

Event Replicator Target Adapter

Event Replicator Target Adapter Concepts

Version 3.10

September 2025

This document applies to Event Replicator Target Adapter Version 3.10 and all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2025 Software GmbH, Darmstadt, Germany and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software GmbH product names are either trademarks or registered trademarks of Software GmbH and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software GmbH and/or its subsidiaries is located at <https://softwareag.com/licenses>.

Use of this software is subject to adherence to Software GmbH's licensing conditions and terms. These terms are part of the product documentation, located at <https://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third-Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software GmbH Products / Copyright and Trademark Notices of Software GmbH Products". These documents are part of the product documentation, located at <https://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Use, reproduction, transfer, publication or disclosure is prohibited except as specifically provided for in your License Agreement with Software GmbH.

Document ID: ART-AARTCONCEPTS-310-20250929

Table of Contents

Preface	v
1 About this Documentation	1
Document Conventions	2
Online Information and Support	2
Data Protection	3
2 What You Can Do	5
3 Components	7
4 How It Works	9
Requirements	10
General Processing	11
Table Structure	13
MU and PE Field Support	14
RDBMS Recovery Processing	15
5 Data Type Conversions	17
6 Character Encoding Considerations	19
7 Samples and Examples	21
8 Getting Started	23
9 Best Practices	25
Improving Overall Performance	26
Processing Initial-State Data	26
Index	29

Preface

This document provides an overview of the Event Replicator Target Adapter.

<i>What You Can Do</i>	Describes what you can do using the Event Replicator Target Adapter.
<i>Components</i>	Describes the components of the Event Replicator Target Adapter.
<i>How It Works</i>	Describes the architecture of the Event Replicator Target Adapter.
<i>Data Type Conversions</i>	Explains how data type conversions occur in the Event Replicator Target Adapter.
<i>Character Encoding Considerations</i>	Describes some character encoding issues you should consider when using the Event Replicator Target Adapter.
<i>Samples and Examples</i>	Describes the samples and examples supplied with the Event Replicator Target Adapter.
<i>Getting Started</i>	Lists the steps to get started using the Event Replicator Target Adapter.
<i>Best Practices</i>	Describes best practices for using the Event Replicator Target Adapter, including practices that will improve performance best methods of processing initial-state data, and how best to handle processing between multiple sources and targets.

1

About this Documentation

■ Document Conventions	2
■ Online Information and Support	2
■ Data Protection	3

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
Monospace font	Identifies service names and locations in the format <i>folder.subfolder.service</i> , APIs, Java classes, methods, properties.
<i>Italic</i>	Identifies: Variables for which you must supply values specific to your own situation or environment. New terms the first time they occur in the text. References to other documentation sources.
Monospace font	Identifies: Text you must type in. Messages displayed by the system. Program code.
{ }	Indicates a set of choices from which you must choose one. Type only the information inside the curly braces. Do not type the { } symbols.
	Separates two mutually exclusive choices in a syntax line. Type one of these choices. Do not type the symbol.
[]	Indicates one or more options. Type only the information inside the square brackets. Do not type the [] symbols.
...	Indicates that you can type multiple options of the same type. Type only the information. Do not type the ellipsis (...).

Online Information and Support

Product Documentation

You can find the product documentation on our documentation website at <https://documentation.softwareag.com>.

Product Training

You can find helpful product training material on our Learning Portal at <https://learn.software-ag.com>.

Tech Community

You can collaborate with Software GmbH experts on our Tech Community website at <https://tech-community.softwareag.com>. From here you can, for example:

- Browse through our vast knowledge base.
- Ask questions and find answers in our discussion forums.
- Get the latest Software GmbH news and announcements.
- Explore our communities.
- Go to our public GitHub and Docker repositories at <https://github.com/softwareag> and <https://hub.docker.com/publishers/softwareag> and discover additional Software GmbH resources.

Product Support

Support for Software GmbH products is provided to licensed customers via our Empower Portal at <https://empower.softwareag.com>. Many services on this portal require that you have an account. If you do not yet have one, you can request it at <https://empower.softwareag.com/register>. Once you have an account, you can, for example:

- Download products, updates and fixes.
- Search the Knowledge Center for technical information and tips.
- Subscribe to early warnings and critical alerts.
- Open and update support incidents.
- Add product feature requests.

