

SeeBeyond ICAN Suite

SeeBeyond ICAN Suite Upgrade Guide

Release 5.0.1



The information contained in this document is subject to change and is updated periodically to reflect changes to the applicable software. Although every effort has been made to ensure the accuracy of this document, SeeBeyond Technology Corporation (SeeBeyond) assumes no responsibility for any errors that may appear herein. The software described in this document is furnished under a License Agreement and may be used or copied only in accordance with the terms of such License Agreement. Printing, copying, or reproducing this document in any fashion is prohibited except in accordance with the License Agreement. The contents of this document are designated as being confidential and proprietary; are considered to be trade secrets of SeeBeyond; and may be used only in accordance with the License Agreement, as protected and enforceable by law. SeeBeyond assumes no responsibility for the use or reliability of its software on platforms that are not supported by SeeBeyond.

SeeBeyond, e*Gate, and e*Way are the registered trademarks of SeeBeyond Technology Corporation in the United States and select foreign countries; the SeeBeyond logo, e*Insight, and e*Xchange are trademarks of SeeBeyond Technology Corporation. The absence of a trademark from this list does not constitute a waiver of SeeBeyond Technology Corporation's intellectual property rights concerning that trademark. This document may contain references to other company, brand, and product names. These company, brand, and product names are used herein for identification purposes only and may be the trademarks of their respective owners.

© 2003 by SeeBeyond Technology Corporation. All Rights Reserved. This work is protected as an unpublished work under the copyright laws.

This work is confidential and proprietary information of SeeBeyond and must be maintained in strict confidence.

Version 20031121145328.

Contents

Chapter 1

Introduction	5
Organization of Information	5
Intended Audience	6
Writing Conventions	6
Additional Conventions	6
Online Documents	7
The SeeBeyond Web Site	7

Chapter 2

Upgrading e*Gate 4.x	8
eGate 5.0 Upgrade Overview	8
Schema Runtime Environment Upgrade	8
General Features of the SRE	10
SRE Supported Platforms	11
Benefits of Upgrading to eGate 5.0	11
Improved Monitoring	11
Schema and Project Integration	11
Upgrade Procedure	12
Pre-Installation Steps	12
Installation Steps	12
Post-Upgrade Deployment Strategies	13
Using a Common SRE JMS IQ Manager	14
SRE Configuration	14
J2EE Environment Configuration	14
Using a Common J2EE IQ Manager	15
SRE Configuration	15
J2EE Environment Configuration	16

Chapter 3

Upgrading eInsight	17
Overview	17

eInsight Upgrade Procedure	17
Integrate Existing Business Processes and Schemas	18
Connect Using the Schema Runtime Environment	18

Chapter 4

Upgrading eIndex	19
Upgrade Overview	19
System Requirements	19
Upgrade Tasks	20
Upgrading the Schema to a Project	20
Upgrading the Outbound Queue	20
Upgrading the GUI	20
Upgrading to the EDM	21
Upgrading Security	21
Upgrading the Administrative Configuration	21
Upgrading the Database	22
Methodology	22
Migration Phases	22
How Person Data is Migrated	22
Obsolete Tables	23
Database Customizations	23
About eGate Queuing	23
Upgrading the Java API for eIndex Active Integration	23

Appendix A

Changes to Monk	24
Changes to Monk's Behavior	24
Glossary	26
e*Gate 4.x Terms in eGate 5.0	33
Index	35

Introduction

Welcome to the *SeeBeyond ICAN Suite Upgrade Guide*. This document includes information about upgrading to the SeeBeyond® Integrated Composite Application Network (ICAN) Suite™.

This chapter includes

- [“Organization of Information” on page 5](#)
- [“Intended Audience” on page 6](#)
- [“Writing Conventions” on page 6](#)
- [“Online Documents” on page 7](#)
- [“The SeeBeyond Web Site” on page 7](#)

1.1 Organization of Information

This document provides information about upgrading to eGate Integrator 5.0 and includes the following chapters:

- **Chapter 1 “Introduction”** describes the purpose of the *SeeBeyond ICAN Suite Upgrade Guide* including writing conventions and a list of related documents.
- **Chapter 2 “Upgrading e*Gate 4.x”** describes the available options for current e*Gate 4.x customers wishing to upgrade to eGate version 5.0.
- **Chapter 3 “Upgrading eInsight”** describes the available options for current e*Insight 4.x customers wishing to upgrade to eInsight version 5.0.
- **Chapter 4 “Upgrading eIndex”** describes the available options for current e*Index 4.x customers wishing to upgrade to eIndex version 5.0.
- **Chapter A “Changes to Monk”** describes the changes to the behavior of the Monk environment after installing the Schema Runtime Environment.

In addition, refer to the [Glossary](#) on page 26 for a list of eGate-related terminology.

Note: *e*Gate 4.x terms and their eGate 5.0 equivalents are listed in [Table 4 on page 33](#).*

1.2 Intended Audience

This guide is intended for experienced PC users who have the responsibility of helping to set up and maintain a fully functioning ICAN Suite system. This person must also understand any operating systems on which eGate will be installed (Windows and UNIX) must be thoroughly familiar with Windows-style GUI operations.

1.3 Writing Conventions

The following writing conventions are observed throughout this document.

Table 1 Writing Conventions

Text	Convention	Example
Button, file, icon, parameter, variable, method, menu, and object names.	Bold text	<ul style="list-style-type: none"> ▪ Click OK to save and close. ▪ From the File menu, select Exit. ▪ Select the logicalhost.exe file. ▪ Enter the timeout value. ▪ Use the getClassname() method. ▪ Configure the Inbound File eWay.
Command line arguments and code samples	Fixed font. Variables are shown in <i>bold italic</i> .	<code>bootstrap -p <i>password</i></code>
Hypertext links	Blue text	For more information, see "Writing Conventions" on page 6.

Additional Conventions

Windows Systems

For the purposes of this guide, references to "Windows" will apply to Microsoft Windows Server 2003, Windows XP, and Windows 2000.

Path Name Separator

This guide uses the backslash ("\\") as the separator within path names. If you are working on a UNIX system, please make the appropriate substitutions.

1.4 Online Documents

The documentation for the SeeBeyond ICAN Suite is distributed as a collection of online documents. These documents are viewable with the Acrobat Reader application from Adobe Systems. Acrobat Reader can be downloaded from:

<http://www.adobe.com>

1.5 The SeeBeyond Web Site

The SeeBeyond Web site is your best source for up-to-the-minute product news and technical support information. The site's URL is:

<http://www.seebeyond.com>

Upgrading e*Gate 4.x

This chapter describes the options available for e*Gate 4.x customers upgrading to eGate 5.0.

