Event Replicator for Adabas

Using the Event Replicator Target Adapter

Version 4.1.0

September 2026

ADABAS & NATURAL

This document applies to Event Replicator Target Adapter Version 4.1.0 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 © 2026 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-AARTUSER-410-20260925

Table of Contents

Preface v
1 About this Documentation 1
Document Conventions 2
Online Information and Support 2
Data Protection 3
2 Starting the Event Replicator Target Adapter 5
Starting Event Replicator Target Adapter from the Administration Tool 6
Restarting the Event Replicator Target Adapter 6
Starting Event Replicator Target Adapter in Windows 7
Starting Event Replicator Target Adapter in UNIX 7
3 Activating Event Replicator Target Adapter Processing 9
4 Deactivating Event Replicator Target Adapter Processing 13
5 Shutting Down the Event Replicator Target Adapter 15
Stopping Event Replicator Target Adapter from the Administration Tool 16
Shutting Down Event Replicator Target Adapter in Windows 16
Shutting Down Event Replicator Target Adapter in UNIX 17
6 Requesting Initial-State Data 19
7 Replicating Data to JMS Queues or Topics 21
Requirements 23
Implementation Steps 23
XML Message Structure for JMS 24
Managing the JMS User Target Plugin 29
8 Replicating Data to Adabas 43
Requirements 44
Adabas-Specific Notes 45
9 Replicating Data to Adabas in Binary Mode 47
Requirements 48
Adabas-Specific Notes 48
10 Replicating Data to Terracotta Caches 51
Requirements 52
Implementation Steps 53
Terracotta-specific Notes 54
11 Replicating Data to Terracotta 57
Requirements 58
Terracotta-specific Notes 59
12 Using the DB2 CONNECT IMPORT Command to Load Initial-State Data 61
Prerequisites 62
Limitations 62
Enabling DB2 Database Initial-State Data Loading 62
Processing Output 64
13 Using the Microsoft SQL Server Bulk Copy (bcp) Utility to Load Initial-State Data 67
Prerequisites 68

Limitations 68
Enabling bcp Utility Initial-State Processing 68
Processing Output 70
14 Using the Oracle SQL*Loader to Load Initial-State Data 73
Prerequisites 74
Enabling Oracle SQL*Loader Initial-State Processing 74
Processing Output 75
15 Using the PostgreSQL psql Utility to Load Initial-State Data 77
Prerequisites 78
Enabling PostgreSQL SQL Initial-State Processing using psql 78
Processing Output 79
16 Using the Teradata MultiLoad (MLoad) Utility to Load Initial-State Data 81
Prerequisites 82
Limitations 83
Enabling Teradata MLoad Utility Initial-State Processing 83
Processing Output 84
17 Requesting RDBMS Data Refreshes 87
18 Dropping RDBMS Tables 89
19 Replicating ADALOD Data to the RDBMS Tables 91
20 Managing Event Replicator Target Adapter Log Files and Console Messages 93
Log File Location 94
Reviewing the sqlrep.log File 94
sqlrep.log File Retention 94
Setting Log Levels 95
21 Controlling Event Replicator for Adabas Destination Definitions 97
22 Samples 99
Generated Field Table (GFFT) 100
Schema and Replicated Data 101
Index 105

Preface

This document describes how to use the Event Replicator Target Adapter to transform and apply replicated Adabas mainframe data to a relational database, such as DB2, MySQL, Oracle, SQL Server, Sybase, or Teradata.

It covers the following topics:

Starting the Event Replicator Target Adapter	Describes how to start the Event Replicator Target Adapter in Windows and UNIX environments.
Activating Event Replicator Target Adapter Processing	Describes how to activate Event Replicator Target Adapter processing.
Deactivating Event Replicator Target Adapter Processing	Describes how to deactivate Event Replicator Target Adapter processing.
Shutting Down the Event Replicator Target Adapter	Describes how to shut down the Event Replicator Target Adapter in Windows and UNIX environments.
Requesting Initial-State Data	Describes how to request that initial-state data be applied to your RDBMS using the Adabas Event Replicator Subsystem.
Replicating Data to JMS Queues or Topics	Describes how to use the Event Replicator Target Adapter to replicate data to JMS queues and topics that are managed by webMethods Broker JNDI providers.
Replicating Data to Adabas	Describes how to use the Event Replicator Target Adapter to replicate data to Adabas.
Replicating Data to Adabas in Binary Mode	Describes how to use the Event Replicator Target Adapter to replicate data to Adabas using raw binary data.
Replicating Data to Terracotta Caches	Describes how to use the Event Replicator Target Adapter to replicate data to Terracotta in memory caches.
Replicating Data to TerracottaDB	Describes how to use the Event Replicator Target Adapter to replicate data to TerracottaDB.
Using the DB2 CONNECT IMPORT Command to Load Initial-State Data	Describes how to use the DB2 CONNECT IMPORT command to load initial-state data into a DB2 database.
Using the Microsoft SQL Server Bulk Copy (bcp) Utility to Load Initial-State Data	Describes how to use the Microsoft SQL Server bulk copy (bcp) utility to load initial-state data into a Microsoft SQL Server database.
*Using the Oracle SQL*Loader to Load Initial-State Data*	Describes how to request that initial-state data be applied to your Oracle database using the Oracle SQL*Loader.
Using the PostgreSQL psql Utility to Load Initial-State Data	Describes how to use the PostgreSQL psql utility to process and load initial-state data into a PostgreSQL database.
Using the Teradata MultiLoad (MLoad) Utility to Load Initial-State Data	Describes how to use the Teradata MultiLoad (MLoad) utility to load initial-state data into a Teradata database.
Requesting RDBMS Data Refreshes	Describes how to request RDBMS data refreshes.

Dropping RDBMS Tables	Describes how to drop RDBMS tables.
Replicating ADALOD Data to the RDBMS Tables	Describes how to replicate ADALOD data to your RDBMS tables.
Managing Event Replicator Target Adapter Log Files and Console Messages	Describes how to manage and review the Event Replicator Target Adapter log files and console messages.
Controlling Event Replicator for Adabas Destination Definitions	Describes batch utilities provided in the Event Replicator Target Adapter that you can use to activate, deactivate, open and close Event Replicator for Adabas destination definitions.
Samples	Provides a sample of Event Replicator Target Adapter use.

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 *`folder.subfolder.service`*, APIs, Java classes, methods, properties.
`Italic`	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.softwareag.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 Starting the Event Replicator Target Adapter

- Starting Event Replicator Target Adapter from the Administration Tool 6
- Restarting the Event Replicator Target Adapter 6
- Starting Event Replicator Target Adapter in Windows 7
- Starting Event Replicator Target Adapter in UNIX 7

This chapter describes how to start Event Replicator Target Adapter in Windows and in UNIX as well as how to start it from the Administration tool.

Notes:

1. The Event Replicator Target Adapter is automatically stopped and restarted every time an Event Replicator Target Adapter definition is saved or deleted using the Administration tool. For this reason, we recommend that you shut down the Event Replicator Target Adapter prior to making changes to the definitions and restart the Event Replicator Target Adapter when all definition changes are complete.
2. Event Replicator Target Adapter only supports replicating data to one relational database at a time.

Starting Event Replicator Target Adapter from the Administration Tool

Note: At least one source and one target definition must be configured before you can start the Event Replicator Target Adapter using the Administration tool.

To start Event Replicator Target Adapter from the Administration tool:

1 Start up the Administration tool, as described in *Starting the Administration Tool*, in the *Event Replicator Target Adapter Administration Guide*.

The Administration tool window appears.

2 Click on the **Start Target Adapter** button ( ) on the Administration tool window.

Or:

On the **Run** menu, select **Start Target Adapter**.

The Event Replicator Target Adapter starts up.

Restarting the Event Replicator Target Adapter

The Event Replicator Target Adapter can be restarted from the Administration tool.

To restart Event Replicator Target Adapter from the Administration tool once it is already started:

- Click on the **Restart Target Adapter** button ( ) on the Administration tool window.

Or:

On the **Run** menu, select **Restart Target Adapter**.

The Event Replicator Target Adapter stops and restarts.

Starting Event Replicator Target Adapter in Windows

To start Event Replicator Target Adapter using the *startupart* batch file in Windows environments:

1 Navigate to the following directory: *<installation-dir>/Event Replicator Target Adapter/Vv.r.s/bin*, where `<installation-dir>` is the name of the directory in which the Event Replicator Target Adapter program files were installed and `v.r.s` represents the current version, release, and system modification level of the Event Replicator Target Adapter.

   Or:

   From the Windows Start menu, select **All Programs**, then select **Software AG Event Replicator Target Adapter** `v.r.s` (where `v.r.s` represents the current version, release, and system modification level of the Event Replicator Target Adapter), and finally select **Event Replicator Target Adapter Command Prompt**. A command prompt window appears, already placed in the correct directory.

2 Run the *startupart.bat* file.

   The Event Replicator Target Adapter starts up.

   Even if there are no target, source, or target database option definitions defined for the Event Replicator Target Adapter, it will still start. If no definitions are defined, no processing will occur and no replication data can be received, but the Event Replicator Target Adapter will start.

Starting Event Replicator Target Adapter in UNIX

To start Event Replicator Target Adapter in UNIX:

1 Navigate to the following directory: *<installation-dir>\Event Replicator Target Adapter\Vvrs\bin*, where `<installation-dir>` is the name of the directory in which the Event Replicator Target Adapter program files were installed and `vrs` represents the current version, release, and system modification level of the Event Replicator Target Adapter.

2 Run the *startupart.sh* script at the UNIX prompt.

The Event Replicator Target Adapter starts up.

Even if there are no target, source, or target database option definitions defined for the Event Replicator Target Adapter, it will still start. If no definitions are defined, no processing will occur and no replication data can be received, but the Event Replicator Target Adapter will start.

3 Activating Event Replicator Target Adapter Processing

Once Event Replicator Target Adapter has been installed and your relational database has been set up properly, you can activate Event Replicator Target Adapter processing.

To activate Event Replicator Target Adapter processing:

1. Verify that you have proper authorization privileges to maintain the RDBMS tables.

   Effective with Event Replicator Target Adapter 2.7, 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. Verify that appropriate target definitions and target database processing option definitions have been created in the Event Replicator Target Adapter Administration tool. One target definition should be created for each RDBMS target to which you want data replicated. A target database processing option definition is only necessary if the settings in the default target database processing option definition are not sufficient. For more information, read *Configuring Target Definitions for Event Replicator Target Adapter* and *Specifying Target Database Processing Option Definitions*, in the *Event Replicator Target Adapter Administration Guide*.

3. Verify that appropriate Event Replicator Server queue definitions have been set up for the messaging system (IBM EntireX or WebSphere MQ) you are using. For more information about maintaining queue definitions, read the documentation in one of the following chapters:

   - To maintain queue definitions using DDKARTE statements, read *IQUEUE Parameter* in *Event Replicator for Adabas Reference Guide*.
   - To maintain queue definitions using the Adabas Event Replicator Subsystem, read *Maintaining Input Queue (destination) Definitions* in *Adabas Event Replicator Subsystem User's Guide*.

4 Verify that you have corresponding source definitions created in the Event Replicator Target Adapter Administration tool for every IQUEUE definition you need to use to replicate data from your mainframe Adabas to your relational database. For information on configuring source definitions in the Event Replicator Target Adapter Administration tool, read *Configuring Source Definitions for Event Replicator Target Adapter*, in the *Event Replicator Target Adapter Administration Guide*.

5 Verify that appropriate Event Replicator Server subscription definitions have been set up for data replication to your relational database. For more information about maintaining subscription definitions, read the documentation in one of the following chapters:

- To maintain subscription definitions using DDKARTE statements, read *SUBSCRIPTION Parameter* in *Event Replicator for Adabas Reference Guide*.
- To maintain subscription definitions using the Adabas Event Replicator Subsystem, read *Maintaining Subscription Definitions* in *Adabas Event Replicator Subsystem User's Guide*.

6 Verify that a global format buffer and field table (GFFT) have been *generated* for the subscription definition. To use Event Replicator Target Adapter, you must use generated global format buffers. To generate a global format buffer and field table (GFFT) using the Adabas Event Replicator Subsystem, read *Generating a GFB* in *Adabas Event Replicator Subsystem User's 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 Save the subscription definitions.

8 Edit the destination definitions used by the subscriptions you modified in the previous steps.

9 Verify that the destination class (DCLASS parameter) is specified as "SAGTARG" for each destination.

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.

Optionally, specify keyword "NOSPRE" for the destination class parameter (DCLASSPARM parameter) if you do not want the subscription name to prefix the names of the tables produced by the Event Replicator Target Adapter. When "NOSPRE" is specified, the schema file name (Predict view name) alone is used for the table names; when "NOSPRE" is *not* specified, the subscription name prefixes the Predict view name in the table names.

For more information about specifying the destination class in a destination definition, read the documentation in one of the following chapters:

- To modify the destination definition using DDKARTE statements, read *DESTINATION Parameter* in *Event Replicator for Adabas Reference Guide*.
- To modify the destination definition using the Adabas Event Replicator Subsystem, read *Maintaining Destination Definitions* in *Adabas Event Replicator Subsystem User's Guide*.

10 Save the destination definitions.

11 Stop and restart the Event Replicator Server or refresh the destination and subscription definitions, as described in *RPLREFRESH Command*, in *Event Replicator for Adabas Administration and Operations Guide*.

12 Activate the destination definitions and the subscription definitions you modified in the previous steps. For more information, read one or more of the following chapters:

- To activate destination and subscription definitions using Adabas Online System (AOS), read *Managing Replication Definitions from AOS* in *Adabas Event Replicator Subsystem User's Guide*.
- To activate destination definitions using an Event Replicator Target Adapter batch utility, read ***Controlling Event Replicator for Adabas Destination Definitions***, elsewhere in this chapter.

If the subscription definitions are already activated, deactivate them first and then reactivate them to pick up the changes made in these steps.

13 Verify that the message systems referenced by your Event Replicator Server queue definitions and your Event Replicator Target Adapter source definitions are started or that their associated source definitions in the Administration tool have **Autostart** selected (checked).

14 If it is not already started, start the Event Replicator Target Adapter.

Event Replicator Target Adapter processing is activated and replicated Adabas data will be transformed and applied to your RDBMS.

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.

4 Deactivating Event Replicator Target Adapter Processing

To deactivate Event Replicator Target Adapter processing:

- Deactivate the subscription definitions that manage the replicated data you are transforming and applying to your relational database. To deactivate subscription definitions using Adabas Online System (AOS), read *Managing Replication Definitions from AOS* in *Adabas Event Replicator Subsystem User's Guide*.

  Event Replicator Target Adapter processing is deactivated.

5 Shutting Down the Event Replicator Target Adapter

- Stopping Event Replicator Target Adapter from the Administration Tool 16
- Shutting Down Event Replicator Target Adapter in Windows 16
- Shutting Down Event Replicator Target Adapter in UNIX 17

This chapter describes how to shut down the Event Replicator Target Adapter in Windows and UNIX environments as how to stop it from the Administration tool.

Note: The Event Replicator Target Adapter is automatically stopped and restarted every time an Event Replicator Target Adapter definition is saved or deleted using the Administration tool. For this reason, we recommend that you shut down the Event Replicator Target Adapter prior to making changes to the definitions and restart the Event Replicator Target Adapter when all definition changes are complete.