Data Protection

Software AG products provide functionality with respect to processing of personal data according to the EU General Data Protection Regulation (GDPR). Where applicable, appropriate steps are documented in the respective administration documentation.

2 What You Can Do

The Event Replicator Target Adapter can be used to transform and apply replicated data from your Adabas mainframe database to a relational database, such as DB2, Microsoft SQL Server, MySQL, Oracle, PostgreSQL, Teradata; a JMS queue, or a Software AG Terracotta cache. For more details on Software AG's Terracotta product, please refer to the “IBO and Big Data” documentation on Software AG's [Empower](#) web site.

Once appropriate Event Replicator subscription and destination definitions are activated and the Event Replicator Target Adapter is started, normal processing for the Event Replicator Target Adapter will transform and apply replicated Adabas data to a relational database or Terracotta cache as the data is processed by the subscription. In addition, using the Adabas Event Replicator Subsystem, you can manually submit requests to the Event Replicator Target Adapter that:

- initiate an initial-state request to populate the target tables or cache.
- clear (refresh) data in the target tables or cache.
- delete (drop) relational tables.

For information on using the Adabas Event Replicator Subsystem to do this, read *Submitting Event Replicator Target Adapter Requests* in *Adabas Event Replicator Subsystem User Guide*.

In addition, you can set up the ADALOD utility so that data loaded or updated by the utility is also replicated to the relational database tables. However, read *Replicating ADALOD Data to the RDBMS Tables*, in the *Event Replicator Target Adapter User Guide* for important information on problems that might arise when you do this.

Finally, you can use the Event Replicator Target Adapter to replicate data to JMS queues and topics that are managed by webMethods Broker JNDI providers. The data in these JMS queues and topics can then be made available to webMethods Broker message consumers (such as webMethods Integration Server). For more information on setting up JMS support in the Event Replicator Target Adapter, read *Replicating Data to JMS Queues or Topics*, in the *Event Replicator Target Adapter User Guide*. For complete information about JMS support in webMethods Broker, read the

EntireX Message Service Provider (MSP), available in the IBM EntireX documentation on Software AG's [Empower](#) web site.

**Notes:**

1. User authorization to maintain any new RDBMS tables via Event Replicator Target Adapter is inherited from the site's privilege settings for the database. Authorization is managed by the user's RDBMS privileges and not by Event Replicator Target Adapter. Event Replicator Target Adapter will no longer grant RDBMS privileges to the user. Therefore, if you want to use Event Replicator Target Adapter to maintain tables in an RDBMS, verify that your RDBMS authorization privileges are correct for the maintenance you want to perform.
2. The Event Replicator Target Adapter will not process updates for a field that is the primary key for an RDBMS table or for any fields that comprise a composite key for the table.

3 Components

The Event Replicator Target Adapter makes use of the following component products, in addition to its own code that allows you to transform and apply replicated data to a relational database.

- The Administration tool.

You can use the Administration tool to configure your Event Replicator Target Adapter installation. For more information, read *Event Replicator Target Adapter Administration* in the *Event Replicator Target Adapter Administration Guide*.

- The Data Mapping Tool.

You can use the Data Mapping Tool to generate a global format buffer (GFB) and field table (GFFT) for use with the Event Replicator and the Event Replicator Target Adapter. For more information, read *Using the Data Mapping Tool* in the *Data Mapping Tool User Guide*.

- Replication Monitoring.

Replication Monitoring is a tool that allows users to monitor the status and progress of data being replicated by the Event Replicator for Adabas. For more information, read *Replication Monitoring* in the *Event Replicator for Adabas Replication Monitoring Guide*.

- An internal component called the Software AG ERTA Administration Service.

This service enables running the Event Replicator Target Adapter as a system service and also allows communication with a remote administration console (for example on platforms, such as HP and zLinux, where the Event Replicator Target Adapter Administration tool cannot be installed)

4

How It Works

■ Requirements	10
■ General Processing	11
■ Table Structure	13
■ MU and PE Field Support	14
■ RDBMS Recovery Processing	15

Requirements

The Event Replicator Target Adapter requires the use of:

- An Event Replicator for Adabas subscription for which a global format buffer (GFB) and field table (GFFT) has been *generated*. If no field table has been generated, the Event Replicator Target Adapter will not work. For more information about generating a global format buffer (GFB) and field table (GFFT), read *Generating a GFB using the Adabas Event Replicator Subsystem* (in the *Adabas Event Replicator Subsystem User Guide*) or *Generating a GFB and Field Table Using the Data Mapping Tool* (in the *Data Mapping Tool User Guide*).
- At least one Event Replicator for Adabas destination definition with a destination class type of "SAGTARG". This destination must be used by the subscription. For information on maintaining destination definitions, read *Maintaining Destination Definitions using the Adabas Event Replicator Subsystem* (in *Adabas Event Replicator Subsystem User Guide*).