This chapter contains

- [“eGate 5.0 Upgrade Overview” on page 8](#)
- [“Benefits of Upgrading to eGate 5.0” on page 11](#)
- [“Upgrade Procedure” on page 12](#)
- [“Post-Upgrade Deployment Strategies” on page 13](#)

2.1 eGate 5.0 Upgrade Overview

e*Gate 4.x customers integrate their 4.x systems with eGate 5.0 by installing the Schema Runtime Environment (SRE). Installing the SRE upgrades an e*Gate 4.x system to the eGate 5.0 Schema Runtime Environment.

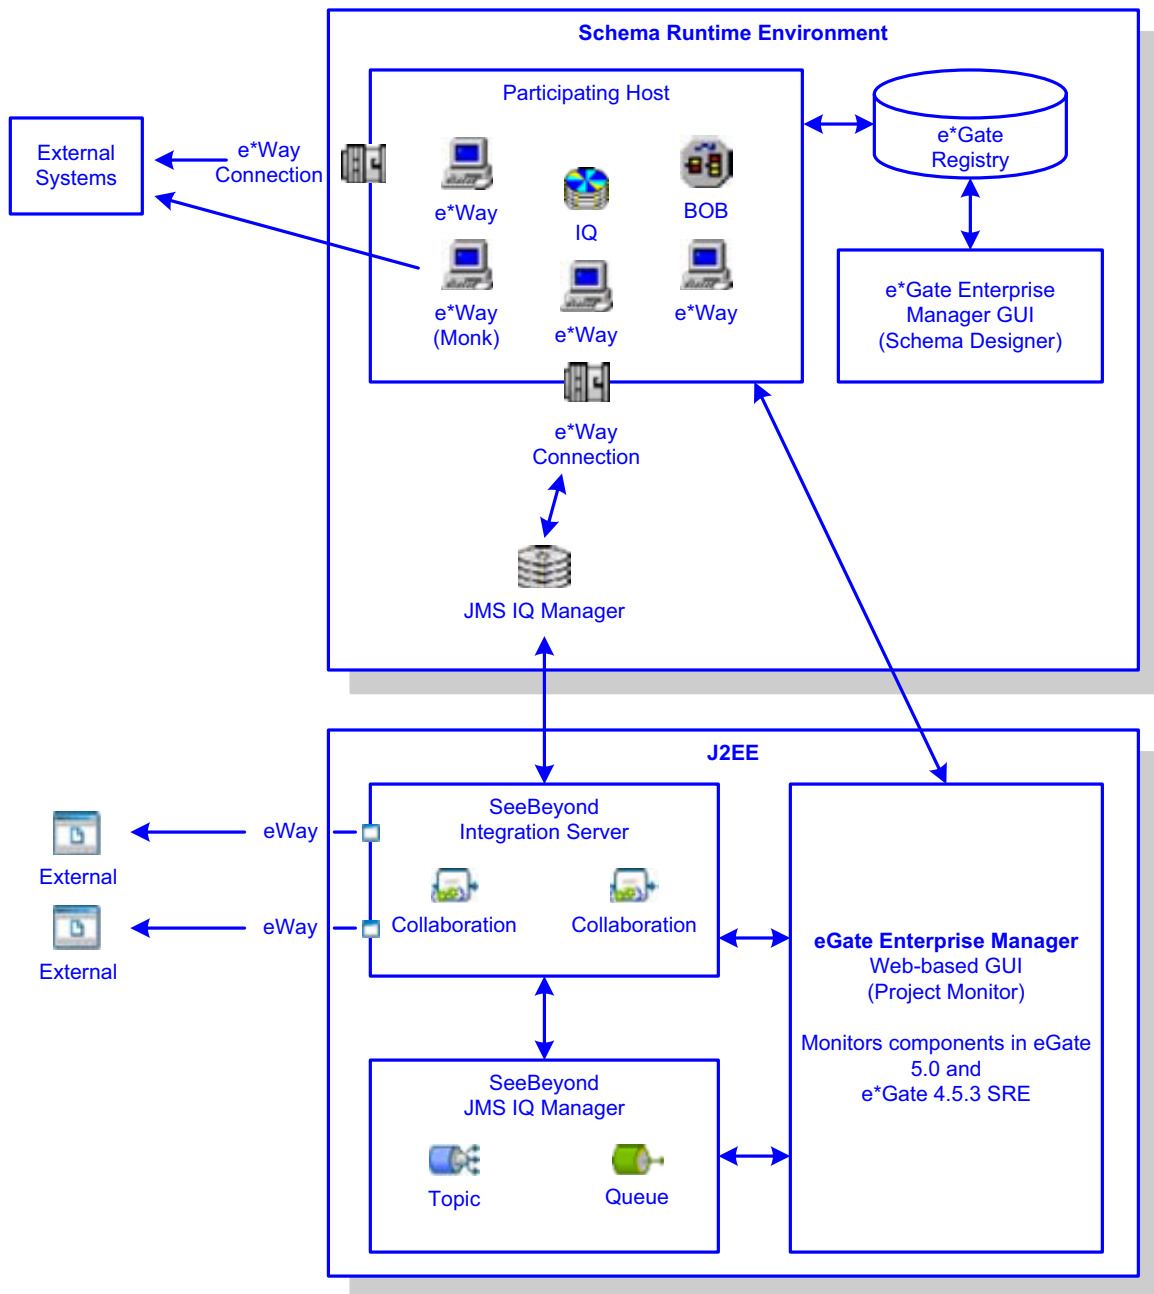
2.1.1. Schema Runtime Environment Upgrade

The Schema Runtime Environment upgrades e*Gate version 4.x systems to take advantage of eGate 5.0 tools, such as the monitoring features in the Enterprise Manager web-based GUI.

The Schema Runtime Environment (SRE) makes it possible for upgraded e*Gate Schemas to interact directly with an eGate 5.0 J2EE Project. This enables SRE Collaborations to publish and subscribe to J2EE JMS IQ Managers. Similarly, J2EE Collaborations can publish and subscribe to JMS IQ Managers running in the SRE. This inter-connectivity is achieved when J2EE Projects and SRE Schemas publish and subscribe to each other’s message services. This inter-operability can be established without re-writing existing Java and Monk Collaborations.

