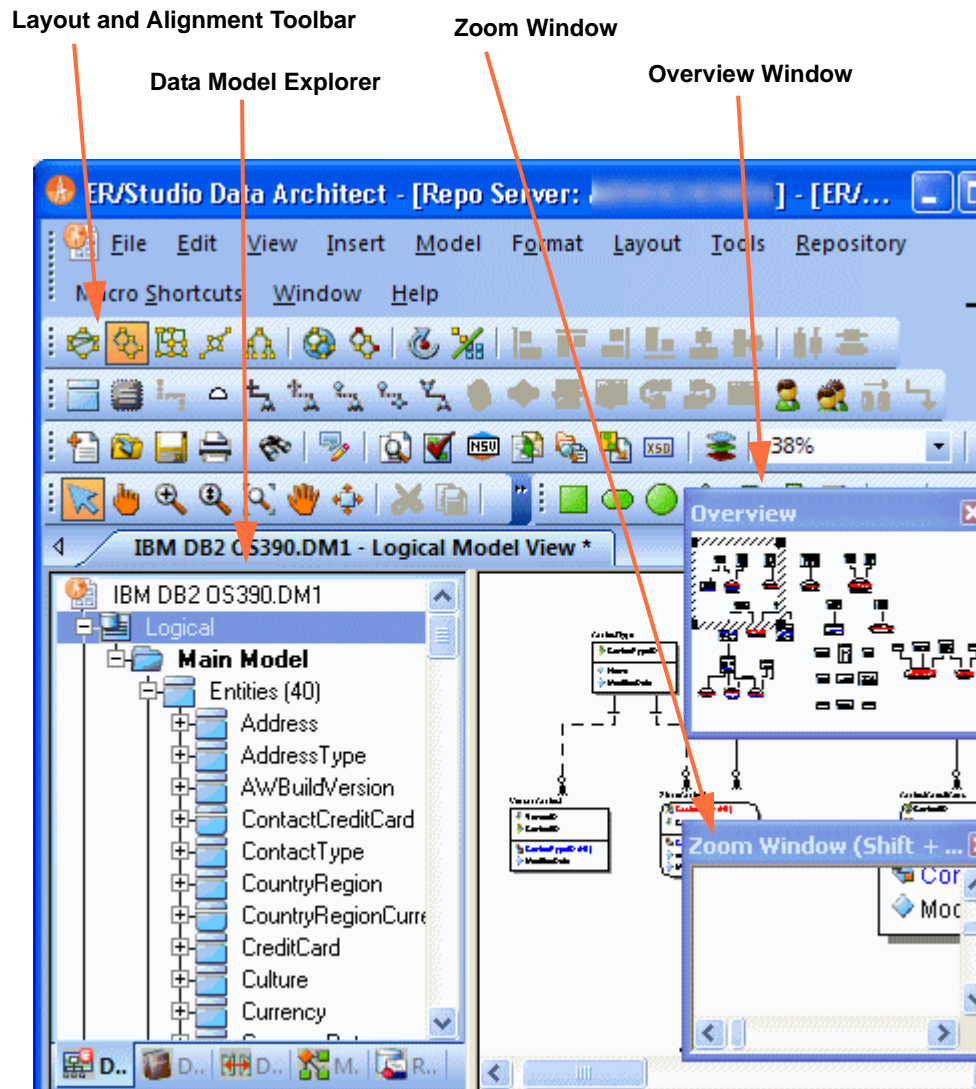
For Windows Vista:

C:\Users*<user>*\AppData\Roaming\Embarcadero\ERStudio\SQLCodeSample

- 5 In the **Select the target database platform** list, click **IBM DB2 UDB for OS /390 9.x**.

6 Click **OK**.

Finished! Once the SQL Script is finished importing (as depicted below) the following items will assist you in leveraging a variety of Auto Layout and Navigation Features.



- **Layout and Alignment Toolbar:** Use any of the four Auto Layout styles to change the layout of the diagram with the click of a button. The auto layout styles are all entirely customizable styles. You can also customize the diagram layout via the Layout Properties pages that can be launched by clicking **Layout > Layout Properties**.
- **Data Model Explorer:** Click on any object in the **Data Model** explorer and it will automatically be selected in the diagram and focused in both the **Zoom Window** and **Overview** window.
- **Overview Window:** Use this window as a thumbnail of your model to pan the entire model or zoom in and out. It can also pan and zoom the diagram if grabbed or sized. If the **Overview** window is not already visible, press **F9** to activate it.

- **Zoom Window:** Use this window as a magnifying glass to enlarge any diagram objects under your mouse cursor. You can also press **SHIFT+F8** to freeze the zoom window to keep a single object frozen while you continue to pan around the diagram. If the **Zoom** window is not already visible, press **F8** to activate it.