Note: Destination class types of SAGTARG (DCLASS=SAGTARG) cannot be specified for Adabas or file type destinations (DTYPE=ADABAS or FILE). The SAGTARG destination class type is only valid for IBM EntireX, WebSphere MQ, or NULL destinations. When DCLASS=SAGTARG is specified, the ADARUN RPLPARMS parameter must be set to "FILE" or "BOTH" to provide access to any field table (GFFT) definitions.

- The subscription and destination definitions listed above must be activated. To activate subscription and destination definitions using Adabas Online System (AOS), read *Managing Replication Definitions from AOS* in *Event Replicator for Adabas Administration and Operations Guide*.
- Proper user authorization privileges to maintain the target application.

User authorization to maintain any new RDBMS tables via Event Replicator Target Adapter is inherited from the site's privilege settings for the database. Authorization is managed by the user's RDBMS privileges and not by Event Replicator Target Adapter. Event Replicator Target Adapter will no longer grant RDBMS privileges to the user. Therefore, if you want to use Event Replicator Target Adapter to maintain tables in an RDBMS, verify that your authorization privileges are correct for the maintenance you want to perform.

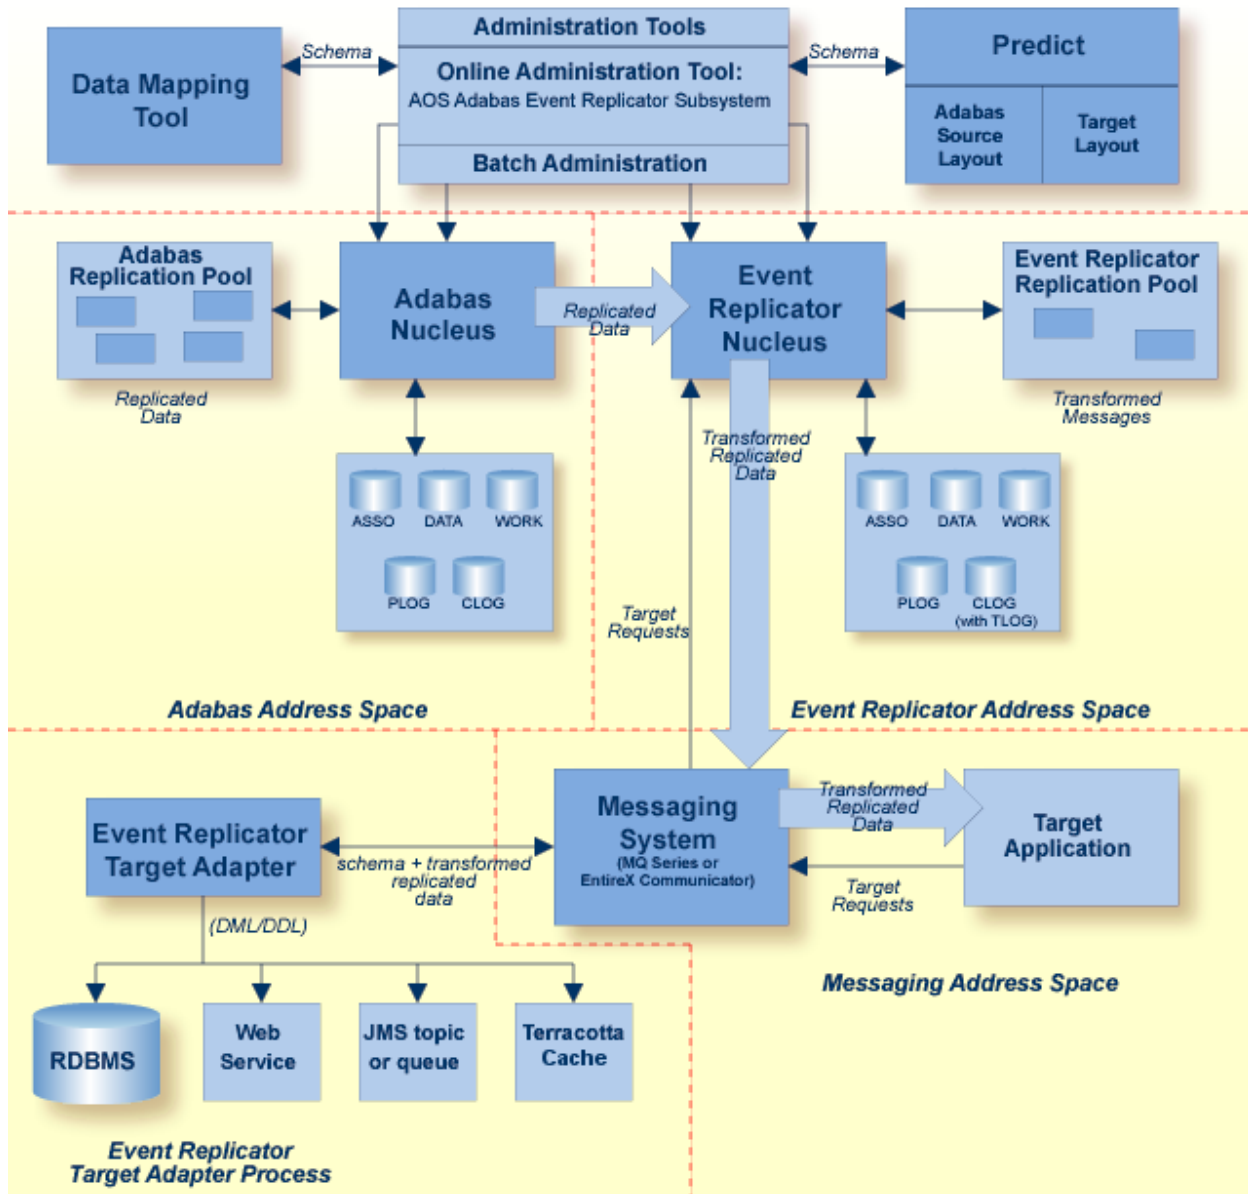
For more information on Event Replicator for Adabas subscription and destination definitions, read *Definition Descriptions* in *Event Replicator for Adabas Administration and Operations Guide*.