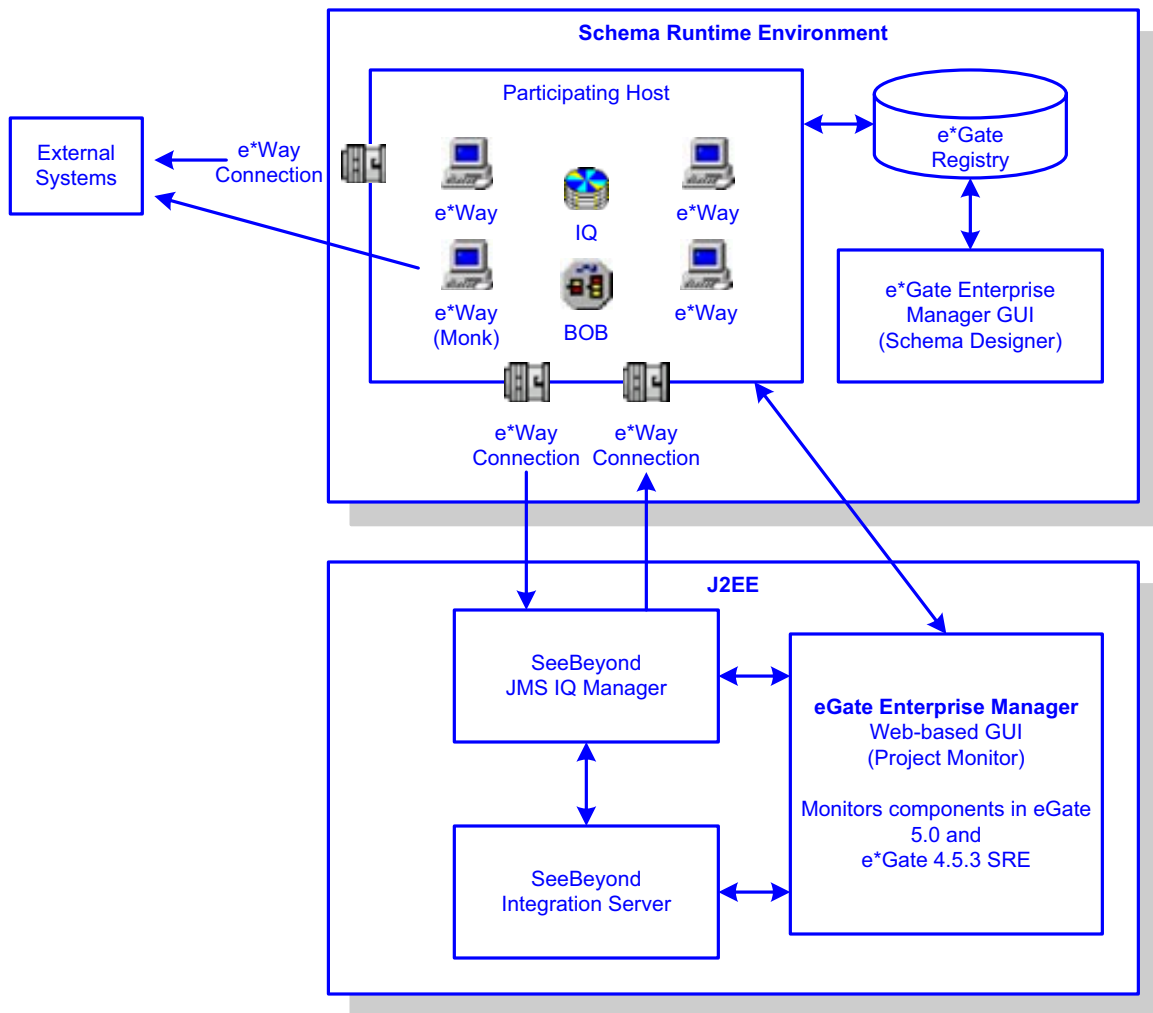
Figure 1 and Figure 2 illustrate the two ways Projects and Schemas can interface with a common message service. In the first example, the J2EE components publish and subscribe to the SRE JMS IQ Manager.

Figure 1 J2EE Interfacing with the SRE JMS IQ Manager



In the second example (see Figure 2), components in the SRE publish and subscribe to the J2EE JMS IQ Manager through the use of e*Way Connections. These e*Way Connections provide a bidirectional flow between Projects and Schemas.

Figure 2 SRE Interfacing with the J2EE JMS IQ Manager



Although either of these scenarios provide similar functionality, it is generally recommended to configure the J2EE Projects to publish and subscribe to the SRE environment (as shown in [Figure 1 on page 9](#)).

For more information on how to configure Projects and Schemas to use a common message service, see [“Post-Upgrade Deployment Strategies” on page 13](#).

General Features of the SRE

Figure 1 and Figure 2 show how SRE components running in the Schema Runtime Environment interact with the eGate J2EE system.

- Components on the Participating Host in the SRE interact with each other and with external systems with the same functionality that existed prior to upgrading.
- The Schema Runtime Environment upgrade includes the familiar GUIs and editors that were used to create the Schema components in e*Gate 4.x.
- Components running in the SRE can be monitored via the eGate J2EE browser-based Enterprise Manager. The Enterprise Manager displays the status of the

components running in the SRE. Additionally, users can start and stop components using the control features of the eGate 5.0 Enterprise Manager.

SRE Supported Platforms

The Schema Runtime Environment upgrade for eGate 5.0 is supported on the following operating systems:

- Windows XP and Windows 2000 SP2 and SP3
- Solaris 8 and 9
- AIX 5.1 and 5.2
- HP-UX 11 and 11i (RISC)
- HP Tru64 V5.1A
- Red Hat Linux Advanced Server 2.1 (Intel)

2.2 Benefits of Upgrading to eGate 5.0

Upgrading to eGate 5.0 provides improved monitoring capabilities and enables Schema/Project integration. Additionally, customers using all of the products in the SeeBeyond ICAN Suite, such as eInsight and eXchange, will enjoy the benefits of the tight integration between all products in the version 5.0 suite.

For a complete list of enhancements in eGate 5.0, see the SeeBeyond web site at:

<http://www.seebeyond.com>

2.2.1. Improved Monitoring

The eGate 5.0 Enterprise Manager provides enhanced monitoring functionality of Project and Schema components. Additionally, the Enterprise Manager uses your Internet Explorer browser and does not require any additional client applications to be installed.

2.2.2. Schema and Project Integration

Upgrading to the eGate 5.0 Schema Runtime Environment makes it possible for:

- SRE components to publish to and subscribe from J2EE Projects.
- J2EE components to publish to and subscribe from SRE Schemas.

2.3 Upgrade Procedure

Upgrading the e*Gate 4.x components is achieved by installing the Schema Runtime Environment. Install the SRE upgrade on your e*Gate 4.0x Registry Host, Participating Hosts, and GUI machines.

The SRE installation program:

- Installs the eGate 5.0 SRE program files, such as the Schema Designer GUI (formerly known as the e*Gate Enterprise Manager).
- Upgrades all of your existing components to run in the eGate 5.0 Schema Runtime Environment.
- Installs all of the add-on program files including all of the e*Ways and libraries licensed for your enterprise.

2.3.1. Pre-Installation Steps

Prior to installing the Schema Runtime Environment upgrade, you must complete the following steps:

- 1 Verify that your servers meet the requirements specified in **“SRE Supported Platforms” on page 11**.
- 2 Analyze all Schemas to determine what add-on products are installed. This will be critical when upgrading the add-ons in step 4 of the **“Installation Steps”** procedure.
- 3 Check for sufficient disk space to run the upgrade installation.
- 4 Export all Schemas from the e*Gate Registry and backup the export files.
- 5 Shut down all publishing Collaborations and allow the subscribing Collaborations to empty all queues.
- 6 Create a backup of your entire e*Gate 4.x system to be used in case the upgrade process is interrupted.