DIAGRAM AESTHETICS

One of the tremendous benefits of building data models is the wide range of audiences that can realize value from them. Part of this relies on what information is displayed on the diagram. Depending on the audience you may want to limit or expand what is displayed. For example, developers may benefit from looking at a model that displays data type, null option, and unique and non-unique index information, while business analysts may just need the entity name and the definition. ER/Studio DA offers many display properties that can be customized exactly for this purpose.


Continuing with the previous section, we will use the DB2 model that was built to demonstrate some of the ways to customize the appearance of the model.

We will use the Diagram and Object Display Options dialog on the Diagram toolbar to further customize the view of the logical and physical models.

NOTE: You can use the Colors & Fonts tool to customize the look and feel further of each model.

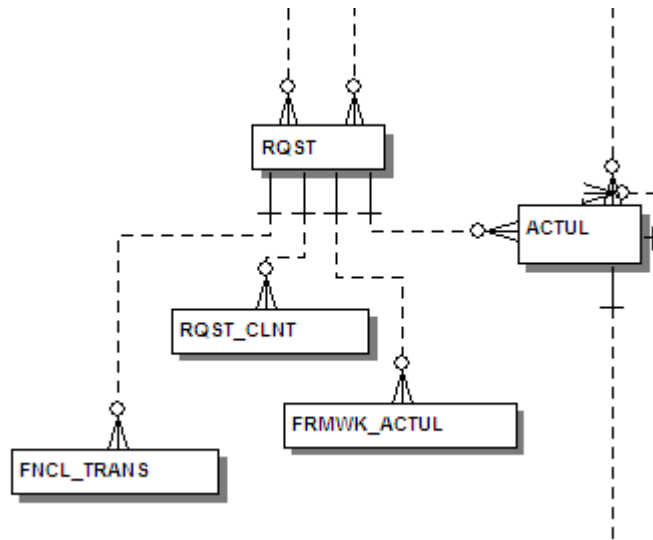
SETTING THE LOGICAL MODEL DISPLAY

- 1 Select the logical model and then on the **Diagram Toolbar**, click the **Diagram and Object Display Options** tool.



- 2 In the **Diagram and Object Display Options** dialog, click the **Entity** tab, and then in the **Display Level** area, select **Entity**.

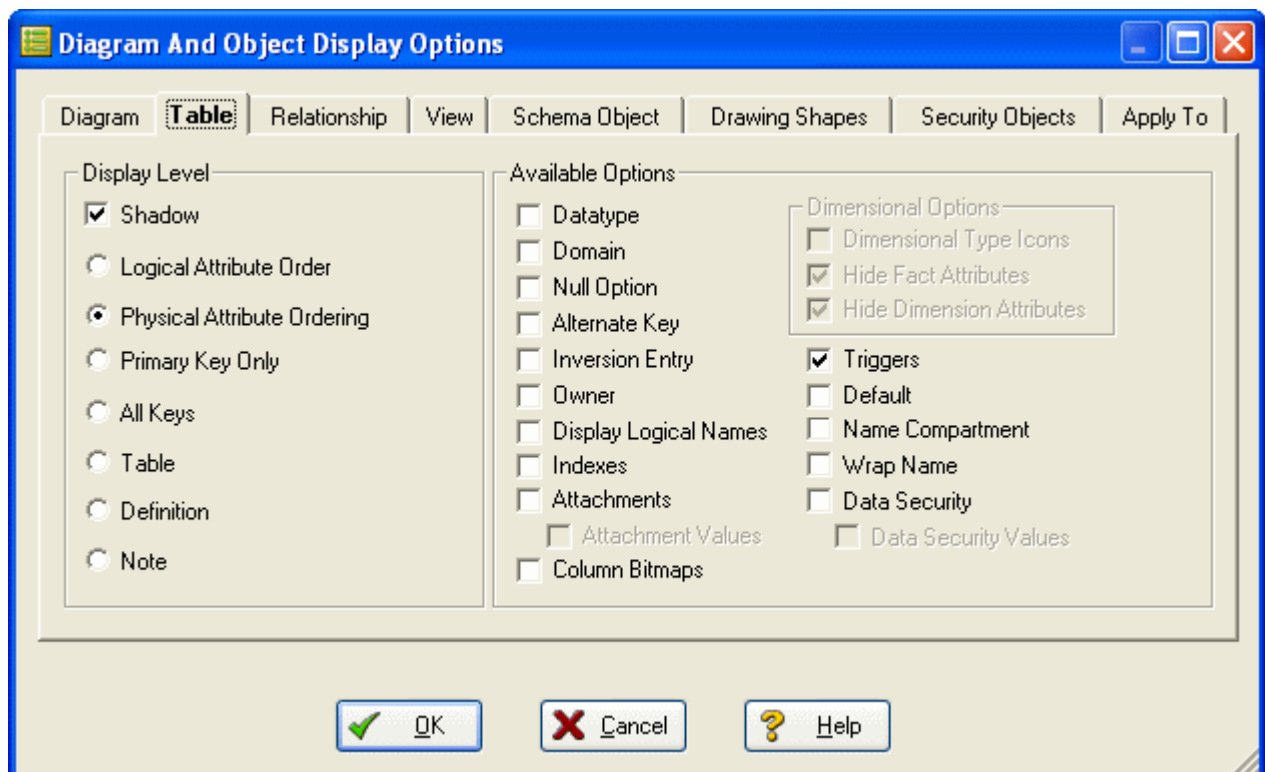
- 3 Click **OK**.