General Processing

When a subscription definition and one or more of its associated destination definitions have been defined in the manner described in [Requirements](#), earlier in this section, and if they are activated, the Event Replicator for Adabas automatically creates a schema that maps the replicated data.

The Event Replicator Target Adapter uses the schema to transform and apply the replicated data to your relational database. It will dynamically create tables if they don't exist and populate the tables with Adabas data using insert, update, and delete processing as these processes occur in near real-time in the replicated Adabas file. For more information about activating the Event Replicator Target Adapter, read *Activating Event Replicator Target Adapter Processing*, in the *Event Replicator Target Adapter User Guide*.


Event Replicator for Adabas and Event Replicator Target Adapter high-level processing are depicted in the following diagram.



The Adabas data is replicated as usual by the Event Replicator using a subscription with Event Replicator Target Adapter processing activated. The subscription transforms the replicated data and creates a schema from the generated field table. The schema and the transformed replicated data are then sent to the Event Replicator Target Adapter. The Event Replicator Target Adapter then processes the data and populates and creates appropriate tables or data in the target, as described next in [Table Structure](#).


Table Structure

Event Replicator Target Adapter creates and populates tables with replicated data in the RDBMS as follows:

- 
Note: The Event Replicator Target Adapter will not process updates for a field that is the primary key for an RDBMS table or for any fields that comprise a composite key for the table.

 - If an SQL table has not already been established for the data, the Event Replicator Target Adapter will create one using the XML schema. XML schemas are identified as "Create" operations in the transferred data. Once the SQL table has been established for the data, the replicated data is inserted, updated, and deleted in the table as specified by the "Insert", "Update", and "Delete" operations in the transferred data. Subsequent XML schemas that might be sent are ignored.

 - The name of the SQL table created by the Event Replicator Target Adapter is determined by combining the name of the subscription with the name of the schema file. For example, if the subscription name is PAYROLL and the schema file name (Predict view name) is EMPLOYEE, the table name will be PAYROLL_EMPLOYEE. The names of the columns in the table are taken from the long names in the generated field table.