2.3.2. Installation Steps

Depending on the number of SeeBeyond Business Integration Suite™ products, add-ons, and libraries used by your environment, there can be as many as twelve CD-ROMs for your Schema Runtime Environment upgrade disc set.

The process of installing the Schema Runtime Environment upgrade includes these high-level steps:

Note: For complete instructions for installing the Schema Runtime Environment upgrade, see the *e*Gate Integrator Installation Guide for Schema Runtime Environment (SRE)*.

- 1 Install the Registry Host upgrade found on **e*Gate Integrator SRE Disc 1** on the Registry Host.

- 2 Install the Participating Host upgrade found on **e*Gate Integrator SRE Disc 1** on all Participating Hosts.
- 3 Install the e*Gate GUI upgrade found on **e*Gate Integrator SRE Disc 2** on all machines using the ETD Editor and Collaboration Editor GUIs.
- 4 Install the appropriate add-on upgrades found on **e*Gate Add-ons/Samples Discs 1 through 5, UN/EDIFACT Templates Discs 1 and 2, and X12 Templates Disc1 through 3** on the appropriate Participating Hosts.

2.4 Post-Upgrade Deployment Strategies

After installing the Schema Runtime Environment upgrade, you will need to decide how you want to publish and subscribe to messages between the SRE Schema and the J2EE Project. The two strategies are illustrated in [Figure 1 on page 9](#) and [Figure 2 on page 10](#). The main difference between the two strategies is the location of the common message server.

- Figure 1 shows a JMS IQ Manager running in the SRE. The JMS IQ Manager is the host for any common messages that are exchanged between the two environments. As the figure shows, the J2EE Project components can publish and subscribe to topics and queues in the SRE JMS IQ Manager *or* in the J2EE JMS IQ Manager.
- Figure 2 shows components in the SRE Schema publishing and subscribing to the J2EE JMS IQ Manager through e*Way Connections. Similarly, the SRE components can publish and subscribe to topics and queues in the SRE JMS IQ Manager *or* in the J2EE JMS IQ Manager.

These two deployment strategies are highlighted in the following table.

Table 2 IQ Manager Configuration Overview

Deployment Strategy	Components to Configure
SRE Schemas and J2EE Projects use a common SRE JMS IQ Manager. For more information, see “Using a Common SRE JMS IQ Manager” on page 14	To configure your eGate Schemas to exchange data with eGate J2EE Projects, the following components must be configured in the <i>SRE</i> : <ul style="list-style-type: none"> ▪ SRE JMS IQ Manager ▪ eWay Connection(s) For specific information, see “SRE Configuration” on page 14
	To configure your eGate J2EE Project components to exchange data with SRE Schema components, the following components must be configured in the <i>eGate J2EE environment</i> : <ul style="list-style-type: none"> ▪ Install stcmessageserver45x.sar. ▪ Configure topics and queues to use the SRE IQ Manager. Fore more information, see “J2EE Environment Configuration” on page 14

Table 2 IQ Manager Configuration Overview (Continued)

Deployment Strategy	Components to Configure
J2EE Projects and SRE Schemas use a common J2EE JMS IQ Manager. For more information, see “Using a Common J2EE IQ Manager” on page 15	To configure your eGate Schemas to exchange data with eGate J2EE Projects, the following components must be configured in the <i>SRE</i> : <ul style="list-style-type: none"> ▪ Configure the eWay Connection(s) to use the J2EE IQ Manager. For more information, see “SRE Configuration” on page 15
	To configure your eGate J2EE Project components to exchange data with SRE Schema components, the following components must be configured in the <i>eGate J2EE environment</i> : <ul style="list-style-type: none"> ▪ J2EE IQ Manager For more information, see “J2EE Environment Configuration” on page 16

2.4.1. Using a Common SRE JMS IQ Manager

[Figure 1 on page 9](#) shows an SRE Schema and a J2EE Project sharing a common JMS IQ Manager residing in the Schema Runtime Environment. In this scenario, messages are exchanged via the SRE JMS IQ Manager. The following sections describe the items that must be configured in the SRE and the J2EE environment.

SRE Configuration

To configure your eGate Schemas to exchange data with J2EE Projects, the following components must be configured in the Schema Runtime Environment:

- SRE JMS IQ Manager: Add a JMS IQ Manager to your Schema. This is the message service that will be used by the SRE and J2EE components.
- eWay Connection(s): Add one or more e*Way Connections and configure them to use the SRE JMS IQ Manager. Configure the e*Way Connection to connect to the JMS IQ Manager you previously added to the SRE.

For details on adding and configuring JMS IQ Managers and e*Way Connections in the SRE, see the *SeeBeyond JMS Intelligent Queue User’s Guide* in the SRE documentation suite.

J2EE Environment Configuration

To configure your J2EE Project to integrate with the SRE JMS IQ Manager, the following actions must be taken:

- If you have not already done so, upload **stcmessageserver45x.sar** to your J2EE Repository. This adds the functionality required for J2EE Projects to communicate with an SRE JMS IQ Manager.

For complete instructions for uploading product .sar files to the Repository, see the *SeeBeyond ICAN Suite Installation Guide* in the ICAN documentation suite.

- Add a new “SeeBeyond Message Server 4.5.x” component to a Logical Host in your J2EE environment. Configure this component to use the host name and port for the JMS IQ Manager that you configured in the SRE.

For details on adding and configuring the SRE JMS IQ Manager, see the *eGate Integrator JMS Reference Guide* in the ICAN documentation suite.

2.4.2. Using a Common J2EE IQ Manager

Figure 2 on page 10 shows a J2EE Project and an SRE Schema sharing a common JMS IQ Manager residing in the J2EE environment. In this scenario, messages are exchanged via the J2EE JMS IQ Manager. The following sections describe the items that must be configured in the SRE and the J2EE environment.

SRE Configuration


To configure your eGate Schemas to exchange data with a J2EE JMS IQ Manager, the following components must be configured in the Schema Runtime Environment:

- **e*Way Connections:** Configure the inbound and outbound e*Way Connections to publish and subscribe to the J2EE JMS IQ Manager. Pay special attention to these parameters in the e*Way Connections:
 - ♦ **HostName:** This is the host name of the Logical Host hosting the J2EE JMS IQ Manager.
 - ♦ **PortNumber:** This is the listener port for the J2EE JMS IQ Manager.
 - ♦ **ConnectionType:** Specify **Queue** or **Topic** to match the configuration of the J2EE JMS IQ Manager.
 - ♦ **DefaultOutgoingMessageType:** Specify **Bytes** or **Text** to match the message type of the messages in the J2EE environment.

Publishing from Non-Java Collaborations

In the SRE, JMS IQs and Standard IQs subscribe to Events in an IQ published by a specific Collaboration.