Stopping Event Replicator Target Adapter from the Administration Tool

To shut down Event Replicator Target Adapter from the Administration tool:

1 Start up the Administration tool, as described in *Starting the Administration Tool*, in the *Event Replicator Target Adapter Administration Guide*.

The Administration tool window appears.

2 Click on the **Stop Target Adapter** button ( ) on the Administration tool window.

Or:

On the **Run** menu, select **Stop Target Adapter**.

The Event Replicator Target Adapter shuts down.

Shutting Down Event Replicator Target Adapter in Windows

To shut down Event Replicator Target Adapter using the *shutdownart* batch file in Windows:

1 Navigate to the following directory: *<installation-dir>/Event Replicator Target Adapter/Vv.r.s/bin*, where `<installation-dir>` is the name of the directory in which the Event Replicator Target Adapter program files were installed and `v.r.s` represents the current version, release, and system modification level of the Event Replicator Target Adapter.

2 Run the *shutdownart.bat* file.

The Event Replicator Target Adapter shuts down.

Shutting Down Event Replicator Target Adapter in UNIX

To shut down Event Replicator Target Adapter in UNIX environments:

1. Navigate to the following directory: *<installation-dir>\Event Replicator Target Adapter\Vvrs\bin*, where `<installation-dir>`is the name of the directory in which the Event Replicator Target Adapter program files were installed and `vrs` represents the current version, release, and system modification level of the Event Replicator Target Adapter.
2. Enter `shutdownart.sh` at the UNIX prompt.

   The Event Replicator Target Adapter shuts down.

6 Requesting Initial-State Data

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 relational database tables. For information on using the Adabas Event Replicator Subsystem to do this, read *Populating the RDBMS Tables* in *Adabas Event Replicator Subsystem User's Guide*.

7 Replicating Data to JMS Queues or Topics

- Requirements 23
- Implementation Steps 23
- XML Message Structure for JMS 24
- Managing the JMS User Target Plugin 29

You can use the Event Replicator Target Adapter to replicate data to JMS queues and topics that are managed by webMethods Broker JNDI providers. Replicate the data to a JMS topic if you want to use publish-subscribe messaging; replicate the data to a JMS queue if you want to use point-to-point messaging. Once the data is published on a topic or queue, message consumers, such as webMethods Integration Server, can retrieve or consume it.

The following diagram depicts the interaction between the Event Replicator Target Adapter and the JNDI providers managed by webMethods Broker.

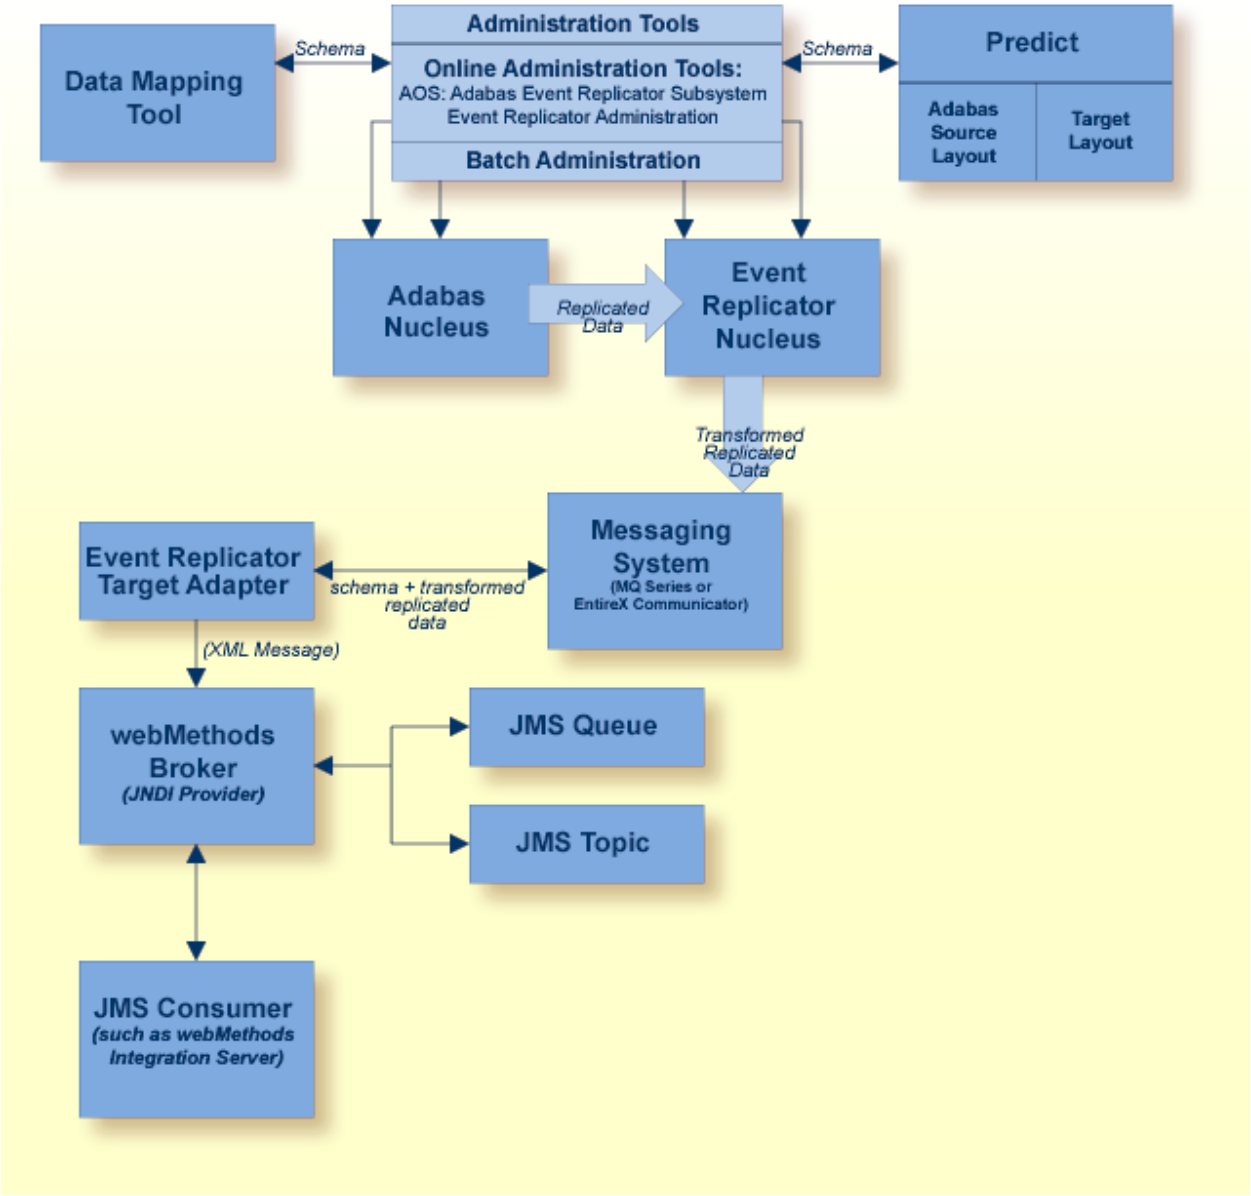

This chapter describes how to set up this JMS support in the Event Replicator Target Adapter and the high-level steps that must be taken in webMethods Broker.

Requirements

webMethods Broker and My webMethods Version 7.1 or later must be installed. In addition, JMS support must be established within webMethods Broker. For complete information, read *EntireX Message Service Provider (MSP)*, available in the IBM EntireX documentation on Software GmbH's **Empower** web site.

Implementation Steps

To set up replication of Adabas data to a JMS topic or queue, complete the following steps:

1 Set up JMS support in webMethods Broker specific for use by the Event Replicator Target Adapter. Be sure to perform the following steps in webMethods Broker:

- Establish a JNDI (Java Naming and Directory Interface) provider for use by the Event Replicator Target Adapter. The root location of the JNDI is called the *initial context*. The JNDI serves as a template for the JMS administered objects (connection factories and destinations), storing their property settings in specified locations relative to the initial context. These settings are applied when the JMS connection factories and destinations are created and used when the JMS application runs.
- Add topic and queue connection factories that configure the connections between JMS programs and webMethods Broker Server.
- Bind and create topic and queue destinations in the JNDI.
- As necessary, update client group access permissions to ensure they have permission to publish and subscribe to the document types used in any durable subscriptions for the JNDI's topics.
- As necessary, create durable subscriptions, linking specific document types to a topic. This is necessary for subscriptions whose data is persisted if the subscriber is disconnected.

For complete information, read *EntireX Message Service Provider (MSP)*, available in the IBM EntireX documentation on Software GmbH's **Empower** web site.

2 Collect the following information about the JMS support in webMethods Broker:

- The URL of the initial context of the JNDI provider established for use by the Event Replicator Target Adapter.
- The lookup name of the JMS topic or queue destination (as defined to webMethods Broker) that should be used by the Event Replicator Target Adapter.
- The name of the JMS topic or queue connection factory (as defined to webMethods Broker) that should be used by the Event Replicator Target Adapter.

3 Using the Event Replicator Target Adapter Administration tool, add a JMS target definition and save it. For more information, read *Maintaining JMS Target Definitions*, in the *Event Replicator Target Adapter Administration Guide*.

4 Start the Event Replicator Target Adapter (if it is not already started), as described in ***Starting the Event Replicator Target Adapter***, elsewhere in this guide.

5 Using the Adabas Event Replicator Subsystem (SYSRPTR), edit the Event Replicator subscriptions you want used 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 GFFT 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.

6 Activate the Event Replicator subscription and destination definitions. For more information, refer to your Event Replicator for Adabas documentation.

XML Message Structure for JMS

This section describes the structure of the XML messages containing replicated data that the Event Replicator Target Adapter publishes on JMS topics and queues. This structure is defined by an XML schema which is included with your Event Replicator Target Adapter.

This section is organized into the following topics:

- XML Schema
- XML Elements
- XML Attributes

- XML Message Examples

XML Schema

The schema defining the JMS message structure is provided with the Event Replicator Target Adapter. It is called *jms.xsd* and can be found in the following locations:

Platform	Directory
Windows Vista	**C:\Program Files\Software AG\Event Replicator Target Adapter\v***v.r.s***\bin**
Other Windows platforms	**C:\Program Files\Software AG\Event Replicator Target Adapter\v***v.r.s***\bin**
UNIX	**$SAG/art/v***vrs***/bin**

XML Elements

The following XML elements are used in the XML messages sent to JMS topics and queues.

Element Name	Hierarchy Level	Attributes	Subelements	Description

`<artdocument/>`	1	version	`<method/>`
			`<tableName/>`
			`<fields/>`
				Identifies the message as an Event Replicator Target Adapter message containing replicated data from an Adabas database.

`<field/>`	3	length name pre type value	-	Identifies an individual field (column) in the table that is affected by the data or table change and provides the details about that change. The attributes that appear in the `<field/>` element vary, based on the setting of the `<method/>` element.
`<fields/>`	2	-	`<field/>`	Surrounds the set of fields affected by the replicated data or table change.
`<method/>`	2	-	-	Identifies the kind of change applied to the database. Valid values are: ■ "create": Indicates that a new table is being created. ■ "delete": Indicates that a field (column) in a table is being deleted. ■ "drop": Indicates that a table is being dropped. ■ "insert": Indicates that a new field (column) is being inserted into a table. ■ "update": Indicates that the value of a field (column) is being updated.

Element Name	Hierarchy Level	Attributes	Subelements	Description
`<tableName/>`	2	-	-	Identifies the name of the table affected by the replicated data or table change.

XML Attributes

The following XML attributes are used in the XML messages sent to JMS topics and queues.

Attribute Name	Element Used In	Prerequisite	Description
length	`<field/>`	`<method/>` element must be "create"	Identifies the length of a field being created.
name	`<field/>`	always required	The name of the field (column). This is the name passed to the Event Replicator Target Adapter in the generated GFB and GFFT used for Event Replicator Target Adapter processing. This name is assigned to a database field when the Adabas database is initially set up. It can be renamed using the Data Mapping Tool before the Data Mapping Tool is used to generate the GFB and associated GFFT.
pre	`<field/>`	`<method/>` element must be "create"	Identifies the precision of a field being created. Alphanumeric fields have a precision of "0" (zero).
type	`<field/>`	`<method/>` element must be "create"	Identifies the data type of a field being created.
value	`<field/>`	`<method/>` element must be "insert", "update", or "delete"	The value of the field being inserted or updated.
version	`<artdocument/>`	-	The version of the Event Replicator Target Adapter schema used to generate this XML message.

XML Message Examples

This section includes examples of all possible types of messages submitted by the Event Replicator Target Adapter to JMS queues and topics.

- Table Created
- Table Dropped
- Field (Column) Inserted
- Field (Column) Value Updated

- Field (Column) Deleted

Table Created

The following is a sample XML message produced when a table is created. Note that the `<method/>` element is set to "create.".

```
<artdocument version="1">
   <method>create</method>
   <tableName>empldb31fnr5</tableName>
   <fields>
      <field name="ISN" type="integer" length="30" pre="0" />
      <field name="PERSONNEL-ID" type="string" length="8" pre="0" />
      <field name="FIRST-NAME" type="string" length="20" pre="0" />
      <field name="MIDDLE-I" type="string" length="1" pre="0" />
      <field name="NAME" type="string" length="20" pre="0" />
      <field name="MIDDLE-NAME" type="string" length="20" pre="0" />
      <field name="MAR-STAT" type="string" length="1" pre="0" />
      <field name="SEX" type="string" length="1" pre="0" />
      <field name="BIRTH" type="date" length="10" pre="0" />
      <field name="CITY" type="string" length="20" pre="0" />
      <field name="ZIP" type="string" length="10" pre="0" />
      <field name="POST-CODE" type="string" length="10" pre="0" />
      <field name="COUNTRY" type="string" length="3" pre="0" />
      <field name="AREA-CODE" type="string" length="6" pre="0" />
      <field name="PHONE" type="string" length="15" pre="0" />
      <field name="DEPT" type="string" length="6" pre="0" />
      <field name="JOB-TITLE" type="string" length="25" pre="0" />
      <field name="LEAVE-DUE" type="decimal" length="2" pre="0" />
      <field name="LEAVE-TAKEN" type="decimal" length="2" pre="0" />
   </fields>
</artdocument>
```

Table Dropped

The following is a sample XML message produced when a table is dropped. Note that the `<method/>` element is set to "drop".

```
<artdocument version="1">
   <method>drop</method>
   <tableName>empldb31fnr5</tableName>
</artdocument>
```

Field (Column) Inserted

The following is a sample XML message produced when a field (column) is inserted into a table. Note that the `<method/>` element is set to "insert".

```
<artdocument version="1">
   <method>insert</method>
   <tableName>empldb31fnr5_INCOME_BONUS</tableName>
   <fields>
      <field name="ISN_INCOME_INDEX" value="1_1" />
      <field name="BONUS" value="23000" />
      <field name="BONUS_IX2" value="1" />
   </fields>
</artdocument>
```

Field (Column) Value Updated

The following is a sample XML message produced when a field (column) value is updated in a table. Note that the `<method/>` element is set to "update".

```
<artdocument version="1">
   <method>update</method>
   <tableName>empldb31fnr5_INCOME_BONUS</tableName>
   <fields>
      <field name="ISN_INCOME_INDEX" value="1_1" />
      <field name="BONUS" value="230" />
      <field name="BONUS_IX2" value="1" />
   </fields>
</artdocument>
```

Field (Column) Deleted

The following is a sample XML message produced when a field (column) is deleted from a table. Note that the `<method/>` element is set to "delete".

```
<artdocument version="1">
   <method>delete</method>
   <tableName>empldb31fnr5_INCOME_BONUS</tableName>
   <fields>
      <field name="ISN_INCOME_INDEX" value="1_1" />
      <field name="BONUS" value="23000" />
      <field name="BONUS_IX2" value="1" />
   </fields>
</artdocument>
```

Managing the JMS User Target Plugin

The JMS User Target Plugin enables the Adabas Event Replicator Target Adapter (ART) to publish replicated data changes to systems that support the standard Java Message Service (JMS) interface.