 - 
Note: If you specify keyword "NOSPRES" for the destination class parameter (DCLASSPARM parameter) in the Event Replicator destination definition, the subscription name will *not* be used to prefix the names of the tables produced by the Event Replicator Target Adapter. When "NOSPRES" is specified (it must be specified in uppercase only), the schema file name (Predict view name) alone is used for the table names. For more information about modifying the destination class in an Event Replicator destination definition, read *DCLASSPARM Parameter* in *Event Replicator for Adabas Reference Guide, Maintaining Destination Definitions Using the Adabas Event Replicator Subsystem* in *Adabas Event Replicator Subsystem User Guide*, or .

 - Event Replicator Target Adapter will process each transaction as a unit-of-work and will only do an SQL commit if all parts of a transaction are successful. For example, if an insert operation is requested and a row with the same primary key already exists in the table, an error occurs and the whole transaction is rejected. Likewise, if an update operation is requested and a row with the specified primary key does not exist in the table, an error occurs and the whole transaction is rejected.
- Error messages are written to the Event Replicator Target Adapter console and the *sqlrep.log* file. Read *Managing Event Replicator Target Adapter Log Files*, in the *Event Replicator Target Adapter User Guide*, for more information.
- All SQL root tables require a primary key. If the schema submitted to the Event Replicator Target Adapter does not specify a primary key (which must be defined to Adabas as a Unique

Descriptor), the Event Replicator Target Adapter uses the ISN as the primary key. You can define as many other unique keys as you like. All keys will be used to create SQL indexes.

- MU and PE fields are supported by the Event Replicator Target Adapter. By default, when MU or PE fields are included in the replicated data, additional tables are created, as described in [MU and PE Field Support](#), next in this section.



Note: Terracotta is not a Relational Database (RDBMS) and data replicated to the cache does not create tables in the same way RDBMS target tables are created. Rather, the cache is populated with data in a structure that mimics the Adabas data structure.

MU and PE Field Support

Multiple-value (MU) and periodic group (PE) fields are supported by the Event Replicator Target Adapter:

- Support for MU fields requires that counters for MU fields be available in the initial Predict data dictionary definition that was used to generate the global format buffer field table for Event Replicator Target Adapter processing. Counters for MU fields are available if the Predict data dictionary definition either specifies fields with field types of MC (MU fields with an automatic counter) or includes explicit CM field (MU counter fields) definitions. For more information about CM, MC, and MU field types, refer to your Predict documentation.
- Support for PE fields requires that counters for PE fields be available in the initial Predict data dictionary definition that was used to generate the global format buffer field table for Event Replicator Target Adapter processing. Counters for PE fields are available if the Predict data dictionary definition either specifies fields with field types of PC (PE fields with an automatic counter) or includes explicit CP field (PE counter field) definitions. For more information about CP, PC, and PE field types, refer to your Predict documentation.

You can control how many occurrences of PE and MU fields are generated in the global format buffer field table, and thus, how many occurrences are available for Event Replicator Target Adapter processing. This is accomplished using a combination of the Predict `Occ` field attribute setting and the **Occurrences used** setting specified when the global format buffer field table is generated. By altering the **Occurrences used** setting, you can specify that the maximum number of occurrences be generated, that no additional occurrences be generated, or that the number of occurrences defined by the Predict `Occ` attribute should be generated. For more information about the `Occ` field attribute, refer to your Predict documentation.

For more information about generating a global format buffer field table, read either *Generating a GFB using the Adabas Event Replicator Subsystem* (in the *Adabas Event Replicator Subsystem User Guide*) or *Generating a GFB and Field Table Using the Data Mapping Tool* (in the *Data Mapping Tool User Guide*).

At this time, the Event Replicator Target Adapter does not support the use of MU or PE fields as primary keys or in composite keys.