Figure 3 SRE Publications and Subscriptions



Messages published by J2EE Collaborations do not contain the information about the publisher that non-Java Collaborations (Monk, Passthrough, or C) require as subscribers. To overcome this, you must add a Multi-mode e*Way or a Business Object Broker (BOB) to the Schema. This e*Way can use an e*Way connection to connect to the JMS IQ Manager as an external system. When the Multi-mode e*Way enqueues the

messages, it will create properly formed e*Gate Events that can be subscribed to by SRE Collaborations.

J2EE Environment Configuration

SRE components can publish and subscribe to the J2EE JMS IQ Manager without any special configuration in the J2EE environment.

Chapter 3

Upgrading eInsight

This chapter discusses the steps you need to upgrade eInsight.

This chapter contains

- [“Overview” on page 17](#)
- [“Integrate Existing Business Processes and Schemas” on page 18](#)
- [“eInsight Upgrade Procedure” on page 17](#)
- [“Connect Using the Schema Runtime Environment” on page 18](#)

3.1 Overview

When considering an upgrade to eInsight 5.0, you should examine the number and complexity of your existing Business Process Models. If you have several complex Business Process Models in production, you may choose to integrate your current configuration with eInsight. This option allows you to leave your existing processes in place, while gaining the latest monitoring abilities and allowing you to create your new Business Process Models in eInsight.

3.1.1. eInsight Upgrade Procedure

The process of installing the Schema Runtime Environment (SRE) includes these high-level steps:

- 1 Install the eInsight Schema Run-time Environment
- 2 Register Components with the Repository
- 3 Connect Schema Components to Projects
- 4 Connect Project Components to Schemas
- 5 Monitor Components with the Enterprise Manager

Note: See the *SeeBeyond ICAN Suite Upgrade Guide* for more information about upgrading your e*Gate Integrator installation.

3.2 Integrate Existing Business Processes and Schemas

To integrate your existing e*Insight Business Processes and related Schemas, you will use the **Schema Runtime Environment (SRE)**. Integrating your Business Processes and Schemas gives you the flexibility to use your existing implementation with the benefits of the new browser-based monitoring tools available in the Enterprise Manager.

3.2.1. Connect Using the Schema Runtime Environment

The SRE allows your systems to take advantage of ICAN tools by providing existing e*Insight Schemas the ability to interact directly with the SeeBeyond JMS IQ Manager. eInsight 5.0 Projects and e*Insight 4.5.x Schemas (running in the SRE) can publish and subscribe to each other's services. This inter-operability is established without re-writing existing Java and Monk Collaborations.

The SRE includes e*Gate and e*Insight GUIs and editors needed to maintain the upgraded components. See the *SeeBeyond ICAN Suite Upgrade Guide* for a list of platforms supported by the Schema Runtime Environment.

Upgrading eIndex

This chapter provides an overview of the process for upgrading an eIndex Global Identifier (eIndex) environment from version 4.5.3 to version 5.0.

This chapter includes

[“Upgrade Overview” on page 19](#)

[“System Requirements” on page 19](#)

[“Upgrade Tasks” on page 20](#)

4.1 Upgrade Overview

Upgrading the eIndex environment consists of two distinct tasks: upgrading the eGate environment (Schemas and eWays) and upgrading the database. While eGate provides the Schema Runtime Environment (SRE) that allows you to run your current 4.5.3 Schemas in an eGate 5.0 environment, we recommend that you create a new Project and Collaborations for eIndex 5.0 to accommodate the new data structure and to take advantage of the new Java API. Upgrading the database involves migrating your production data from a copy of the 4.5.3 production database to a new 5.0 database.

4.2 System Requirements

eIndex 5.0 supports different operating system, database, and integration software platforms than eIndex 4.5.3. If you need to upgrade any of the following components, be sure to complete the upgrades before beginning the eIndex upgrade. The following platforms are supported by eIndex 5.0.

- **Operating Systems**

eIndex 5.0 supports the following operating systems.

- ♦ Windows XP, 2000, and 2003
- ♦ AIX 5.1 and 5.2
- ♦ HP-UX 11, 11i, and TRU64 5.1A
- ♦ Sun Solaris 8 and 9
- ♦ Red Hat Linux 8, Advanced Server 2.1

- **Database Platforms**
eIndex 5.0 supports Oracle 8.1.7 or 9i, but does not support Microsoft SQL Server or Sybase at this time.
- **eGate Integrator**
eIndex 5.0 can only be installed with the eGate Integrator version 5.0. The Oracle eWay is no longer required for eIndex.

Important: Your eIndex environment must be at version 4.5.3 in order to upgrade to version 5.0.

4.3 Upgrade Tasks

Moving your eIndex environment to version 5.0 requires upgrading several components. The following steps are required to complete the eIndex upgrade.

- [Upgrading the Schema to a Project](#) on page 20
- [Upgrading the Outbound Queue](#) on page 20
- [Upgrading the GUI](#) on page 20
- [Upgrading the Database](#) on page 22
- [Upgrading the Java API for eIndex Active Integration](#) on page 23

4.3.1. Upgrading the Schema to a Project

eIndex 5.0 provides a Project that processes data in the same manner as the eIndex 4.5.3 sample schema. To upgrade the schema, install eIndex in the Enterprise Designer, and then customize the new Project based on the customizations you made to the 4.5.3 schema. In eIndex 5.0, you will be using Java methods rather than Monk functions to customize how data is processed.

4.3.2. Upgrading the Outbound Queue

In previous versions, when certain updates were made using the eIndex GUI, eIndex published messages to a database table (the “outbound queue”) with the new information. These messages were removed from the outbound queue by the polling eWay, and forwarded to eGate for routing to external systems. eIndex 5.0 publishes all messages processed through eIndex directly to a JMS Topic for routing. To upgrade the outbound queue, you may need to create Collaborations to publish messages to external systems. The outbound queue in the eIndex 4.5.3 database must be empty before beginning the database portion of the upgrade.

4.3.3. Upgrading the GUI

eIndex 4.5.3 had three user interfaces: eIndex Global Identifier, eIndex Administrator, and eIndex Security. eIndex 5.0 retains only the global identifier GUI, which is now a

web-based browser called the *Enterprise Data Manager* (EDM). The functions previously performed using the eIndex Administrator are either obsolete or are now handled through the XML configuration files in the eIndex Project. Security functions are now handled through the configured Environment.

Upgrading to the EDM

Upgrading to the EDM requires that you have a web browser available to connect to the integration server for eIndex. If the standard appearance of the eIndex Global Identifier 4.5.3 GUI was customized, these customizations must be made manually to the Enterprise Data Manager file in the eIndex Project. This includes showing and hiding fields, modifying labels, modifying field names, defining field formats, and so on.

Upgrading Security

Previously, eIndex handled security through its own security module, which restricted access at the user, function, and action levels. eIndex 5.0 relies on the Environment used in the deployment to handle security. Any security information previously stored in eIndex must be transferred manually to the Environment. eIndex no longer supports the Event Notification function; however, this can be handled through eGate and the e-mail eWay.