NOTE: Only entity names are displayed for each entity. You may also want to re-layout the diagram since the entity sizes have changed.

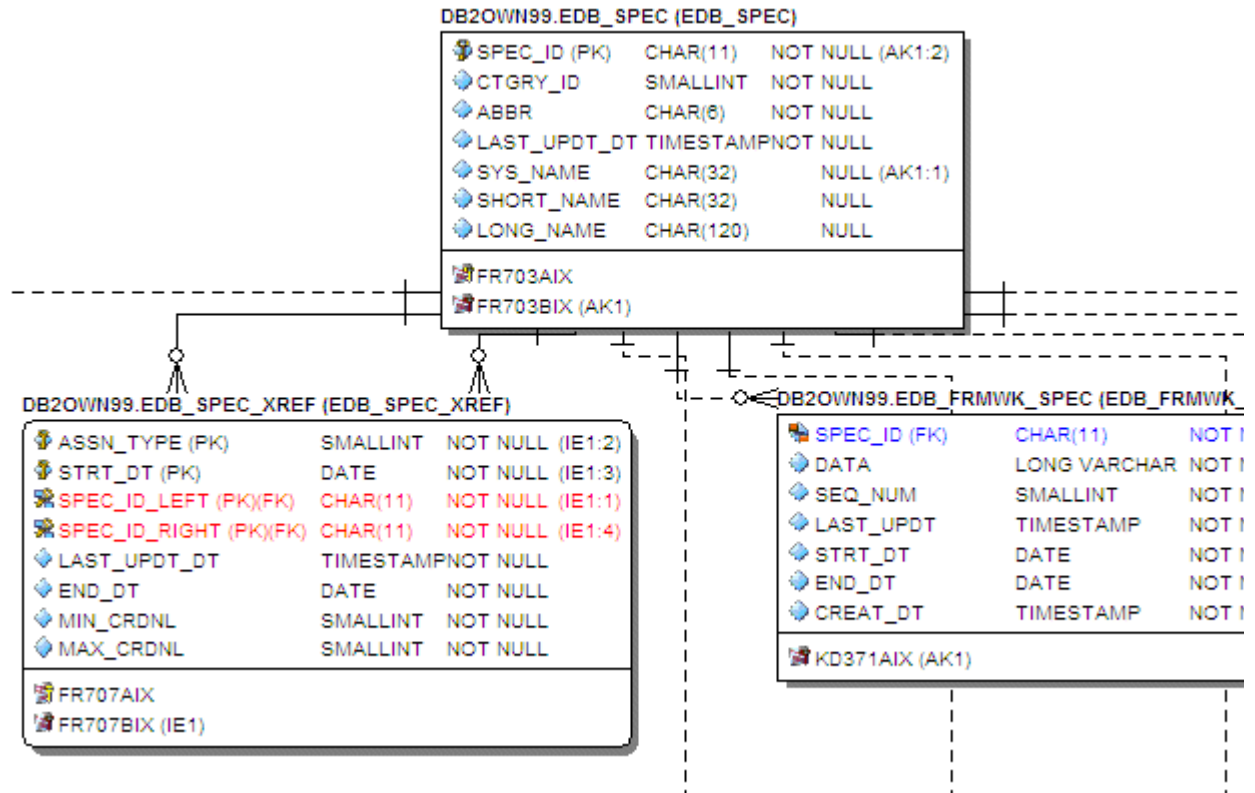
SETTING THE PHYSICAL MODEL DISPLAY

- 1 Select the physical model and then choose **View > Diagram and Object Display Options**.
- 2 In the **Diagram And Object Display Options** dialog, click the **Table** tab.



- 3 In the **Display Level** area, select **Physical Attribute Ordering**.
- 4 In the **Available Options** area, select the specific properties you want to display.
- 5 Click **OK**.

The model should now display more details for the physical model, as seen below.



NOTE: Since the sizes of the objects changed, you may want to change the model layout using one of the ER/Studio DA advanced layout engines. You can also customize the default display properties for new models by choosing **Tools > Options**, and then selecting the desired options on the **Display** tab.

Conclusion

In this session, you have learned how to:

- Import an SQL file and allow ER/Studio DA to automatically create a diagram.
- Use a variety of auto layout and navigation tools to enhance the aesthetic experience of the diagram and to improve the data model navigability.
- Customize the display of both the logical and physical models.