By default, when MU or PE fields are included in the replicated data, additional tables are created, as follows:

Field Type	Tables Generated
MU	<p>A new MU table is generated for each MU field. Each value of the MU field will be a row in the MU table.</p> <p>The name of the MU table is comprised of the name of the root table (for example, PAYROLL_EMPLOYEE_ADDRESS-LINE) and the column name of the MU field.</p> <p>The MU table will consist of two columns:</p> <ul style="list-style-type: none"> ■ The reference key will be made up of the root table primary key, or if there is no primary key, the ISN will be used. ■ The MU column name may be defined as an index.
PE	<p>One table is generated for a PE group field. The elements that comprise the PE group field become columns in the table. For example, if the PE field consists of five subfields, five columns are created.</p> <p>The name of the PE table includes the name of the PE field (for example PAYROLL_EMPLOYEE_INCOME). The columns in the table are given the subfield names.</p>

Using the Data Mapping Tool, you can flatten MU and PE fields in the replicated data so that they are replicated as columns in the resulting table, rather than as separate tables. For more information, read *Flattening Fields for Replication*, in *Data Mapping Tool User Guide*.

RDBMS Recovery Processing

If, at any time during normal Event Replicator Target Adapter processing, the Event Replicator Target Adapter detects that the RDBMS is not operational, it will attempt to reconnect to the database at an interval specified by the **Retry Interval** parameter defined for the Event Replicator Target Adapter source definition. The number of times it attempts to reconnect to the database is determined by the setting of the **Retry Count** option. You can set the values for these options, as described in *Maintaining the Event Replicator Target Adapter Engine Configuration*, in the *Event Replicator Target Adapter Administration Guide*. If you use the default settings of these environment entries, the Event Replicator Target Adapter will attempt to reconnect to the RDBMS database every 30 seconds for a total of five attempts.

If a reconnect attempt is successful, Event Replicator Target Adapter processing continues normally. If, however, all of the reconnect attempts fail, the Event Replicator Target Adapter remains up, but the source from which the message was received is shut down. Unprocessed messages remain

in the message queue, so no transactions are lost. The Event Replicator Target Adapter will not remove a transaction from the message queue unless its associated transaction has been successfully committed to the RDBMS database.

If a SQL error is detected in a given transaction, by default the Event Replicator Target Adapter will write the transaction to the *sqlrep.log* file and will shut down the source from which the message was received. However, you can change the value of an environment entry, called **Error Continue**, that logs the error, but allows Event Replicator Target Adapter to continue processing. For more information about error processing, read *Managing Event Replicator Target Adapter Log Files*, in the *Event Replicator Target Adapter User Guide*,. For more information about processing options, read *Configuring Target Definitions for Event Replicator Target Adapter*, *Specifying Target Database Processing Option Definitions*, and *Configuring Source Definitions for Event Replicator Target Adapter*, in the *Event Replicator Target Adapter Administration Guide*,.

5 Data Type Conversions

The following table describes the data type conversions used by Event Replicator Target Adapter during its processing.

Event Replicator Data Type	JDBC Data Types	RDBMS Data Types				Terracotta Data Type
		DB2	MySQL	Oracle	SQL Server	
binary (length less than or equal to 8)	decimal	decimal(22)		raw (see Note 1)	numeric(22)	byte, byte[]
binary (length greater than 8)	varchar	varchar		raw (see Note 1)	varchar	byte[]
date	date	date			datetime	Date
decimal (with precision)	decimal	decimal(len+pre, pre)		number(len+pre, pre)	decimal(len+pre, pre)	float, double
decimal (without precision)	decimal	decimal(len)		number(len)	numeric(len)	short, int, long, BigInteger
float	float	float				float, double
integer	integer	integer				short, int, long, BigInteger
string	varchar	varchar	varchar	varchar2/ nclob (see Note 2)	varchar	String
time	time	timestamp	datetime	timestamp	datetime	Date



Notes:

1. Note 1: binary fields are stored in field of data type raw when 'Target Database Option' dbRawDatatype is set to "true".
2. Note 2: strings with length greater 253 will be store as nclob data types in Oracle target databases.

6 Character Encoding Considerations

XML data sent to Event Replicator Target Adapter is in EBCDIC character encoding. However, support for UTF-8 encoding of the RDBMS data (not table or column names) is available for internationalization purposes. To send XML messages to Event Replicator Target Adapter using UTF-8 format, make sure the `SACODE` parameter is set to "4091" in the Event Replicator subscription definition used for Event Replicator Target Adapter processing.

Wide-character fields sent to Event Replicator Target Adapter are translated to hexadecimal. However, support for UTF-8 encoding is available for internationalization purposes. To send wide-character fields to Event Replicator Target Adapter using UTF-8 format, make sure the `SWCODE` parameter is set to "4091" in the Event Replicator subscription definition used for Event Replicator Target Adapter processing.