Upgrading the Administrative Configuration

The eIndex Administrator allowed you to maintain several types of data stored in the eIndex database, including:

- Control key values.
- Configurable queries.
- Common table data.
- INTEGRITY rule set information.

Of these, only common table data is stored in the eIndex 5.0 database, and is the only component that is migrated automatically.

Control Key Values

Most control keys have been removed as obsolete. Weight thresholds must be defined manually in the eIndex Manager Service (the Threshold file).

Configurable Queries

The configurable queries define fields to block for matching, and the combinations of fields allowed for a phonetic search. Previously, the configurable queries were modified through eIndex Administrator, and the information was stored in the database. In version 5.0, the configurable blocking queries are customized in the Query Builder (the Candidate Select file), and the information is stored in that file. eIndex 5.0 does not provide the ability to include constants in the “where” clause of the SQL statements, but it provides greater flexibility in choosing fields on which to block.

Extensive searching was previously defined in the configurable query for record processing, and by a control key for GUI searches. In version 5.0, it is defined by the

blocker query and the ALPHA-SEARCH query in the Query Manager. This information must be transferred manually.

Common Table Data

All common table data is automatically migrated into the eIndex 5.0 database.

INTEGRITY Rule Set Files

Rule set information is no longer stored in the eIndex database, making the Rule Set Maintenance function obsolete. Rule set files only need to be upgraded if they have been customized in version 4.5.3. To upgrade the rule set files, the customized 4.5.3 rule set files must be copied to the machine on which the integration server resides, and the path to the rule set files must be specified by the environment variable VTICFG.

4.3.4. Upgrading the Database

Upgrading the database involves migrating your production data from a copy of the 4.5.3 production database to a new 5.0 database. eIndex 5.0 includes several new tables, and makes several of the 4.5.3 tables obsolete.

Methodology

The migration process involves three databases: the production database, an eIndex 4.5.3 source database, and an eIndex 5.0 target database. To begin the upgrade process, data is exported from the production database and then imported into the 4.5.3 source database to create an exact replica of the production database. After this, the production database is not touched for the remainder of the migration and can continue to be used as the production database. The data is migrated from the 4.5.3 source database to the 5.0 target database. The 5.0 target database becomes the production database once the upgrade and testing are complete.

To ensure no information is lost during the upgrade, inbound messages are stored in the eGate 5.0 queue while the upgrade is performed. Once the export of the production database begins, no changes to person data can be made from the eIndex Global Identifier GUI.

Migration Phases

The migration of data from the 4.5.3 source database to the 5.0 target database is performed in two phases.

- 1 Migrate common table data, systems, and current person data
- 2 Migrate transaction history data

How Person Data is Migrated

When you migrate person information from the 4.5.3 source database to the 5.0 target database, current person data is migrated first. This data is inserted into the system object tables and the single best record (SBR) tables. History information is no longer stored as complete person records in the “*_history” database tables. This information is now stored in one table, sbyn_transaction, which stores a history of only the changes

made in each transaction. This means that migrating transaction history data requires the program to translate each history record to a description of the changes that occurred for that transaction. The data migration process inserts each translation into a row in `sbyn_transaction`, with a complete description of the changes made.

Obsolete Tables

The structure of eIndex 5.0 simplifies the database, and several tables are now obsolete. Most of the information previously stored in these tables is now stored in the XML configuration files, which can be updated as needed. The history tables have been consolidated into a single delta table, and the functions handled by both the security tables and the outbound event tables have been transferred to eGate. Rule set information is now only stored in the rule set files on the integration server, and is no longer stored in the database as well.

Database Customizations

If you have customized your eIndex database by creating any of the following items, they will not be transferred to the new database during the migration process. To maintain the functionality of these items, you must add them to the new eIndex 5.0 database manually.

- New database tables
- New columns added to existing tables
- New stored procedures
- New triggers
- New views

About eGate Queuing

When you export your production database, you must queue any incoming transactions to the 4.5.3 Schema to avoid losing any information. At the same time, you also begin queuing the incoming transactions to the 5.0 Project. During the migration process, you can continue to use the eIndex 4.5.3 production database to process daily transactions, but at the same time you must continue queuing those transactions in the 5.0 Project. When the 5.0 database is ready to move into production, you can start processing the stored Events from the 5.0 queue to bring the database current.

4.3.5. Upgrading the Java API for eIndex Active Integration

eIndex 4.1.1 and 4.5.3 provide Java API methods that can be called by a third party to access information in the eIndex database. This API must be upgraded manually for version 5.0. Customers currently on 4.5.3 who use the functions directly from the eGate Collaborations can upgrade by using the method OTDs rather than writing straight Java in the Collaborations.

Changes to Monk

This chapter describes the changes to the behavior of the Monk environment after installing the Schema Runtime Environment.

This chapter contains

- [“Changes to Monk’s Behavior” on page 24](#)

1.1 Changes to Monk’s Behavior

The following table lists the changes that have been made to the Monk engines or specific Monk functions since e*Gate 4.1.1. For additional details of the issues listed below, see the ESR readme.txt referenced for that specific issue.

Table 3 Monk Updates

Version	Description	ESR
4.1.1	Repeated calls to the open-pipe Monk function cause processes to be spawned without shutting down properly. This ESR fixed the issue without changing the behavior of the Monk function.	44338
	This ESR fixes a problem where SWIFT messages produce the default data of optional nodes even when no data is entered into that field.	29844
	Changes were made to the following Monk functions to incorporate BCD (batch data communication) changes: <ul style="list-style-type: none"> ▪ sapbapi-struct-fetch.monk ▪ sapbapi-struct-raise.monk ▪ sapbapi-struct-send.monk ▪ sapbapi-struct-call.monk ▪ saptrfc-struct-call.monk ▪ saptrfc-send-idoc4.monk ▪ saptrfc-receive-idoc4.monk The copy-strip command caused parsing errors and truncated any data fields ending in characters considered to be spaces by copy-strip . All copy-strip commands were also removed from these Monk functions.	28037
	The file stc_monksap.dll was modified so that it would not interpret BLOB data as C-string data, which would otherwise result in the premature ending of all BLOBs by a NULL character (such as from a BCD field). The net error would be that fields defined after a BCD field would be disregarded, causing the BAPI to malfunction.	28037

Table 3 Monk Updates (Continued)