The plugin works with JMS-compliant implementations such as Apache ActiveMQ Artemis and other JMS brokers.

- Prerequisites
- Installing the JMS User Target Plugin
- Enabling File Browser Access
- Registering the Plugin
- Configuring Source Definitions
- Creating a JMS Target
- JMS Communication Properties File
- Execution Flow
- Common Pitfalls
- Best Practices
- Optional Configuration Parameters
- Finalizing Configuration
- Usage
- Notes and Best Practices
- Migrating User Targets to Adabas Manager

Prerequisites

Before deploying the JMS User Target Plugin, ensure that the following prerequisites are met:

- Obtain the JMS client JAR for your JMS provider, for example, *artemis-jms-client-all-2.44.0.jar* for Apache ActiveMQ Artemis.
- Copy the JMS client JAR into the following directory:

```
<sag_install_dir>/EventReplicatorTargetAdapter/art/TargetAdapter/webapps/sqlrep/WEB-INF/lib/
```

Note: The plugin requires a JMS implementation that fully supports the standard Java JMS API.

Installing the JMS User Target Plugin

To install the JMS User Target Plugin:

1. Place the plugin JAR in a directory accessible by the Target Adapter.
2. Log in to Adabas Manager.

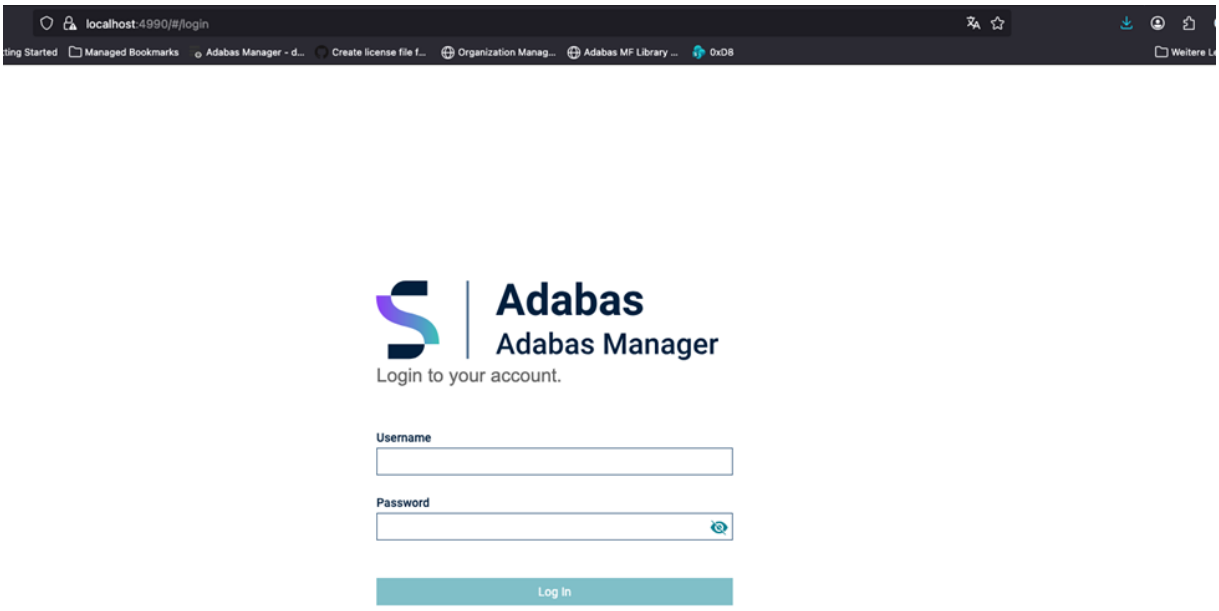

3. Navigate to **Event Replicator Target Adapter** and select **User Targets**.

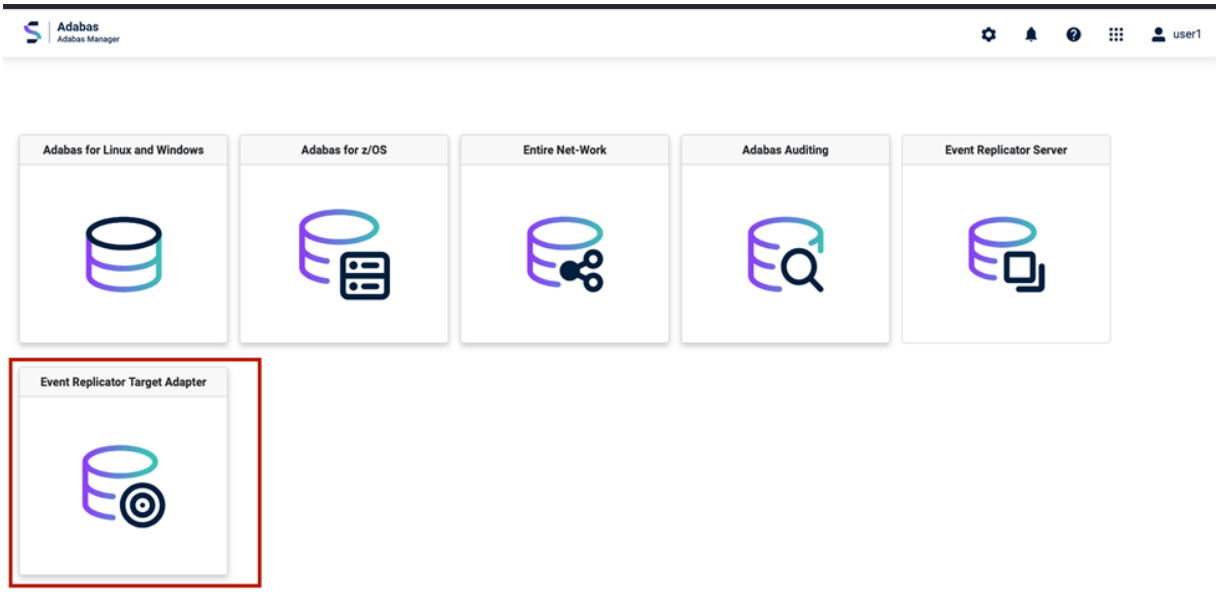

4. Click **Add User Target**.

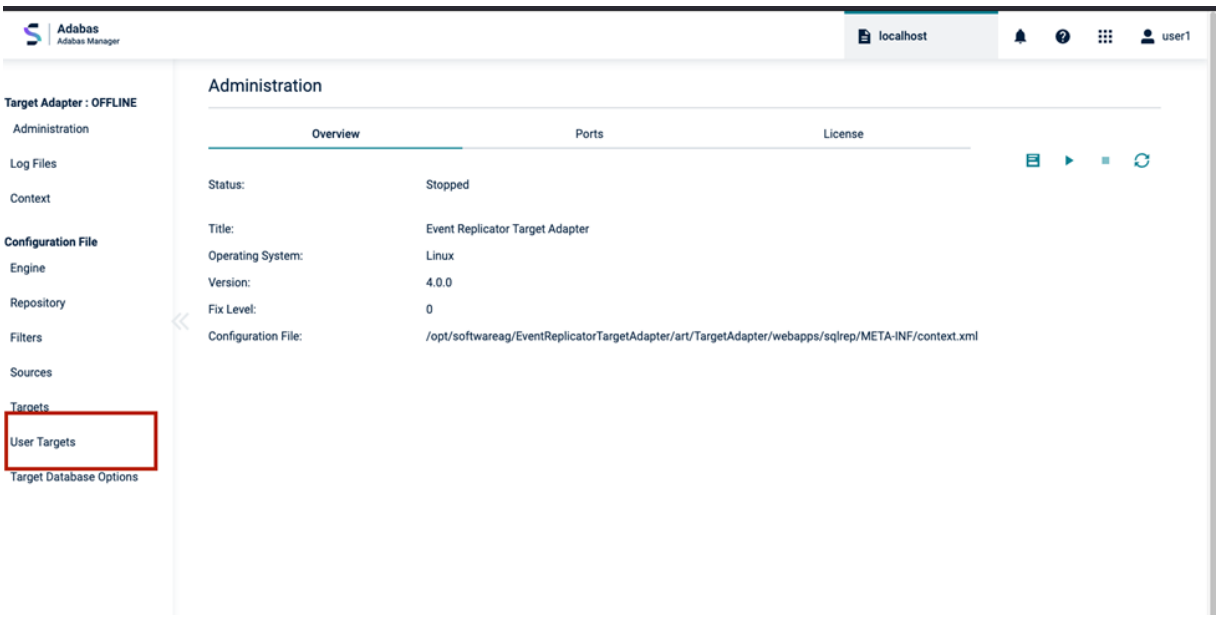

Enabling File Browser Access

To use the Browse button in the user interface, configure the directories that users are allowed to browse.

Edit the following file:

```
/EventReplicatorTargetAdapter/art/program/browsable-dirs.json
```

Add the required directories to the JSON array:

```
[
"/home/saguser",
"/home/testuser"
]
↩
```

Note: Restart the ART REST server after making changes for them to take effect.

Registering the Plugin

You can register the plugin using one of the following methods.

Option A - Browse

Use the Browse button in the Register User Targets form to navigate to and select the plugin JAR file from the file system.

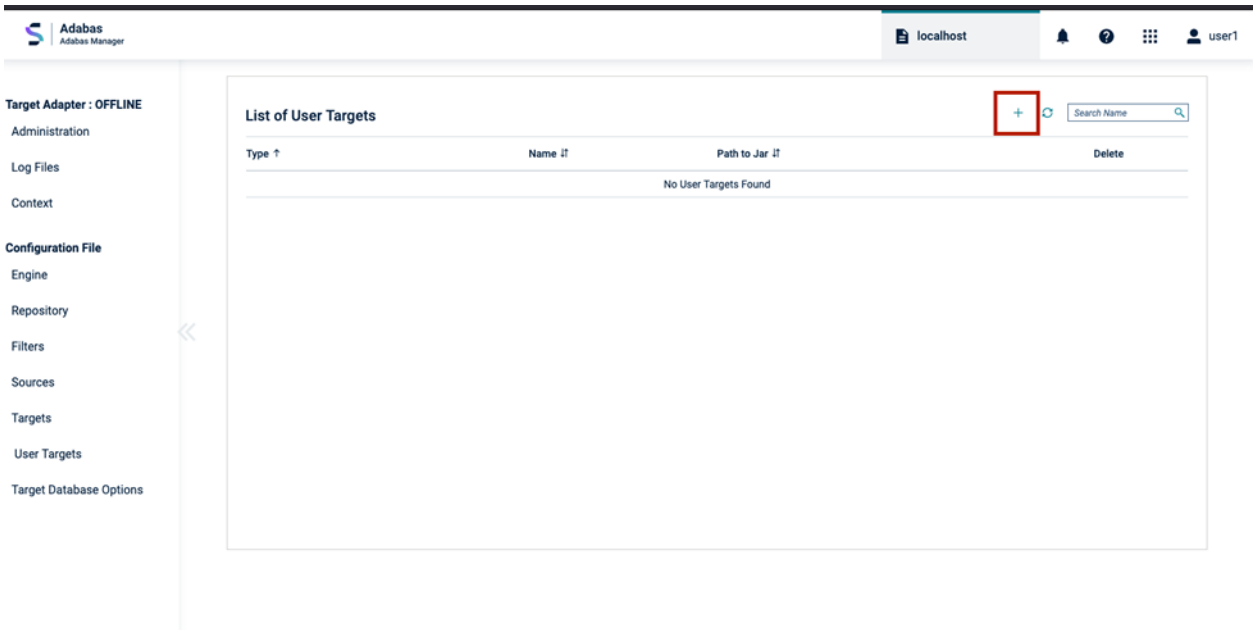

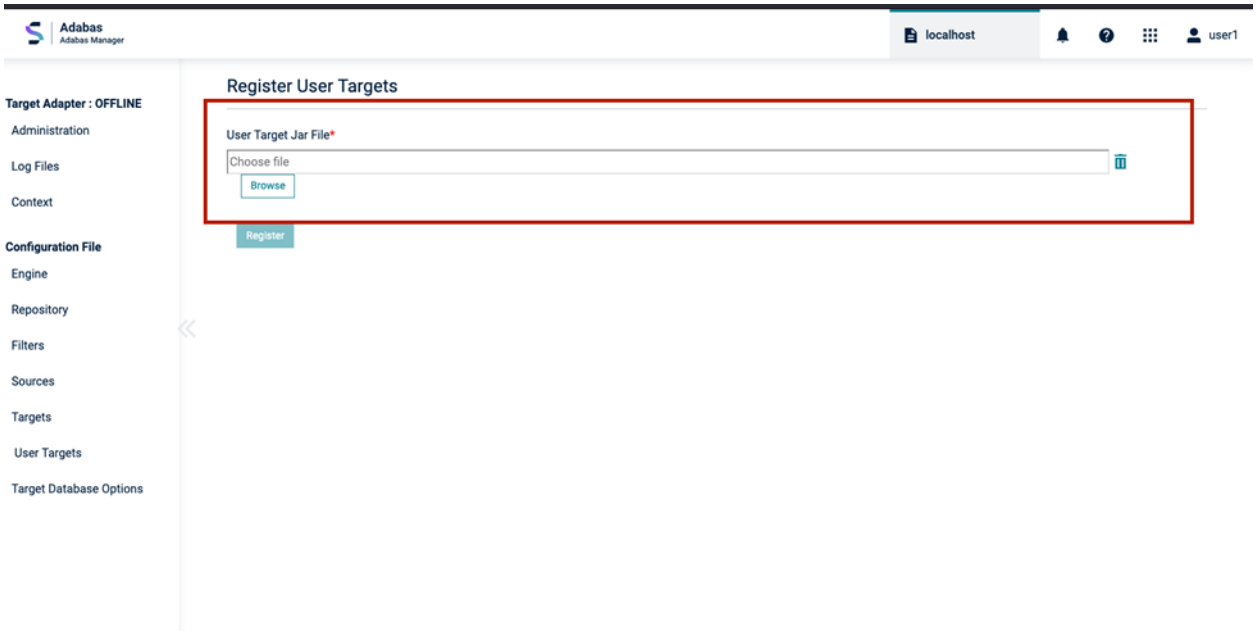
Adabas
Adabas Manager
localhost
user1
Target Adapter : OFFLINE
Administration
Log Files
Context
Configuration File
Engine
Repository
Filters
Sources
Targets
User Targets
Target Database Options
Register User Targets
User Target Jar File*
Choose file
Browse
Register