For more information about these parameters, read *SACODE Parameter* and *SWCODE Parameter*, in *Event Replicator for Adabas Reference Guide*.



Note: If you want to use UTF-8 character encoding, you must verify that your GFB field lengths are increased as required to accommodate the character set referenced by the code page you select and the data requested in the GFB. You can increase these field lengths manually by editing the GFB itself or by editing the Predict file or data definition module (DDM) used when the GFB is generated.

7 Samples and Examples

A sample of a generated field table, its corresponding XML schema, some replicated data, and the resulting SQL table are provided in *Samples*, in the *Event Replicator Target Adapter User Guide*.

8

Getting Started

➤ To get started using the Event Replicator Target Adapter:

- 1 Install the Event Replicator Target Adapter, as described in *Installing the Event Replicator Target Adapter*, in the *Event Replicator Target Adapter Installation Guide*,.

Make sure you address any Event Replicator Target Adapter installation considerations provided for your relational database, as described in the Event Replicator Target Adapter installation documentation.

- 2 Set up the preferences in the Administration tool so that you can specify target, target option, and source definitions for Event Replicator Target Adapter processing. These preferences identify, among other things, the location of the *context.xml* file used to store the Event Replicator Target Adapter definitions, the locations of the Event Replicator Target Adapter startup and shutdown commands, and the location of the Event Replicator Target Adapter log files.



Caution: We do not recommend that you alter Event Replicator Target Adapter definitions in the *context.xml* file using a text editor or any other editor other than the Event Replicator Target Adapter Administration tool. If you do, errors may result.

For more information, read *Setting Configuration File and Preferences*, in the *Event Replicator Target Adapter Administration Guide*.

- 3 Use the Administration tool to configure the target definitions for your installation. Target definitions define the RDBMS data base or web service targets to which the Event Replicator Target Adapter should direct the replicated data. For more information, read *Configuring Target Definitions for Event Replicator Target Adapter*, in the *Event Replicator Target Adapter Administration Guide*.
- 4 Use the Administration tool to set up any target database processing option definitions needed for Event Replicator Target Adapter processing. Target database processing option definitions specify options regarding how the Event Replicator Target Adapter processes the data sub-

mitted to the RDBMS data base and web service targets. For more information, read *Specifying Target Database Processing Options*, in the *Event Replicator Target Adapter Administration Guide*

- 5 Use the Administration tool to configure the source definitions for your installation. Source definitions define the messaging system to be used (IBM EntireX or WebSphere MQ) and associated target definition to use which the Event Replicator Target Adapter should direct the replicated data acquired from the source definition. For more information, read *Configuring Source Definitions for Event Replicator Target Adapter*, in the *Event Replicator Target Adapter Administration Guide*.
- 6 Optionally, use the Administration tool to configure any filter definitions for Event Replicator Target Adapter processing. For more information, read *Specifying Filter Definitions*, in the *Event Replicator Target Adapter Administration Guide*
- 7 Start the Event Replicator Target Adapter (if it is not already started), as described in *Starting the Event Replicator Target Adapter*, in the *Event Replicator Target Adapter User Guide*.
- 8 Edit the Event Replicator subscriptions you want to use to transform and submit data to the Event Replicator Target Adapter. For each subscription, verify that:
 - A global format buffer (GFB) and field table (GFFT) are *generated* for at least one SFILE definition in the subscription. The GFB and field table must be generated either by Event Replicator for Adabas using Predict file definitions or using the Data Mapping Tool.
 - At least one destination definition used by the subscription specifies a value of "SAGTARG" for the Destination Class field (DCLASS parameter).



Note: DCLASS=SAGTARG cannot be specified for Adabas or file destinations (DTYPE=ADABAS or FILE). It is only valid for IBM EntireX, WebSphere MQ, or NULL destinations. When DCLASS=SAGTARG is specified, the ADARUN RPLPARMS parameter must be set to "FILE" or "BOTH" to provide access to any field table (GFFT) definitions.

Be sure to save the subscription definitions and any destination definitions you modified.

For more information on Event Replicator for Adabas subscription, SFILE, GFB, and destination definitions, read *Definition Descriptions* in *Event Replicator for Adabas Administration and Operations Guide*.