Version	Description	ESR
4.1.1 (Continued)	Added ETD templates to support updated EDIFACT libraries (d95a to d00a).	26257
4.1.2	Updated the EDIFACT ETD templates to remove the periods in some of the node names. Monk interprets the periods to be separators for subnodes and parses the nodes incorrectly.	27704
	After applying ESR 50507, the path-valid? function no longer works. This ESR corrects that problem.	52154
	After applying ESR 40201, the open-random-access-file function no longer works. This ESR corrects that problem.	48403
	After applying ESR 41845, any e*Ways based on stcewipmp.exe crash when their Collaborations execute the \$event-clear function. This ESR corrects that problem.	47918
	This ESR fixes a problem where the path-nodedepth function returns the incorrect depth when passing it an iteration of a repeating node.	34585
	This ESR provides an enhancement. After applying this ESR, a Monk exception will be raised when no data is mapped for a required node which has an optional parent node.	33931
	This ESR adds additional Monk functions for converting data between ASCII and EBCDIC systems. For a complete list of the conversion functions, see the Monk Utility Functions chapter in the Monk Developer's Reference. Additional conversion functions (although not part of this ESR) are listed in the International Conversion Functions chapter.	31801 and 32344
4.5	This ESR fixes a problem where the addition operator (+) incorrectly returns non-integer sums when two integers are added.	35862
4.5.1	A memory leak occurred after applying ESR 43731. This ESR fixed that problem.	46257
	This ESR resolves a problem with the copy-strip function. The copy-strip function did not return the correct value when the length of the byte in multi-byte data.	43060
	This ESR resolves a problem where the set! Monk function does not resolve the index of repeating nodes when it is run against a path (for example, set! current_var ~in%root.child[<i>].data).	42500
4.5.2	None	
4.5.3	The file-exists and open-input-file functions were enhanced to accommodate files larger than 2 GB.	
	This ESR fixes a problem with positional based escape delimiters.	54189
	This ESR fixes a problem where a Monk exception was being generated when a local variable defined in a let expression is passed as an argument to another function because the variable is unresolved.	53526

Glossary

BI

Business integration (also Business Intelligence).

Collaboration

See “**Service**” and “**Collaboration Definition**”.

Collaboration Definition

The encoding of business rules, in Java or XSLT format. Typically, the encoding consists of operations on OTDs (see “**OTD**” on page 30). Several Collaborations can have the same Collaboration Definition.

Connection

Consists of the configuration information that enables an eWay to connect to an external system.

Connectivity Map

Contains business logic and routing information about the data transmission. A Connectivity Map usually includes one or more Collaborations, Passthrough Collaborations, topics, queues, and eWays. A Connectivity Map is created under a Project. A Project may have multiple Connectivity Maps.

Constants

A name or value pair that is visible across a Project.

CRM

Customer Relations Management

Data Cleansing

Data must be cleansed of errors in structure and content before it is useful in data warehousing and integration; this means transforming data for accurate and effective use in a database or data management system by cleansing “dirty” or redundant data.

Data Dictionary

Defines the organization of a database and lists all files in the database, the number of records in each file, and the names and types of each field. The data dictionary is often hidden from end users. Although the dictionary doesn’t contain actual data, it does contain essential information for managing the database.

Data Integrity

Refers to the accuracy and validity of data. Data integrity can be compromised in many ways, including human error through data entry, or through faulty logic in

programming. Computer viruses, software bugs and many other factors can also compromise data integrity.

Data Mapping

In relational databases (RDBMSs) data mapping is the relationship and data flow between source and target objects. Mapping involves structuring the relationship between source and target objects.

Data Mart

A smaller, focused, database designed to help managers make business decisions. (A data warehouse is a larger, enterprise, database(s).)

Data Mining

Used to synthesize or isolate unique data patterns to predict future behaviors or to filter data to select patterns that help discover previously unknown relationships among data. Commonly used by marketers who acquire and distill consumer information.

Data Transformation

Data transformation is necessary after extracting data from legacy data formats, or any format that requires cleansing. Data is transformed for efficient use for Business-to-Business Enterprise Data Integration.

Data Warehouse

A copy or view of enterprise transaction data (sometimes non-transaction data) that is used for reporting. The data is often summarized and always structured for queries and analysis.

Deployment Profile

Contains the information about how the Project components will be deployed in an Environment. A Project can have multiple Deployment Profiles, but only one Deployment Profile can be activated for a Project in any one Environment.

Derived Collaboration

Collaboration that inherits operations from another, according to standard object-oriented practice.

Dimension Table

Dimension tables describe the business entities of an enterprise; also called lookup or reference tables.

Dirty Data

Dirty data contains, but is not limited to, incorrect data including spelling errors, punctuation errors, incorrect data referencing, incomplete, inconsistent, outdated, and redundant data.

Drill Down

To move from summary to more detailed data by “drilling down” to get it. In database terminology this might mean starting with a general category and drilling down to a specific field in a record.

eGate System

See “Project”.

Environment

A collection of physical resources and their configurations that are used to host Project components. An Environment contains logical hosts and external systems.

EPR

Enterprise Resource Management

ETL

Extract, Transform, Load. Extract is the process of reading data from a source database and extracting the desired subset of data. Transform is the process of converting the extracted data from its previous form into the desired form. Load is the process of writing the data into a larger database.

eView Manager Service

A component of an eView master index that provides an interface to all components and includes the primary functions of the master index.

eWay

A link between a Collaboration and an external connection including the message server connection (topic or queue) or external application.

External Application

A logical representation in an eGate Project of an external application.

External System

A representation in an eGate Project of an external application system.

Extraction

Data are extracted from a source using software tools. This first step in ETL initially “gets” the data.

Fact Table

A fact table typically contains two types of columns: those containing facts and those that contain foreign keys to dimension tables. Fact tables contain detail facts and/or summary facts.

ICAN Suite

The SeeBeyond Integrated Composite Application Network Suite.

Integration Server

J2EE software platform that houses the business logic container used to run Collaborations and JCA connectors (eWays). Provides transaction services, persistence, and external connectivity.

JMS IQ Manager

JMS-compliant, guaranteed delivery store, forwarding, and queueing service.

Join

Matches records, which are joined by a common field, in two tables in a relational database. Often part of a Select query.

Link

The JMS Connection between a Collaboration and a topic or queue in a JMS-compliant message server.

Linked Message Destination

A reference to a Message Destination defined in another Connectivity Map.

Logical Host

An instance of the eGate runtime Environment that is installed on a machine. A Logical Host contains the software and other installed components that are required at runtime, such as application and message servers.

Management Agent

Uses J2EE technology to manage and monitor an eGate 5.0 deployment that may contain other application servers in addition to the SeeBeyond Integration Server. Defines management interfaces and services designed for distributed environments, focusing on providing functionality for managing networks, systems, and applications.

master index

A database application that stores and cross-references information on specific objects in a business organization, regardless of the computer system from which the information originates. Also called *enterprise-wide master index*.

Matching Service

A component of an eView master index that contains the logic for the matching process.

Message Destination

A general term for a topic or queue. Two or more Projects can share a message destination that has the same name and is deployed on the same message server. A single Project may also have a single message destination referenced in multiple Connectivity Maps.

Metadata

“Data about data.” Metadata describes “how,” “when,” and “who” about structure and format, of a particular set of data. ETL tools are used to generate and maintain a central metadata repository.

Non-normalized Data

Non-normalized data cannot be cross-referenced accurately, if at all, and causes manageability issues. Non-normalized data may be converted to normalized data.