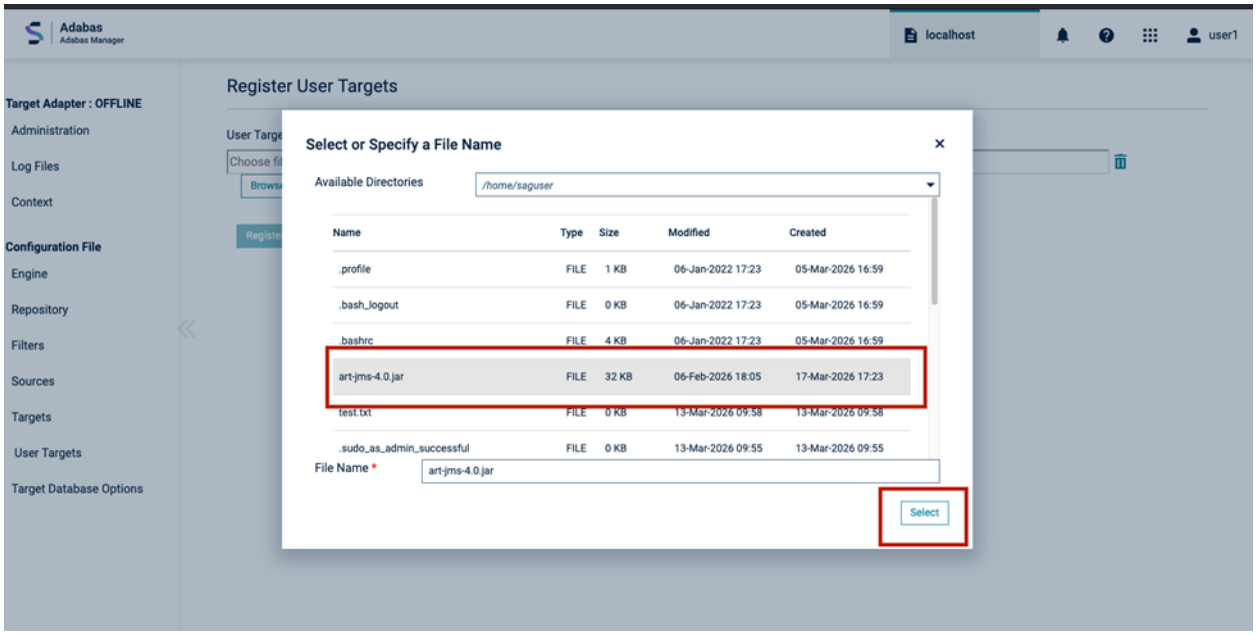
Adabas
Adabas Manager
localhost
user1
Target Adapter : OFFLINE
Administration
Log Files
Context
Configuration File
Engine
Repository
Filters
Sources
Targets
User Targets
Target Database Options
Register User Targets
Select or Specify a File Name
Available Directories
/home/saguser
Name
Type
Size
Modified
Created
.profile
FILE
1 KB
06-Jan-2022 17:23
05-Mar-2026 16:59
.bash_logout
FILE
0 KB
06-Jan-2022 17:23
05-Mar-2026 16:59
.bashrc
FILE
4 KB
06-Jan-2022 17:23
05-Mar-2026 16:59
art-jms-4.0.jar
FILE
32 KB
06-Feb-2026 18:05
17-Mar-2026 17:23
test.txt
FILE
0 KB
13-Mar-2026 09:58
13-Mar-2026 09:58
.sudo_as_admin_successful
FILE
0 KB
13-Mar-2026 09:55
13-Mar-2026 09:55
File Name *
art-jms-4.0.jar
Select

Option B - Manual Path Entry

Manually enter the full file system path to the plugin JAR in the User Target Jar File field, then click **Register** to deploy the plugin.

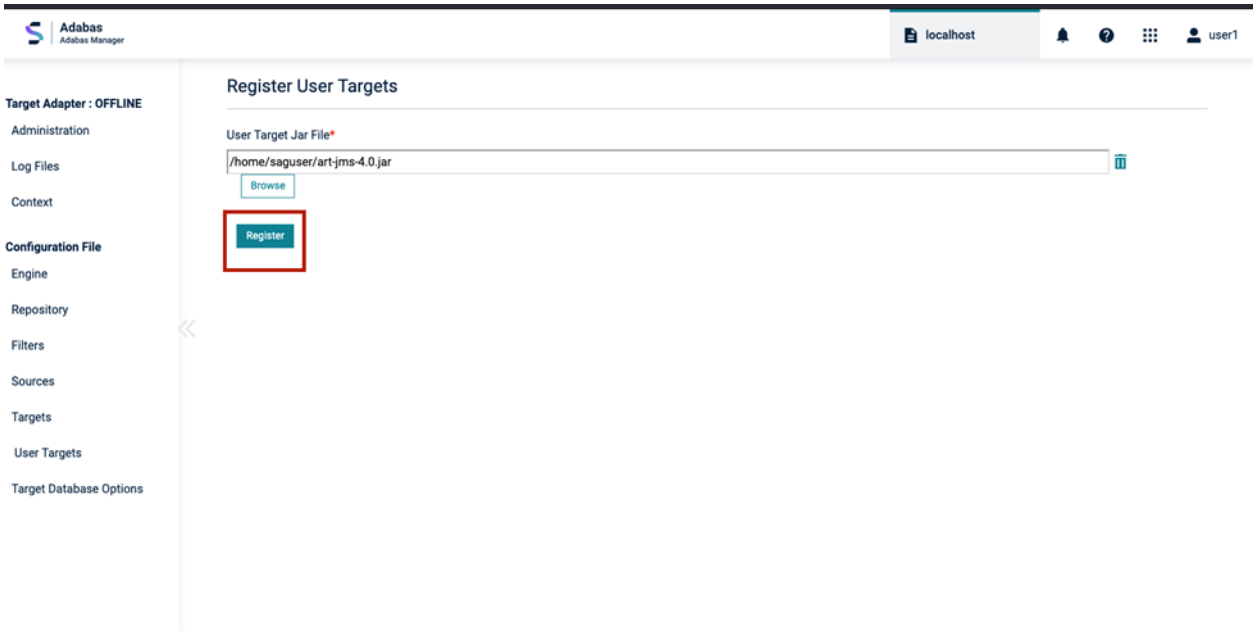

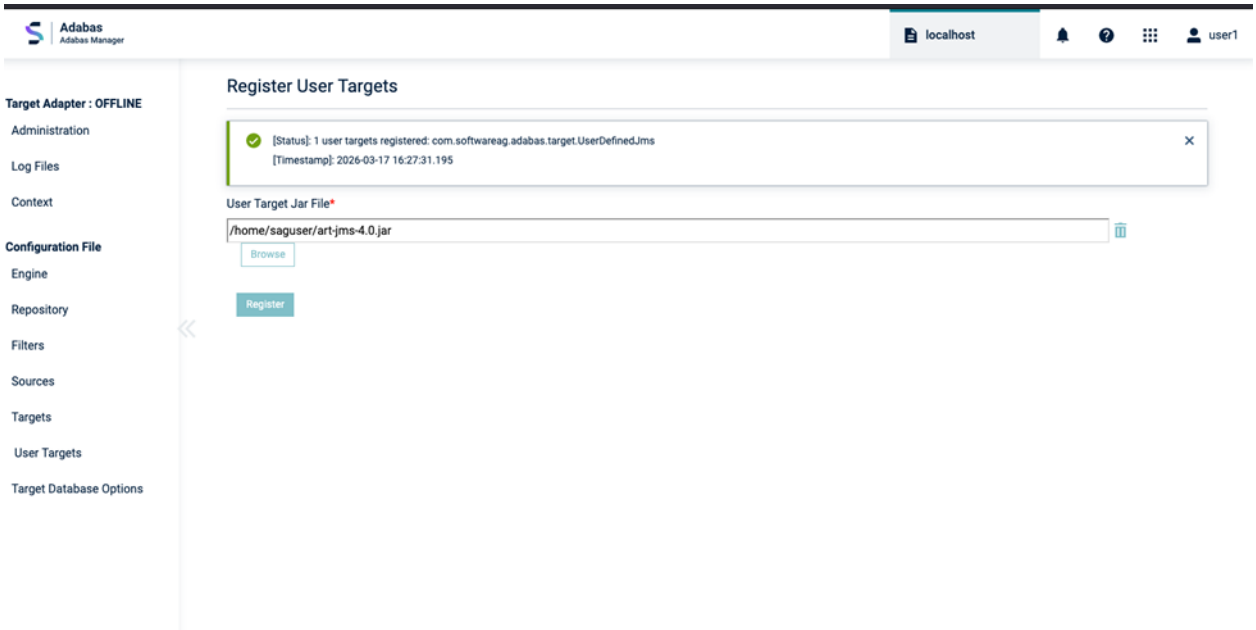

Configuring Source Definitions

For information about configuring Event Replicator Target Adapter source definitions, refer to the **Configuring Source Definitions for Event Replicator Target Adapter documentation.**

Creating a JMS Target

To create a JMS target:

1. Navigate to Targets and create a new target.

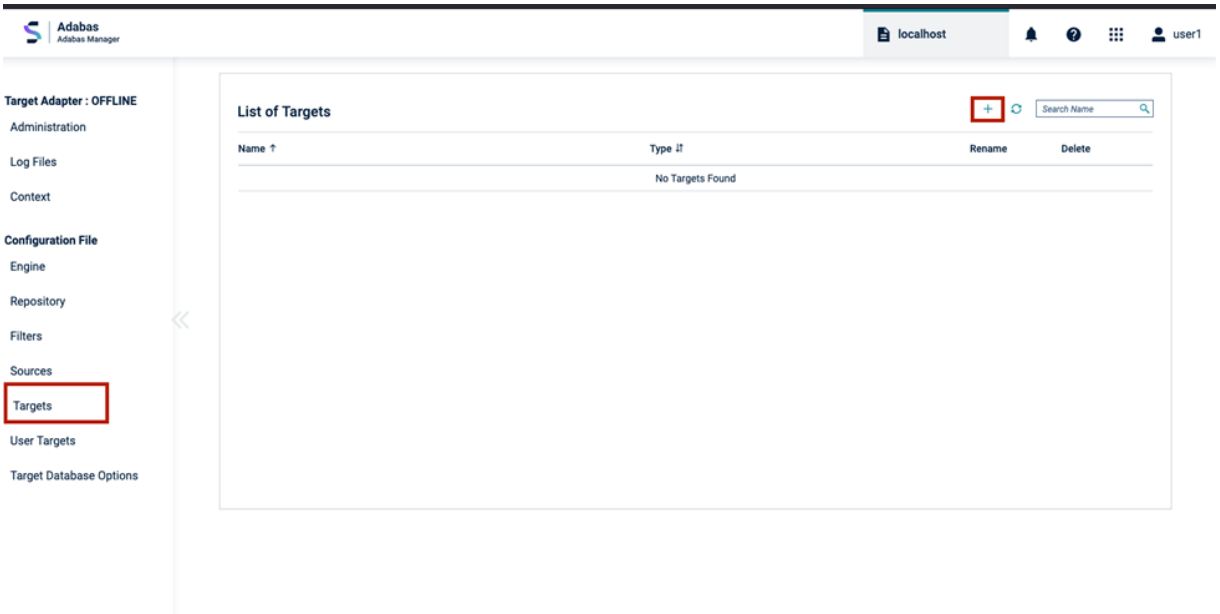

2. Select "UserDefinedJms" as the target type.

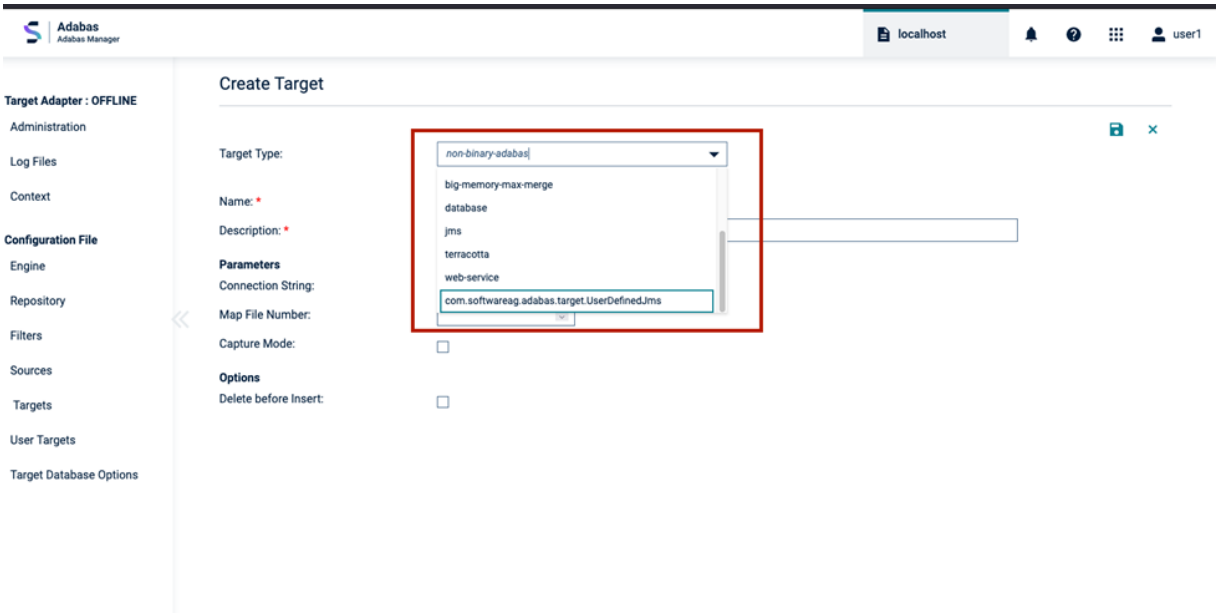

3. Specify the required values, including the name, description, and JMS communication properties file path.

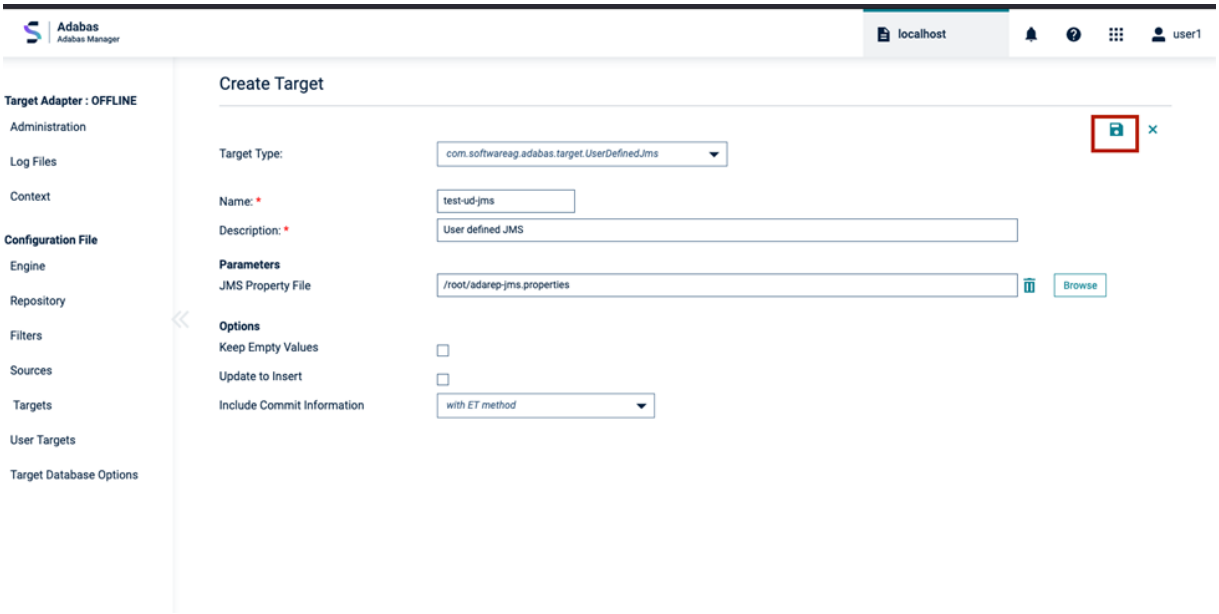

JMS Communication Properties File

The JMS connection is defined through a properties file. The properties file specifies the connection details, authentication settings, and JMS resources used by the JMS provider.