- 9 Activate the Event Replicator subscription and destination definitions. To activate subscription and destination definitions using Adabas Online System (AOS), read *Managing Replication Definitions from AOS* in *Event Replicator for Adabas Administration and Operations Guide*.

9

Best Practices

■ Improving Overall Performance	26
■ Processing Initial-State Data	26

This chapter describes best practices for using the Event Replicator Target Adapter.

Improving Overall Performance

To improve the performance of your Event Replicator Target Adapter installation consider the following suggestions.

1. Set the compression for data that comes from any IBM EntireX sources to zero (0). This turns off compression entirely. Decompressing the compressed data, especially large quantities of data, will slow replication processing down. For more information about setting this compression value, read *Defining a New IBM EntireX Source Definition*, in the *Event Replicator Target Adapter Administration Guide*.
2. Try to maintain a one-to-one relationship between the number of Event Replicator Target Adapter engine threads and the number of active source definitions.

For more information about maintaining the configuration engine definition, read *Maintaining the Event Replicator Target Adapter Engine Configuration*, in the *Event Replicator Target Adapter Administration Guide*.

3. Install the Event Replicator Target Adapter on a separate server from the relational database you will be updating.
4. Install the Event Replicator Target Adapter on a moderately powerful machine. The minimum requirements for this machine are described in *Space Requirements* and *Memory Requirements*, in the *Event Replicator Target Adapter Installation Guide*.
5. Install the Event Replicator Server, the Event Replicator Target Adapter, and the target relational database on the same network segment.
6. If your target database is running in a cloud environment (e.g. AWS), you will get better performance if you also run your Event Replicator Target Adapter in the same cloud service.

In addition, read the other sections in this Best Practices documentation for information on improving the performance of initial-state processing.

Processing Initial-State Data

■ Performance Tips

- When Initial-State Processing Fails While Using a Loader

Performance Tips

If possible, use one of the Loaders provided with the Event Replicator Target Adapter for initial-state processing. At this time, support is provided for the Oracle SQL *Loader, the Microsoft SQL Server Bulk Copy (bcp) Utility, Teradata MultiLoad Utility, the DB2 CONNECT IMPORT Command and the PostgreSQL psql Utility.

In addition, if you are using a IBM EntireX source, the initial-state process will run faster if the persistent store is turned off during the run. For more information on doing this, refer to your IBM EntireX documentation.

When Initial-State Processing Fails While Using a Loader

➤ If an initial-state process fails between the Event Replicator Server and the Event Replicator Target Adapter while using a Loader, complete the following steps:

- 1 Determine why the initial-state process failed and correct it. The log file or console messages will show errors after the initial-state process has started. No errors and the "Initial-State completed" message indicates that all is well.
- 2 Depending on the reason the initial-state process failed, you may need to drain the queue between the Event Replicator Server and the Event Replicator Target Adapter. To do this, you can use the *etbivp* or *mqsivp* tools, or any other methods for draining queues provided by IBM EntireX Broker or MQ Series.
- 3 Depending on the reason the initial-state process failed, you may need to drop the relational database (RDBMS) tables using the appropriate RDBMS tool. Only the tables created or used by the particular initial-state process will need to be dropped.
- 4 Restart the initial-state process.

Index

B

best practices, 25

C

character encoding, 19
conversions of data types, 17

D

data type conversions, 17

E

encoding, 19
Event Replicator Target Adapter
 character encoding considerations, 19
 data type conversions, 17
 examples, 21
 general processing, 11
 getting started, 23
 how it works, 9
 log file, 13
 MU and PE field support, 14-15
 names of SQL tables, 13
 overview, 5
 primary keys of SQL tables, 14
 processing flow, 12
 RDBMS recovery processing, 15
 requirements, 10
 sample, 21
 table structure, 13
 tasks, 5
examples, 21

G

getting started
 with Event Replicator Target Adapter, 23

I

improving performance, 26-27
initial-state processing
 performance tips, 27

M

MU field support, 14-15

P

PE field support, 14-15
performance tips
 initial-state processing, 27
 overall, 26
processing overview, 11

R

RDBMS
 recovery processing, 15
 table structure, 13
recovery processing, 15
requirements, 10

S

samples, 21
SQL tables
 names of, 13
 primary key, 14
sqlrep.log file, 13

T

table structure, 13