Normalized Data

Normalization is a common database design process used to remove redundant or incorrect organization and data. The design and normalization of the database will create a maintainable data set that can be cross-referenced.

Normalized data is not only easier to analyze but also easier to expand. Normalization involves removing redundancy and correcting incorrect data structure and organization.

OLAP

Online analytical processing.

OTD

An acronym for Object Type Definition. OTDs contain the data structure and rules that define an object. An OTD is used in Java Collaboration Definitions for creating data transformations and interfacing with external systems.

Project

Contains a collection of logical components, configurations, and files that are used to solve business problems. A Project organizes the files and packages and maintains the settings that comprise an eGate system in SeeBeyond's Enterprise Designer.

Query

A request for information from a database. There are three query methods:

Choose – With this easy-to-use method, the database system presents a list of parameters from which you can choose. This method is not as flexible as other methods.

Query by example (QBE) – With this method, the system lets you specify fields and values to define a query.

Query language – With this method, you have the flexibility and power to make requests for information in the form of a stylized query using a query language. This is the most complex and powerful method.

Query Builder

A component of an eView master index that defines how queries are processed.

Queue

A JMS queue is a shareable object that conforms to the *point-to-point* (p2p, or PTP) messaging domain, where one sender delivers a message to exactly one receiver. When the SeeBeyond JMS IQ Manager sends a message to a queue, it ensures it is received once and only once, even though there may be many receivers "listening" to the queue. This is equivalent to the subscriber pooling in other queue implementations. You can reference a queue that exists in another Connectivity Map or Project.

Raw Data

Data that has not been turned into "information," through processing. Although factual and "real," raw data is unorganized.

Relational Database (RDBMS)

Short for Relational Database Management System, most often referred to as RDBMS. Data is stored in related tables. Relational databases can be viewed in many different ways.

In this system a single database can be spread across several tables. (RDBMS differs from flat-file databases where each database is self-contained as a single file or table.)

Repository

Stores and manages the setup, component, and configuration information for eGate Projects. The Repository also provides monitoring services for Projects, which include version control and impact analysis.

Schema Runtime Environment

An add-on in eGate 5.0 that provides the upgrade path for e*Gate 4.x users to upgrade to eGate 5.0. Also known as the SRE.

Service

Contains the information about executing a set of business rules. These business rules can be defined in a Java Collaboration Definition, XSLT Collaboration Definition, Business Process, eTL Definition, or other service. A Service also contains binding information for connecting to JMS Topics, Queues, eWays, and other services.

Single Best Record

Also known as the SBR, this is the best representation of an entity's information in an eView master index. The SBR is populated with information from all source systems based on the survivor strategies defined for each field.

Staging Data

Data that is to be processed before entering the warehouse.

Subproject

An independent Project that is included as part of another Project and listed on the Enterprise Explorer tree beneath the main Project icon.

Table

Refers to data arranged in rows and columns, like a spreadsheet. In relational database management systems, all information is stored in tables.

Topic

A JMS topic is a shareable object that conforms to the *publish-and-subscribe* (pub/sub) messaging domain, where one publisher broadcasts messages to potentially many subscribers. When the SeeBeyond JMS IQ Manager publishes a message on a topic, it ensures that all subscribers receive the message.

Transformation

Data that are extracted from databases are transformed into a desired form, using various tools that cleanse, merge, purge, aggregate, calculate, audit, remove redundancy, standardize, etc.

XSLT

An acronym for Extensible Stylesheet Language Transformations. A file format used in eGate to generate Collaboration Definitions.

Update Manager

A component of an eView master index that contains the Java classes and logic that determine how records are updated and how the single best record is populated.

Warehouse

See "Data Warehouse".

e*Gate 4.x Terms in eGate 5.0

Table 4 provides definitions for the terms that are new with eGate release 5.0, as well as equivalent terms from eGate release 4.x.

Table 4 eGate 5.0 Terms

5.0 Term	4.x Equivalent Term
Collaboration	Collaboration
Collaboration Definition	Collaboration Definition
Connection	eWay Connection
Connectivity Map	Closest: Network View of an entire Schema
Deploy	Run the Control Broker
Deployment	<none>
Deployment Profile	Closest: Schema
Enterprise Designer	Enterprise Manager
Enterprise Manager	Enterprise Monitor
Environment	Schema (except only includes physical information, not business logic)
eWay	eWay Connection eWay
eWay Configuration	eWay Connection Configuration
External Application	eWay Connection
External System	eWay Connection
JMS Connection	eWay Connection
Integration Server	<none>
Link	JMS eWay Connection
Linked Message Destination	<none>
Logical Host	Participating Host
Message Destination	Topic or queue
Message Server	JMS IQ Manager
Object Type Definition (OTD)	Event Type Definition (ETD)
Process Manager	Control Broker
Project	Schema (except not including physical layer)
Queue	JMS queue
Repository	Registry
Subproject	Schema

Table 4 eGate 5.0 Terms (Continued)

5.0 Term	4.x Equivalent Term
Topic	JMS topic
XSLT	<none>

Index

B

BI 26

C

Collaboration 26, 33

 derived 27

Collaboration definition 26, 33

Connection 33

 connection 26, 33

Connectivity Map 26, 33

constants 26

Control Broker 33

conventions

 path name separator 6

 Windows 6

CRM 26

D

data cleansing 26

data dictionary 26

data integrity 26

data mapping 27

data mart 27

data mining 27

data transformation 27

data warehouse 27, 32

deploy 33

Deployment 33

Deployment Profile 27, 33

derived Collaboration 27

dimension table 27

dirty data 27

document

 conventions 6

drill down 27

E

eGate system 28

Enterprise Designer 33

Enterprise Manager 33

Enterprise Monitor 33

Environment 28, 33

EPR 28

ETD 33

ETL 28

Event Type Definition 33

eWay 28, 33

eWay Configuration 33

eWay Connection 33

eWay Connection Configuration 33

external

 application 28, 33

 system 28, 33

extraction 28

F

fact table 28

I

ICAN Suite 28

Integration Server 28, 33

J

JMS

 connection 33

 e*Way Connection 33

 IQ Manager 33

 queue 33

 topic 34

JMS IQ Manager 28

join 29

L

link 29, 33

linked message destination 29, 33

Logical Host 29, 33

M

Management Agent 29

message

 destination 29, 33

 server 33

metadata 29

N

network view 33

non-normalized data 29

normalized data 29

Index

O

Object Type Definition 30, 33
OLAP 30
OTD 30, 33

P

Participating Host 33
Process Manager 33
Project 30, 33

Q

query 30
queue 30, 33

R

raw data 30
rdbm 30
Registry 33
relational database 30
Repository 31, 33

S

Schema 33
Schema Runtime Environment 31
Security Server 31
SRE 31
staging data 31
subproject 31, 33

T

table 31
topic 31, 33, 34
transformation 31

W

warehouse 32
writing conventions 6

X

XSLT 31, 34