- Property Definitions
- Example Properties File

Property Definitions

`java.naming.factory.initial` specifies the fully qualified class name of the JNDI Initial Context Factory.

`java.naming.provider.url` defines the network endpoint of the JMS broker.

`java.naming.security.principal` specifies the user name used to authenticate with the JMS provider.

`java.naming.security.credentials` specifies the password used for authentication.

`connectionFactory` defines the JNDI name of the connection factory.

`queue.queue/<queueName>` defines a JMS queue and binds it to a JNDI name.

Example Properties File

```
              ↩
java.naming.factory.initial=org.apache.activemq.artemis.jndi.ActiveMQInitialContextFactory

               java.naming.provider.url=tcp://172.17.0.3:61616
               java.naming.security.principal=artemis
               java.naming.security.credentials=artemis
               connectionFactory=ConnectionFactory
               queue.queue/ordersQueue=ordersQueue
             ↩
```

Execution Flow

When the JMS User Target is active, the following processing sequence occurs:

1. Initialize the JNDI context.
2. Authenticate with the JMS provider.
3. Look up the connection factory and queue.
4. Establish the connection and session.
5. Publish replication messages to the configured queue.

Common Pitfalls

- Connection Failure - Verify that the provider URL is correct and that the broker is accessible.
- Class Not Found - Ensure that the JMS client JAR is placed in the correct directory.
- Authentication Error - Verify that the user name and password match the broker configuration.
- Queue Not Found - Verify that the JNDI name mapping in the properties file is correct.
- Silent Message Loss - Verify that the queue name matches the broker configuration exactly.

Best Practices

- Use consistent and descriptive queue naming conventions.
- Externalize credentials using environment variables or a secrets manager - avoid plaintext in production.
- Validate broker connectivity before deployment.
- Enable ART logging to assist with troubleshooting connection or delivery issues.
- Version-control all configuration files for auditability.
- Ensure version compatibility between the JMS client library and your broker.

Optional Configuration Parameters

Option	Description
Keep Empty Values	Select this option to indicate that all values of periodic (PE) fields or multiple-value (MU) fields within PE fields are replicated, even if they are empty. If this option is not selected, only non-empty values of PE fields or MU fields within PE fields are replicated. This option is available only for JMS targets when the Event Replicator runs with OPTIONS=32 on the mainframe.
Update to Insert	Select this option to convert update commands to insert commands. When enabled, the Event Replicator Target Adapter first deletes the existing row, if present, and then inserts a new row containing the changed information instead of issuing an update command. To use this feature, set the mainframe replication OPTIONS parameter specified in the destination class parameter data (DCLASSPARM parameter) to 32 (OPTIONS=32). **Caution:** Specifying OPTIONS=32 results in the entire Adabas record image being sent to the Event Replicator Target Adapter and can lead to degradation of overall throughput and performance. Default: Leave this option unselected.

Option	Description
	Note: This setting is ignored when a loader processes initial-state data.
Include Commit Information	Select an option from the drop-down menu to control whether and how Adabas commit time information is included in the replicated data for all processed transactions. with ET method - Produces output with a separate ET document. yes - Provides commit information inline with the record document. no - No commit information is provided. Default: no.

Include Commit Information Examples

The following examples show how commit information is included when different options are selected.

Option "with ET method"

This option produces output with a separate ET document.

```
<?xml version="1.0" encoding="UTF-8"?>
<artdocument version="1" committed="2018/10/04-12:46:41.284060">
<method>update</method>
<tableName>EMPLOYEES</tableName>
<fields>
<field name="ISN" value="1130" />
<field name="CITY" value="Auckland2" />
</fields>
</artdocument>

<?xml version="1.0" encoding="UTF-8"?>
<artdocument version="1" committed="2018/10/04-12:46:41.284060">
<method>ET</method>
</artdocument>
↩
```

Option "yes"

This option provides commit information inline with the record document.

```
<?xml version="1.0" encoding="UTF-8"?>
<artdocument version="1" committed="2018/10/04-12:44:42.834695">
<method>update</method>
<tableName>EMPLOYEES</tableName>
<fields>
<field name="ISN" value="1130" />
<field name="CITY" value="Auckland" />
</fields>
</artdocument>
↩
```

Finalizing Configuration

1. Click Save to persist the target configuration.

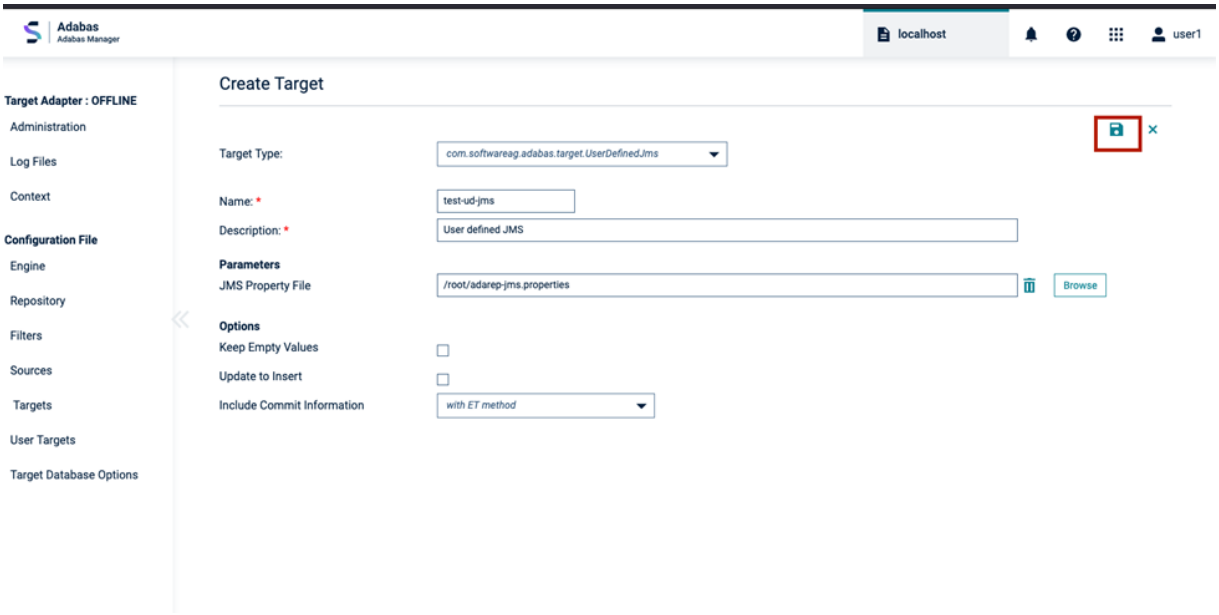

2. The JMS target is now available for selection in replication source definitions.

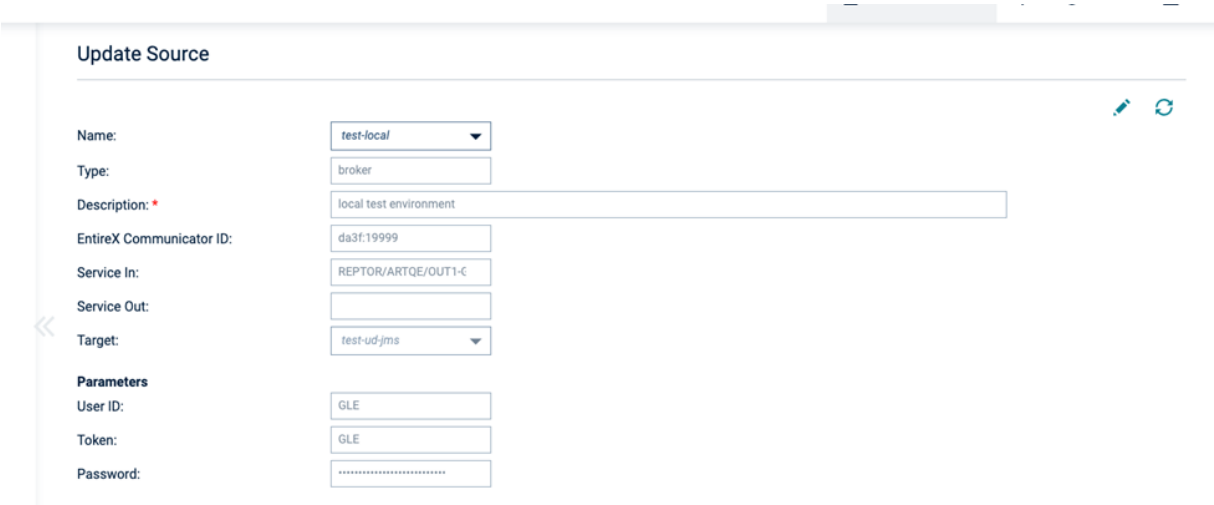

3. Start the Event Replicator Target Adapter to activate data replication.

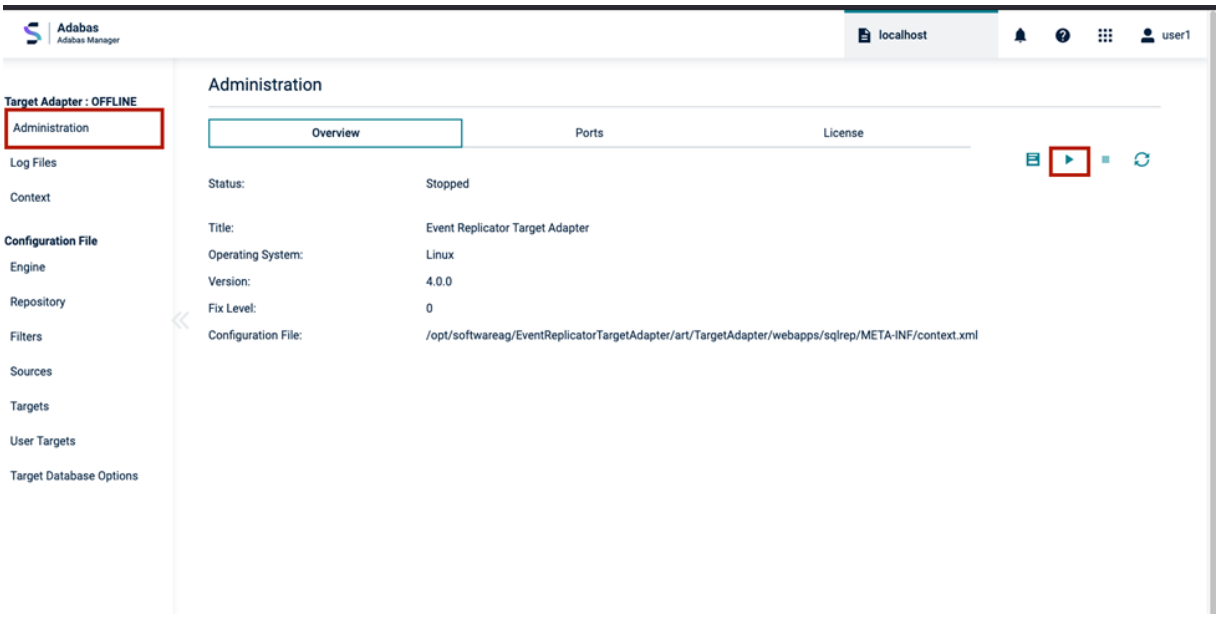

Usage

Once configured, the JMS User Target can be assigned to a replication source within the Event Replicator setup, enabling real-time data propagation to JMS infrastructure.

Replicated Adabas database changes are published automatically to the configured JMS queue as transactions are committed.

Notes and Best Practices

- Ensure version compatibility between the JMS client library and the JMS broker.
- Verify network connectivity between the ART host and the JMS endpoint before going live.
- Secure credentials appropriately and avoid plaintext credentials in production environments where possible.
- Monitor ART logs regularly for connection errors or message delivery issues.

Migrating User Targets to Adabas Manager

If you have migrated a context.xml file from a previous version of the Event Replicator Target Adapter and are now using Adabas Manager to manage adapter configuration, re-register the User Target in the Adabas Manager interface.

8 Replicating Data to Adabas

- Requirements 44
- Adabas-Specific Notes 45

You can use the Event Replicator Target Adapter to replicate data to Adabas.

The following diagram depicts the interaction between the Event Replicator Target Adapter and Adabas.

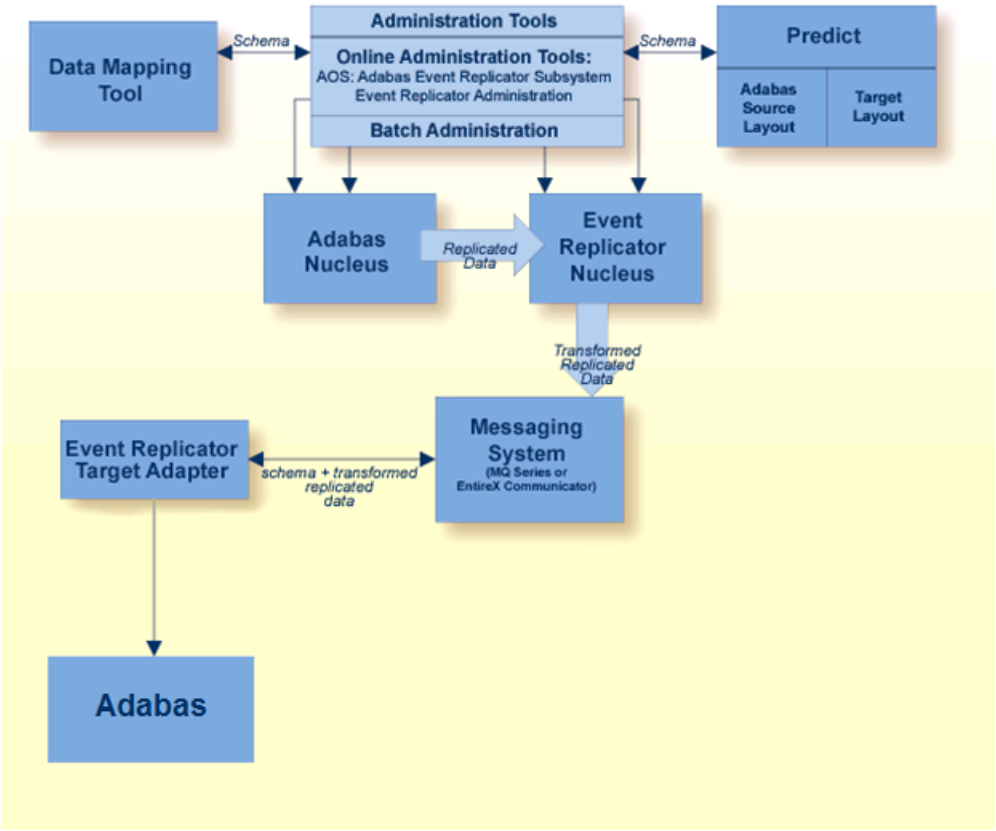

This chapter describes how to set up the Adabas target in the Event Replicator Target Adapter.

Requirements

- The Event Replicator Target Adapter uses the Adabas Client for Java internally to write the data to Adabas. See the requirements for Adabas Client for Java.
- Direct replication using Adabas file short names requires Event Replicator for Adabas Version 3.7 SP1 with zap AZ371002 applied on the mainframe.

Adabas-Specific Notes

Adabas Client for Java and the Map File

The Event Replicator Target Adapter uses Adabas Client for Java to process the data. For details about the concept, especially instructions on creating and configuring a Map File in the target database, see the section *Creating Maps* in the *Adabas Client for Java* documentation.

Depending on the user's desired end result and configuration, there are times when a map file is required and times when it is not.

A Map File and target Map *are not* required when:

- The Adabas source and target file numbers are the same.
- Event Replicator for Adabas Version 3.7 SP1 is installed on the mainframe with zap AZ371002 applied.
- The Event Replicator Destination is defined with `DCLASSPARM=NOSPRE,ADA`.
- The user wants the Event Replicator Target Adapter to create the target file and is not concerned if all extended field attributes are replicated when the target file is created. By default, only UQ and DE attributes are sent as defined in the GFFT, with all fields defined at the target as Null Suppressed. As a workaround to this behavior, the user can preload the target file with an FDT that fully defines the extended field attributes. The file number must still match the source database file number.

When no Map File is defined in the Adabas target definition in Event Replicator Target Adapter Administration, none will be used or created during Initial State processing. However, the requirements outlined above must be met or the Target Adapter will throw an exception and stop processing.

A Map File and target Map *are* required when:

- The Adabas source and target file numbers are different.
- The mainframe replicator configuration does not meet the requirements outlined above (e.g. short names are not being sent in the schema and replicated data).
- The user desires the target file definition with specific field attributes exactly matching the source (e.g. using the same predefined FDT for file and target).

When a Map File is defined in the Adabas Target definition, a Map is always used. It can be:

- Predefined using the Data Designer in Adabas Client for Java by preloading the target file with an FDT that matches the source file FDT or using a systrans'd DDM. See *Creating Maps* in the *Adabas Client for Java* documentation for details. In this case, the name of the target file and the

name of the map must match and be the same as the file name of the source. The target file number need not be the same as the source file number. or

- Created during the Initial State processing, thereby mapping field Long Names from the GFFT definition used by the Event Replicator for Adabas to Adabas short names in the target file "on the fly". If the target database file number that matches the source is not available, the next available file number in the target database will be used.

Note: If the Event Replicator for Adabas is not running Version 3.7 SP1 with AZ371002 applied and with `DCLASSPARM=ADA` defined, the field short names on the target side are assigned by the Target Adapter and will not necessarily agree with those from the target.

9 Replicating Data to Adabas in Binary Mode

- Requirements 48
- Adabas-Specific Notes 48

You can use the Event Replicator Target Adapter to replicate data to an Adabas LUW target using raw binary data sent from the Event Replicator on the mainframe. The Event Replicator Target Adapter uses Entire Net-Work Server to connect to an Adabas LUW target and Adabas direct calls to process the replicated transactions..

This chapter describes how to set up the Adabas target in the Event Replicator Target Adapter.

Requirements

- Direct replication in binary mode requires Event Replicator for Adabas Version 3.8 SP1 (or higher) on the mainframe with the following zaps applied:
  - AZ381005
  - AZ381007
  - AZ381013
- The Destination used in the Event Replicator Subscription must use the Destination Class Parameter (DCLASSPARM) "SAGTRGB". For details, please refer to the *Event Replicator for Adabas* documentation.
- Your target database must be accessed via an Entire Net-Work node (kernel) which registers the database for client access.
- The Entire Net-Work node must register the target with a Adabas Directory Server. The target adapter functions as a Directory Server client to retrieve the appropriate access entry for the target database ID; it then uses the communication path to the appropriate Entire Net-Work node provided by the Directory Server. For details on configuring Entire Net-Work to register Adabas databases, see the *Entire Net-Work* documentation.

Adabas-Specific Notes

Adabas (Binary) as Targets for Target Adapter - Use and Limitations

In the current implementation, there are certain restrictions and limitations to using the Adabas (Binary) target for replication:

- Entire Net-Work and a Directory Server must be used for the Target Adapter to access the Adabas LUW target, even if the target database and Target Adapter are installed on the same system. There is currently no way to provide the Target Adapter a direct path to the target, e.g. an ADATCP port or direct path to an Entire Net-Work Server port.
- Event Replicator Target Adapter Administration is used to configure the source-to-target mapping to be used in replication to an Adabas (Binary) target. It is possible to map multiple source DBIDs and files to multiple target files in a single mapping. For examples and an explanation

of how the mapping is configured, refer to *Event Replicator for Adabas* > *Event Replicator Target Adapter* > *Event Replicator Target Adapter Administration* > *Using Event Replicator Target Adapter Administration* > *Maintaining Adabas Target Definitions*.

- Event Replicator Target Adapter has the ability to create the configured target file in the target database if it does not exist when a schema is received and replicated. How this is done depends on the method used to initially create and populate the target file. The following methods are available:
  - Use the ADALOD utility on the mainframe. Use of this method requires the following:
    1. A parameter setting of DREPLICATEUTI=Y on the Event Replicator Destination definition used.
    2. The ADALOD parameter RPLLOAD=FILE in the ADALOD job on the mainframe.

    When this method is invoked using the specified parameters and the mainframe database is loaded, creation of the configured target file is triggered including the FCB/FDT of the source file and the target file populated as well. Using this approach, the resulting target file is an exact replica of the source. For more information and details on using the ADALOD utility, see the *Adabas for Mainframes* documentation.
  - Using SYSRPTR on the mainframe, perform an Initial State of the file to be replicated (in SYSRPTR, Administrative Functions --> Perform Initial State). This method may also be used to create a target file, *however the resulting FDT of the created file will not be identical to the FDT of the mainframe source*. The extended field attributes of the created fields and super-/sub-descriptors specified in the source FDT will be missing. If the user wishes to use this method and requires these extended attributes on the target file's fields, the file must be pre-loaded with an FDT identical to the one used on the mainframe source file. Refer to the Adabas LUW documentation for information on how to load an FDT for a new file.

- In all replication cases except the ADALOD scenario described above, the actual data replicated to the target Adabas file is controlled by the GFB/GFFT specified in the Event Replicator Subscription that controls replication. Source fields missing from or defined differently than the source field will be replicated as per the active GFB. For example, if field AE in the source is defined as AA,20,A, but defined in the Subscription File Parameters GFB as AA,1,A, only the first byte of the field will be replicated. The user should take care that the GFB/GFFT defined in the Subscription describes the data desired for replication. For best results, when defining and generating the GFB/GFFT, we recommend use of the *Event Replicator Target Adapter Data Mapping Tool* to create optimized GFBs for use in your replication scenario.
- For a listing of certain restrictions and limitations on the use of the Adabas (Binary) Target feature, refer to the Release Notes and product and fix Readme files.

10 Replicating Data to Terracotta Caches

- Requirements 52
- Implementation Steps 53
- Terracotta-specific Notes 54

You can use the Event Replicator Target Adapter to replicate data to Terracotta in memory caches.

The following diagram depicts the interaction between the Event Replicator Target Adapter and Terracotta.

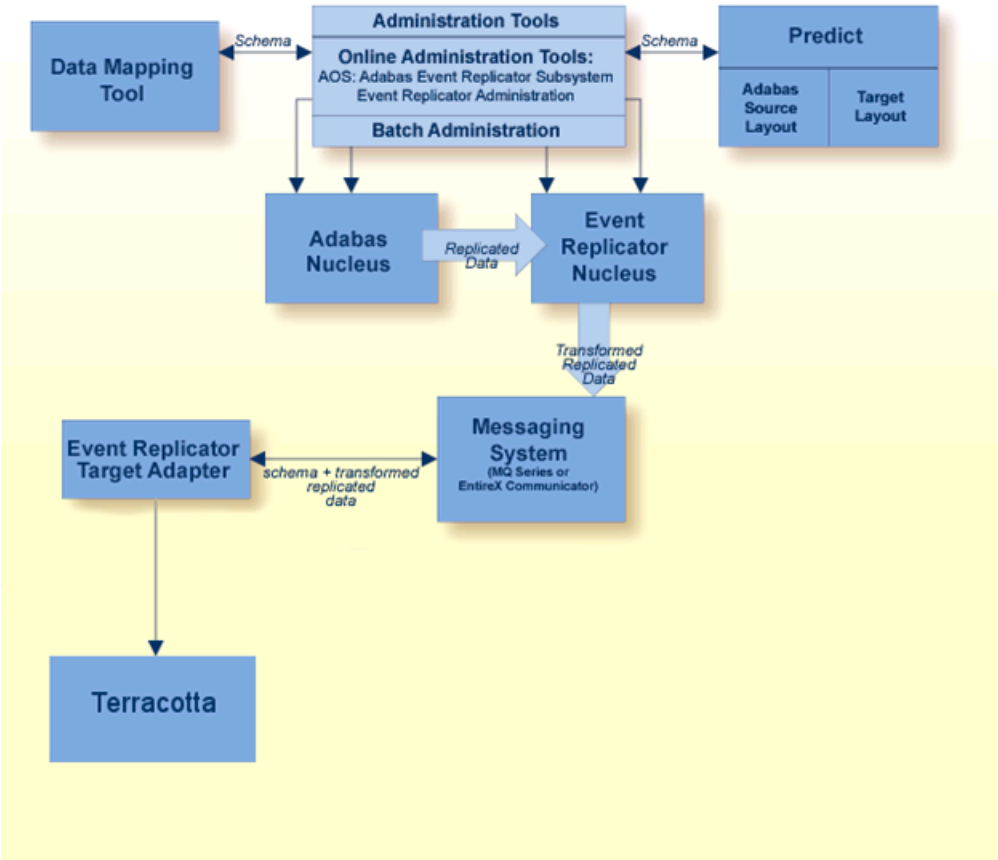

This document describes how to set up this Terracotta support in the Event Replicator Target Adapter.

Requirements

Terracotta Version 4.3 must be installed in your enterprise. For complete information, read *Terracotta*, available in the "IBO and Big Data" documentation on Software GmbH's **Empower** web site. The necessary components to connect to the Terracotta target are provided with the delivery of Event Replicator Target Adapter, except for the licenses described under Implementation Steps below.

Implementation Steps

Note: In the steps described below, the value of <*Root*> is the upper level directory where Software GmbH technology products are installed. On Windows, the default is `C:\SoftwareAG` and on Linux/UNIX systems the default is `/opt/softwareag`. You may have installed your Software GmbH products to a different root, in which case you would replace <*Root*> with this location.

Licenses

Event Replicator Target Adapter must be licensed to communicate with the Terracotta target. This is an additional license to the one you used when installing v. In addition, if you are connecting to/using the enterprise features of the Terracotta cache server, you must also provide a copy of that license to the v environment.

To license Event Replicator Target Adapter for connection to the Terracotta server:

- Obtain a license for the Terracotta target from your Software AG representative. This license will likely be provided to you in the form of an .xml file with a name like <*numericprefix*>*_arttc.xml*.
  - Rename this file to *arttc.xml*.
  - Manually copy this file to <*Root*>/common/conf.

You will be able to connect to the Terracotta server without this license file, but your ability to use Terracotta features will be limited to 30 minutes of use in a single Event Replicator Target Adapter session.

To use the Terracotta enterprise features:

To use the features of the enterprise versions of the Terracotta products a Terracotta license is required. This is the same license you were provided in order to use the enterprise features of your Terracotta server installation. A copy of this license must be manually placed in a resource location available to Event Replicator Target Adapter:

- **Recommended**: simply copy the file to:

  Windows: <*Root*>\EventReplicatorTargetAdapter\art\data\webapps\sqlrep\WEB-INF\classes.

  UNIX: <*Root*>/EventReplicatorTargetAdapter/art/TargetAdapter/webapps/sqlrep/WEB-INF/classes.

  This is a resource location already “well known” to Event Replicator Target Adapter and the application will load the license from this resource location.

- **Optional** (not recommended as it involves manually editing an installed file):

  Set the path to the license in the CATALINA_OPTS section of the ERTA.bat (Windows)/ERTA.sh (Unix) file. This file will be found in:

  Windows: <*Root*>\EventReplicatorTargetAdapter\art\program\bin

  UNIX: <*Root*>/EventReplicatorTargetAdapter/art/bin.

  Add the property:

  ```
  -Dcom.tc.productkey.path=/path/to/terracotta-license.key
  ```

  …at the end of the line that sets CATALINA_OPTS.

  For example:

  Windows: set CATALINA_OPTS=-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=%ART_JMX_PORT% -Dcom.sun.management.jmxremote.ssl=false -Dcom.sun.management.jmxremote.authenticate=false -Dcom.tc.productkey.path=/path/to/terracotta-license.key

  UNIX: CATALINA_OPTS="-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=${ART_JMX_PORT} -Dcom.sun.management.jmxremote.ssl=false -Dcom.sun.management.jmxremote.authenticate=false" -Dcom.tc.productkey.path=/path/to/terracotta-license.key

For other possibilities on specifying the location of this file, check the Terracotta documentation.

Terracotta-specific Notes

Cache Configuration File

The Ehcache cache configuration file (*ehcache.xml*) contains the configuration for the CacheManager that is used by the client. It contains all information about the environment, parameters, caches, etc. At a minimum, the *ehcache.xml* file used by Event Replicator Target Adapter should contain the following:

```
<?xml version="1.0" encoding="UTF-8"?>

<ehcache xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:noNamespaceSchemaLocation="ehcache.xsd" updateCheck="false"
 monitoring="autodetect" dynamicConfig="true">

 <!-- Mandatory Default Cache configuration. These settings will be applied
  to caches created programmtically using CacheManager.add(String cacheName).
  The defaultCache has an implicit name "default" which is a reserved cache
  name. -->
```

```
 <defaultCache eternal="true"
  overflowToDisk="false" maxEntriesLocalHeap="1">
  <terracotta/>
 </defaultCache>

 <terracottaConfig url="terracottahost:port" rejoin="true"/> This line is only ↩
needed for interaction with the enterprise version of Terracotta.
</ehcache>
```

For more information about the Cache Configuration File, read *Terracotta*, available in the "IBO and Big Data" documentation on Software GmbH's **Empower** web site.

Cache Definition

Although not necessary, cache definitions can be defined in the Cache Configuration File (*ehcache.xml*). The path to this file is specified using Event Replicator Target Adapter Administration in the your target configuration.

The main reason to use a cache definition in *echache.xml* is to enable a client application to search by a particular key. Event Replicator Target Adapter does not require your cache be defined in this file; however, specifying the cache definition in *echache.xml* does let you control which Adabas descriptor fields are to be marked in the Terracotta cache as keys (searchable fields). If no cache definition exists in *echache.xml*, the cache will be dynamically created when a CREATE statement (schema) is sent from Event Replicator for Mainframe and all fields specified in the Adabas DDM as descriptors will be marked as searchable in the Terracotta cache. For more information see Searchable Attributes below.

AdabasObject

A Terracotta cache is in general a Key-Value store for Java objects. To store Adabas records in such a Key-Value store the record must be transformed to a Java object (value). The key for this object is the key that is provided by the Event Replication. Unless you specified a custom Primary Key in the replication GFB/GFFT, the ISN by default will be used as key. The Java object that is stored in the cache is an instance of AdabasObject, which is a simple Java class. To retrieve the fields from this objects the method Object `evaluateValue(String fieldName)` must be used, where fieldname is the long name for this field as defined in the GFFT (e.g. NAME, Address, etc.). The result type depends on the type of the Adabas field:

Adabas Structure	Adabas Object	Example
Field	Java type	String, int, float,
Multiple field	List of Java type	List <String>, List<int>
Group	Adabas Object	AdabasObject
Periodic group	List of Adabas Object	List<AdabasObject>

Note: Groups are not supported by the Event Replication Target Adapter.

Searchable Attributes

The data in the Terracotta cache is accessible via the key. To use the Terracotta search capabilities the attributes must be marked as searchable. By default, when a cache is created with receipt of a schema from Event Replicator Target Adapter, any Adabas field marked as a Descriptor in the GFFT is created as a searchable key in the Terracotta cache. For information on how to use Searchable Attributes to query the cache, please refer to the *Terracotta* documentation in **Empower**.

11 Replicating Data to Terracotta

- Requirements 58
- Terracotta-specific Notes 59

You can use the Event Replicator Target Adapter to replicate data to Terracotta.

The following diagram depicts the interaction between the Event Replicator Target Adapter and Terracotta.

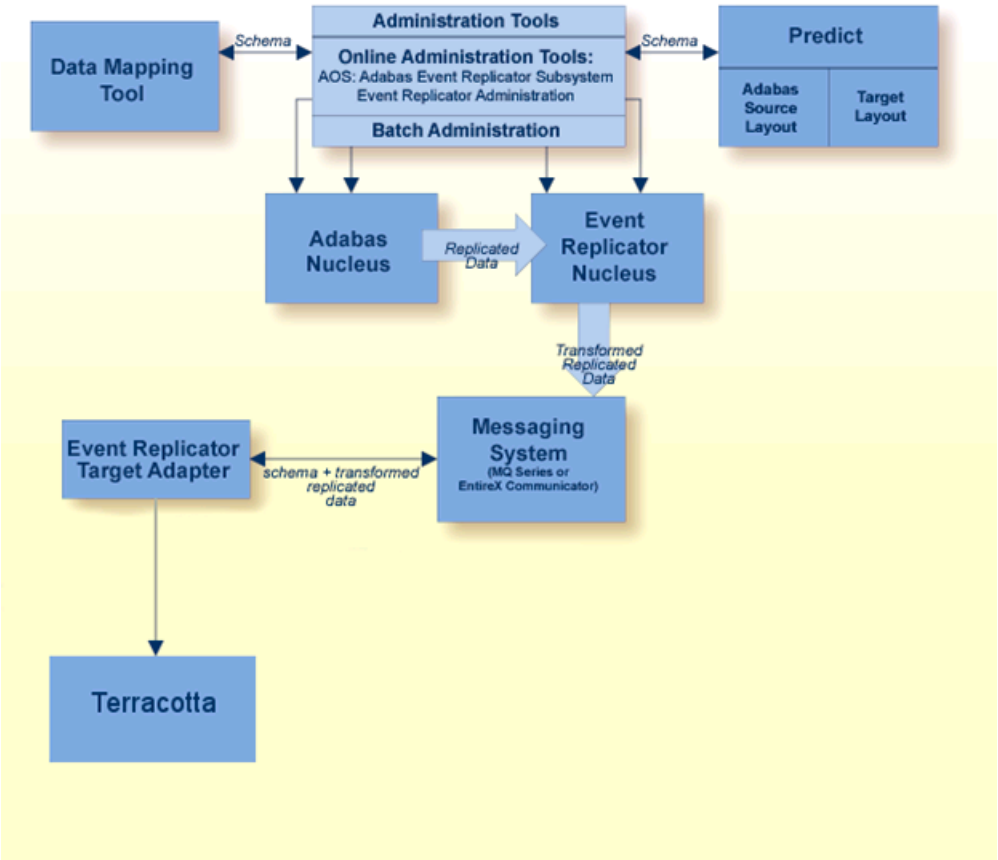

This chapter describes how to set up the Terracotta target in the Event Replicator Target Adapter.

Requirements

Terracotta Version 10.1 or higher must be installed in your enterprise. For complete information, read *TCStore API Developer Guide*, available in the *Terracotta* documentation on Software GmbH's **Empower** web site. The necessary components to connect to the Terracotta target are provided with the delivery of Event Replicator Target Adapter. Replication to Terracotta requires a license.

Terracotta-specific Notes

Sub-Datasets

Terracotta is very similar to a relational database. This means that Adabas multiple fields (MU) and periodic groups (PE) are represented as their own sub-datasets, similar to the creation of child tables in relational databases.

12 Using the DB2 CONNECT IMPORT Command to Load Initial-State Data

- Prerequisites 62
- Limitations 62
- Enabling DB2 Database Initial-State Data Loading 62
- Processing Output 64

You can use the DB2 CONNECT IMPORT command to process and load initial-state data into aDB2 database. Only ASCII codepages are supported.

Note: This support is provided only in Windows environments at this time.

Prerequisites

Before the DB2 CONNECT IMPORT command can be used to load initial-state data to a DB2 database, via the Event Replicator Target Adapter, DB2 CONNECT for Windows must be installed. In addition, using the Configuration Assistant from DB2 CONNECT, make sure you configure the alias that should be used as a database name when importing the initial-state data occurs.

Limitations

The following limitations exist in this version:

- This support is provided only in Windows environments at this time.
- Only ASCII codepages are supported.

Enabling DB2 Database Initial-State Data Loading

To enable DB2 database initial-state data loading, complete the following steps:

1. Once all prerequisites have been met, use the Administration tool to configure a target definition for your DB2 database. For more information, read *Configuring Target Definitions for Event Replicator Target Adapter* in the *Event Replicator Target Adapter Administration Guide*.
2. Using the Administration tool, create a specific target processing option definition for your DB2 database. For more information about setting up target processing option definitions, read *Specifying Target Database Processing Option Definitions*, in the *Event Replicator Target Adapter Administration Guide*.

   In the target processing option definition for your DB2 database:

   Note: No value need be specified in the **Environment File** field.

   - Select the **Use Loader** check box (check it).
   - Enter the alias name of the DB2 database in the **Oracle Service, Teradata Tdpip, Microsoft SQL Schema Name or DB2 Alias** field.

- Provide a value for the **Path to Loader Executable** field, identifying the full path of the DB2 executable, *db2cmd.exe*.
- Specify any of the following options, as required for your installation, in the Options field.

  When specified, these options must be specified in the order shown in the table and they must be separated by a space. The Event Replicator Target Adapter does not validate these command options; it just passes them to the DB2 CONNECT IMPORT command for processing. For complete information about these options, refer to your *DB2 Command Reference*.

Option	Description
ALLOW WRITE ACCESS	This option causes the Event Replicator Target Adapter to run the import process in online mode. Specify this option only when you want to update your Adabas database at the same time as running an initial-state. It will avoid the database table locks required by the IMPORT command. Otherwise, this option is only required when you import data to a table with a nickname (in which case the COMMITCOUNT option must also be supported with an integer setting, not the literal "AUTOMATIC"). When this option is specified, an intent exclusive (IX) lock on the target table is acquired when the first row is inserted. This allows concurrent readers and writers to access table data. The import operation will periodically commit inserted data to prevent lock escalation to a table lock and to avoid running out of active log space. These commits will be performed even if the COMMITCOUNT option is not specified for the run. During each commit, the import process will lose its IX table lock, and will attempt to reacquire it after the commit.
COMMITCOUNT { *n* | AUTOMATIC }	This option controls when commits of imported records occur. Valid values are an integer representing the number of imported records or the literal "AUTOMATIC". To reduce network overhead, Event Replicator Target Adapter import processes use NOT ATOMIC compound SQL to insert the data, with 100 SQL statements attempted with each import. If your SQL transaction log is not sufficiently large, the import operation will fail. It must be large enough to accommodate the number of rows specified by the COMMITCOUNT option or the number of rows in the data file if COMMITCOUNT is not specified. We therefore recommend that the COMMITCOUNT option be specified to avoid transaction log overflows. When a number (*n*) is specified, the import process performs a commit after that many records have been imported. When compound inserts are used, a user-specified commit frequency of *n* is rounded up to the first integer multiple of the compound count value. When the literal "AUTOMATIC" is specified, the import process internally determines when a commit should be performed. The utility will commit the imported data for one of two reasons: ■ To avoid running out of active log space; or

Option	Description
	■ To avoid a lock escalation from row level to table level. If the ALLOW WRITE ACCESS option is specified, and the COMMITCOUNT option is not specified, the import process will perform commits as if COMMITCOUNT AUTOMATIC had been specified.
WARNINGCOUNT *n*	This option controls when the import operation stops because of warnings. Valid values are integers beginning with zero (0). Set this parameter if no warnings are expected, but you want to verify that the correct file and table are being used. If the import file or the target table is specified incorrectly, the import utility will generate a warning for each row that it attempts to import, which will cause the import to fail. If a value of zero is specified for this option or if this option is not specified, the import operation will continue regardless of the number of warnings issued.
NOTIMEOUT	This option specifies that the import utility will not timeout while waiting for locks.

Save the target processing option definition with the same name as the target definition you created in Step 1.

3 When all processing options are specified, start initial-state processing, directing the output to the Event Replicator Target Adapter via a GFFT you created (using the Data Mapping Tool or the Adabas Event Replicator Subsystem) for the DB2 database. For more information on this, read ***Requesting Initial-State Data***, elsewhere in this guide.. The Event Replicator Target Adapter will collect the initial-state records and load them into your DB2 database.

Processing Output

The output from this processing includes:

- Various DB2 CONNECT IMPORT files used for its processing;
- DB2 CONNECT IMPORT log files from the processing; and
- DAT files containing the initial-state data received by the Event Replicator Target Adapter prior to being imported.

All output files are written to the *logs* subdirectory of your Event Replicator Target Adapter installation.

- Processing Files

- Log Files

Processing Files

The Event Replicator Target Adapter's load processing for DB2 creates two processing files per table loaded. The following processing files are stored in the *\logs* subdirectory of your Event Replicator Target Adapter installation:

- A single *xxxx_yyyymmdd_hhmmss.BAT* file for all DB2 CONNECT processing (where `xxxx` is the base table name, `yyyymmdd` is the date, and `hhmmss` is the time: This .BAT file is used to run the load utility.
- One or more *xxxx_yyyymmdd_hhmmss.DEL* files (where `xxxx` is the table name, `yyyymmdd` and `hhmmss` are the timestamp identifying when the file was written, and DEL indicates the file is a delimited ASCII file). These files contain the data to be loaded to the database tables. These files can take a good deal of space, so make sure that there is enough space to accommodate them and be sure to manually delete them when initial-state processing has completed successfully.

Log Files

Two kinds of log files are produced by Event Replicator Target Adapter import processing for DB2: a primary log file and multiple processing log files. All log files are stored in the *\logs* subdirectory of your Event Replicator Target Adapter installation.

1. The DB2 load processing primary log files make references to the processing log files. The primary log files have names in the format *xxxx_yymmdd_hhmmsss.out*, where `xxxx` is the table name, `yymmdd` is the date of the log file, and `hhmmsss` is the time (to the microsecond) when the log file was last updated.

   We recommend that you review these primary log files carefully to ensure that the import process had no problems restoring primary and foreign keys, indexes, or the unique index of tables. If import processing fails for some reason, we recommend reviewing these primary `.out` log files first when resolving the problem, proceeding then to the processing log files, as necessary.
2. Multiple processing log files are created, one for each import process executed by the Event Replicator Target Adapter. These processing log files have names in the format *xxxx_yyyymmdd_hhmmss.MSG*, where `xxxx` is the table name, `yyyymmdd` and `hhmmss` are the timestamp identifying when the file was written, and MSG indicates the file is a log file.

13 Using the Microsoft SQL Server Bulk Copy (bcp) Utility to Load Initial-State Data

- Prerequisites 68
- Limitations 68
- Enabling bcp Utility Initial-State Processing 68
- Processing Output 70

You can use the bcp utility to process and load initial-state data into a Microsoft SQL Server database. The bcp utility can only be used with the Event Replicator Target Adapter to load initial-state data; it cannot be used for transactional data.

Prerequisites

The following prerequisites must be met before the bcp utility can be used to load initial-state data to a Microsoft SQL Server database, via the Event Replicator Target Adapter:

- The Microsoft SQL Server utilities supporting the bcp utility must be installed on the same machine as the Event Replicator Target Adapter. The bcp utility version 9 or greater must be installed.
- Be sure there is enough space to accommodate the *.BCPDAT files that will be created by this processing.

Limitations

You can use the bcp utility to load initial-state data only if composite keys are *not* used in the data.

Enabling bcp Utility Initial-State Processing

To enable bcp utility initial-state processing, complete the following steps:

1. Using the Administration tool, configure a target definition for your Microsoft SQL Server database. For more information, read *Configuring Target Definitions for Event Replicator Target Adapter* in the *Event Replicator Target Adapter Administration Guide*.
2. Using the Administration tool, create a specific target processing option definition for your Microsoft SQL Server database.
   - In the Loader options area, check the box **Use Loader** and provide a value for the **Path to Loader Executable** field, identifying the full path of the Microsoft SQL Server bcp utility executable. For example:

```
C:\Program Files\Microsoft SQL Server\90\Tools\Binn\bcp.exe
```

- Optionally, in the Loader options area of the target processing option definition panel, specify the MS SQL Server schema name if your schema name is different from the MS SQL Server default schema name. Prior to the release of MS SQL Server 2005, all tables created by MS SQL Server used the user login name as the default schema name; in MS SQL Server 2005 and later versions, Microsoft uses "dbo" as the default schema name (which is different from the user login name). If you have used any MS SQL Server schema name that is not the same as the MS SQL Server default schema name, specify your schema name in the **Oracle Service or Teradata Tdpip** field.
- Specify the following options, as required for your installation, in the **Options** field.

Option	Description
-b *batchsize*	Use the -b (batch size) option in this field to specify the number of rows per batch of imported data. Each batch is imported and logged as a separate transaction that imports the whole batch before being committed. By default, all rows in the data file are imported as one batch. To distribute the rows among multiple batches, use the -b option to specify a batch size that is smaller than the number of rows in the data file. If the transaction for any batch fails, only insertions from the current batch are rolled back. Batches already imported by committed transaction are unaffected by a later failure. Do not use this option in conjunction with the -h ("ROWS_PER_BATCH=bb") option.

Note: The **Environment File**, and **Use Stream** fields in the Loader options area have no effect on the bcp utility, so they should not be specified.

Save the target processing option definition with the same name as the target definition you created in Step 1.

3 When all processing options are specified, start initial-state processing, directing the output to the Event Replicator Target Adapter via a GFFT you created (using the Data Mapping Tool or the Adabas Event Replicator Subsystem) for the Microsoft SQL Server database. For more information on this, read *Requesting Initial-State Data*, elsewhere in this guide.. The Event Replicator Target Adapter will collect the initial-state records and send them to the bcp utility to be loaded to your database.

Processing Output

The output from this processing includes:

- Various bcp utility processing files (*.*FMT* files)
- bcp utility log files from the processing (*.*BCPERR* and *.*BCPOUT* files)
- Binary data files (*.*BCPDAT* files) containing the initial-state data received by the Event Replicator Target Adapter prior to being processed by the bcp utility, and possibly *.*DBGDAT* files (if the debug logging option is turned on)
- An output file (*.*out* file) containing output messages from the final loader process run.

All output files are written to the *\logs* subdirectory of your Event Replicator Target Adapter installation.

- Processing Files
- Log Files

Processing Files

The Event Replicator Target Adapter's bcp utility processing creates three processing files per table loaded. The following processing files are stored in the *\logs* subdirectory of your Event Replicator Target Adapter installation:

- A single *xxxx_yyyymmdd_hhmmss.BAT* file for all bcp utility processing (where `xxxx` is the base table name, `yyyymmdd` is the date, and `hhmmss` is the time: This .BAT file is used to run the bcp utility. It generates the format file (*.FMT) for each table and loads the binary *.BCPDAT file into the database.
- One `xxxx_yyyymmdd_hhmmss.FMT` file for each table (where *xxxx* is the table name): This is the format file for the bcp utility processing.
- One or more *xxxx_yyyymmdd_hhmmss.BCPDAT* file for each table (where `xxxx` is the table name, `yyyymmdd` is the date, and `hhmmss` is the time): These files contain the data to be loaded to the database tables by the bcp utility. These files are in compressed binary format, but can still take a good deal of space, so make sure that there is enough space to accommodate them and be sure to manually delete them when initial-state processing has completed successfully.
- If the file debug logging option is turned on, the resulting debugging information is stored in a debug file with a name in the format *xxxx_yyyymmdd_hhmmss.DBGDAT*, where `xxxx` is the table name, `yyyymmdd` is the date of the debug file, and `hhmmss` is the time the debug file was last updated. Multiple debug files may be produced for a single table.

Log Files

Several kinds of log files are produced by Event Replicator Target Adapter bcp utility processing: a primary log file and multiple processing log files. All log files are stored in the *\logs* subdirectory of your Event Replicator Target Adapter installation.

1. The bcp utility primary log file makes references to the processing log files. The primary log file has a name in the format *xxxx_yyyymmdd_hhmmss_nnn.out*, where `xxxx` is the base table name, `yyyymmdd` is the date of the log file, and `hhmmss_nnn` is the time (hours, minutes, seconds and milliseconds) the log file was last updated. A single primary log file is created for all bcp utility processing. If bcp utility processing fails for some reason, we recommend reviewing these primary *.out log files first when resolving the problem, proceeding then to the processing log files, as necessary.
2. Multiple processing log files are created, one for each bcp process executed by the Event Replicator Target Adapter. These processing log files have names in the format *xxxx_yyyymmdd_hhmmss.BCPOUT*, where `xxxx` is the table name, `yyyymmdd` is the processing date (year, month, and day) and `hhmmss` represents the processing time (hours, minutes, and seconds).
3. Any errors encountered during bcp utility processing will be written to error files with names in the format *xxxx_yyyymmdd_hhmmss.BCPERR*, where `xxxx` is the table name, `yyyymmdd` is the processing date (year, month, and day) and `hhmmss` represents the processing time (hours, minutes, and seconds). A single*.BCPERR file is created for each table.

14 Using the Oracle SQL*Loader to Load Initial-State Data

- Prerequisites 74
- Enabling Oracle SQL*Loader Initial-State Processing 74
- Processing Output 75

You can use the Oracle SQL*Loader to process and load initial-state data into an Oracle database. This process is supported in two modes: file mode and stream mode:

- In file mode, initial-state data is received by the Event Replicator Target Adapter and stored in a DAT file prior to being sent to the Oracle SQL*Loader for processing. For more information about the DAT files created as part of this processing, read ***Processing Output***, elsewhere in this chapter
- In stream mode, initial-state data is received by the Event Replicator Target Adapter and stored in memory prior to being sent to the Oracle SQL*Loader for processing. No DAT files are created.

Note: The Oracle SQL*Loader can only be used with the Event Replicator Target Adapter to load initial-state data; it cannot be used for transactional data.

Prerequisites

The following prerequisites must be met before the Oracle SQL*Loader can be used to load initial-state data to an Oracle database, via the Event Replicator Target Adapter.

- The Oracle utilities supporting the Oracle SQL*Loader must be installed on the same machine as the Event Replicator Target Adapter.
- The Oracle Net Manager must be installed and configured if Oracle is running on a different machine than the Event Replicator Target Adapter.
- If file mode is requested, be sure there is enough space to accommodate the DAT files that will be created by this processing.

Enabling Oracle SQL*Loader Initial-State Processing

To enable Oracle SQL*Loader initial-state processing:

1. Using the Administration tool, configure a target definition for your Oracle database. For more information, read *Configuring Target Definitions for Event Replicator Target Adapter*, in the *Event Replicator Target Adapter Administration Guide*.
2. Using the Administration tool, create a specific target processing option definition for your Oracle database. For this definition you must:
   - Select the **Use Loader** check box (check it).
   - Provide a value for the **Path** field, identifying the full path of the Oracle SQL*Loader executable.

- Specify the `rows` option in the **Options** field, specifying the number of rows you want to read from the data file before a data save is performed by the SQL*Loader.
- Save the target processing option definition with the same name as the target definition you created in Step 1.

You should also evaluate and optionally specify values for the following fields in this target processing option definition:

- **Service**, which specifies the Oracle database service name that should be used with the Oracle SQL*Loader.
- **Environment File**, which specifies the file name containing the Oracle environment variables that should be used with the Oracle SQL*Loader.
- **Use Stream**, which indicates whether processing should be performed in stream mode or not.

For more information about these options, read *Specifying Target Database Processing Option Definitions*, in the *Event Replicator Target Adapter Administration Guide*.

3 When all processing options are specified, start initial-state processing, directing the output to the Event Replicator Target Adapter via a GFFT you created (using the Data Mapping Tool or the Adabas Event Replicator Subsystem) for the Oracle database.

For more information on this, read ***Requesting Initial-State Data***, elsewhere in this guide.

The Event Replicator Target Adapter will collect the initial-state records and send them to the Oracle SQL*Loader to be loaded to your Oracle database.

Note: Some errors might prevent the Oracle SQL*Loader from restoring primary and foreign keys, indexes, and the unique index of tables. We recommend that you examine the Oracle SQL*Loader log files to determine if such an error occurred. For more information, read ***Processing Output***, elsewhere in this section.

Processing Output

The output from this processing includes:

- Various Oracle SQL*Loader files used for its processing
- Oracle SQL*Loader log files from the processing
- DAT files containing the initial-state data received by the Event Replicator Target Adapter prior to being processed by the Oracle SQL*Loader.

All output files are written to the *\logs* subdirectory of your Event Replicator Target Adapter installation.

- Log Files
- DAT Files

Log Files

The Oracle SQL*Loader log files are stored in files that have names in the format *xxxx_yymmdd_hhmmss*.log, where *xxxx* is the table name, *yymmdd* is the date of the log file, and *hhmmss* is the time the log file was last written to.

We recommend that you review this log file carefully to ensure that the Oracle SQL*Loader had no problems restoring primary and foreign keys, indexes, or the unique index of tables.

DAT Files

DAT files are only created by the Event Replicator Target Adapter if you have selected Oracle SQL*Loader processing in file mode. If you select stream mode, these files are not created. These files can take a good deal of space, so make sure that there is enough space to accommodate them and be sure to manually delete them when initial-state processing has completed successfully.

The format of the DAT file names is *xxxx_yymmdd_hhmmss*.dat, where *xxxx* is the table name, *yymmdd* is the date of the DAT file, and *hhmmss* is the time the DAT file was last written to.

15 Using the PostgreSQL psql Utility to Load Initial-State Data

- Prerequisites 78
- Enabling PostgreSQL SQL Initial-State Processing using psql 78
- Processing Output 79

You can use the PostgreSQL psql utility to process and load initial-state data into a PostgreSQL database.

Note: The PostgreSQL psql utility can only be used with the Event Replicator Target Adapter to load initial-state data; it cannot be used for transactional data.

This document covers the following topics:

Prerequisites

The following prerequisite must be met before the PostgreSQL psql utility can be used to load initial-state data to a PostgreSQL database, via the Event Replicator Target Adapter.

The PostgreSQL utility "psql" (psql.exe on Windows, psql on Linux) must be installed on the same machine as the Event Replicator Target Adapter.

Enabling PostgreSQL SQL Initial-State Processing using psql

To enable PostgreSQL initial-state processing using psql:

1. Using the Administration tool, configure a target definition for your PostgreSQL database. For more information, read *Configuring Target Definitions for Event Replicator Target Adapter* in the *Event Replicator Target Adapter Administration Guide*.
2. Using the Administration tool, create a specific target processing option definition for your PostgreSQL database. For this definition you must:
   - Select the **Use Loader** check box (check it).
   - Provide a value for the Path field, identifying the full path of the PostgreSQL *psql* executable.
   - Save the target processing option definition **with the same name** as the target definition you created in Step 1.

     For more information about these options, read *Specifying Target Database Processing Option Definitions*, in the *Event Replicator Target Adapter Administration Guide*.
   - When all processing options are specified, start initial-state processing, directing the output to the Event Replicator Target Adapter Adapter via a GFFT you created (using the Data Mapping Tool or the Adabas Event Replicator Subsystem) for the PostgreSQL database.

     For more information on this, read ***Requesting Initial-State Data***, elsewhere in this guide..

     The Event Replicator Target Adapter will collect the initial-state records and send them to the psql utiltiy to be loaded to your PostgreSQL database.

Processing Output

The output from this processing includes:

- Output files containing the output from the psql utility.
- psql utility log files.
- Data files containing the initial-state data received by the Event Replicator Target Adapter prior to being processed by the PostgreSQL psql utility.

All output files are written to the \logs subdirectory of your Event Replicator Target Adapter installation.

This section covers the following topics:

- Log Files
- Processing Files

Log Files

The PostgreSQL psql log files are stored in files that have names in the format *xxxx_yymmdd_hhmmss*.log, where *xxxx* is the table name, *yymmdd* is the date of the log file, and *hhmmss* is the time the log file was last written to. A similarly names file with the extention .out contain the actual output of the psql utility.

In the event of errors received during processing of initial-state data using the psql utility, please examine these files for hints as to the cause of the error.

Processing Files

A single processing file for each table with the extension .sql is created by the Event Replicator Target Adapter to feed into the psql utiltiy. This files can take a good deal of space, so make sure that there is enough space to accommodate it and be sure to manually delete it when initial-state processing has completed successfully. This file contains not only the data loaded into the table(s), but also the actual psql commands passed to the utility.

The format of the processing file name is *xxxx_yymmdd_hhmmss*.sql, where *xxxx* is the table name, *yymmdd* is the date of the file, and *hhmmss* is the time the sql data file was last written to.

16 Using the Teradata MultiLoad (MLoad) Utility to Load Initial-State Data

- Prerequisites 82
- Limitations 83
- Enabling Teradata MLoad Utility Initial-State Processing 83
- Processing Output 84

You can use the Teradata MultiLoad (MLoad) utility to process and load initial-state data into a Teradata database. The MLoad utility can only be used with the Event Replicator Target Adapter to load initial-state data; it cannot be used for transactional data. It can load UTF8 data.

Note: The MLoad utility is supported in Windows environments only. In addition, composite keys are not supported when using the MLoad utility with Event Replicator Target Adapter.

This chapter covers the following topics:

Prerequisites

The following prerequisites must be met before the Teradata MLoad utility can be used to load initial-state data to a Teradata database, via the Event Replicator Target Adapter:

- The Teradata utilities supporting the MLoad utility must be installed on the same machine as the Event Replicator Target Adapter.
- Be sure there is enough space to accommodate the DAT files that will be created by this processing.
- After installing the Teradata MLoad utility, it must be configured. According to Teradata MLoad utility configuration rules, you should add entries in the Windows HOSTS file that map the TDPIDs (Teradata Director Program Identifier) with the Teradata databases and their hosts. When the MLoad utility is invoked for a particular target, the Event Replicator Target Adapter will relay the TDPID provided in the Loader section of the Target Database Options entry definition (in Event Replicator Target Adapter Administration) into the Teradata .login command. For example:

```
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#

127.0.0.1       localhost
n.n.n.n         tdpid tdpidcop1
```

The following substitutions are used in this example:

Term	Description
n.n.n.n	The specific TCP/IP address for the named TDPID where the Teradata database expected to interact with the MLoad utility is running.
tdpid	The Teradata Director Program Identifier (TDPID); usually the first eight characters of the Teradata database name or the Teradata service name. Teradata database names can be up to 30 characters long. However, MLoad restrictions for Teradata releases up to Version 12 require that TDPIDs be eight (8) characters or less.

Term	Description
	Event Replicator Target Adapter current logic dynamically determines the TDPID using the first eight characters of the Teradata service name specified in the Loader options of the database option definition created using the Event Replicator Target Adapter Administration tool. If you do not supply a service name, the first eight characters of the Teradata database name is used. Whatever TDPID is used must also be present in the HOSTS file. Be sure that the specifications in the database option definition and the specification in the HOSTS file match.
`tdpidcop1`	Teradata requires that you specify a second parameter in the format `tdpidcop1`, where `tdpid` is the TDPID you specified above and "cop1" is a literal. For example, if you specify "A2345678" as your TDPID, you would specify "A2345678cop1" for this parameter.

Limitations

You can only use the MLoad utility to load initial-state data if:

- Composite keys are *not* used in the data.
- You are running Event Replicator Target Adapter, Teradata, and the MLoad utility in Windows environments.

Enabling Teradata MLoad Utility Initial-State Processing

To enable Teradata MLoad utility initial-state processing, complete the following steps:

1 Using the Administration tool, configure a target definition for your Teradata database. For more information, read *Configuring Target Definitions for Event Replicator Target Adapter* in the *Event Replicator Target Adapter Administration Guide*.

> **Note:** If you intend to load UTF-8 data to your Teradata database, be sure you have applied the appropriate Event Replicator Target Adapter hot fix and that you have specified the "CHARSET=utf8" parameter in the data source URL. For more information about the hot fixes that support this, refer to the README file provided with Event Replicator Target Adapter.

2 Using the Administration tool, create a specific target processing option definition for your Teradata database. For this definition you must:

- Select the **Convert Hyphen** option (check it) on the target processing option definition panel.

- In the Loader options area of the target processing option definition panel, select the **Use Loader** check box (check it).
- In the Loader options area of the target processing option definition panel, specify the Teradata Director Program Identifier (TDPID) in the **Oracle Service or Teradata Tdpip** field. This is usually the first eight characters of the Teradata database name or the Teradata service name. Teradata database names can be up to 30 characters long. However, MLoad restrictions for Teradata releases up to Version 12 require that TDPIDs be eight (8) characters or less. Event Replicator Target Adapter current logic dynamically determines the TDPID using the first eight characters of the name specified in this field. If you do not supply a name, the first eight characters of the Teradata database name is used. Whatever TDPID is used must also be present in the HOSTS file.
- In the Loader options area, provide a value for the **Path to Loader Executable** field, identifying the full path of the Teradata MLoad utility executable. For example:

```
C:\Program Files\NCR\teradata client\bin\mload.exe
```

Save the target processing option definition with the same name as the target definition you created in Step 1.

3 When all processing options are specified, start initial-state processing, directing the output to the Event Replicator Target Adapter via a GFFT you created (using the Data Mapping Tool or the Adabas Event Replicator Subsystem) for the Teradata database.

For more information on this, read ***Requesting Initial-State Data***, elsewhere in this guide..

The Event Replicator Target Adapter will collect the initial-state records and send them to the Teradata MLoad utility to be loaded to your Teradata database.

Note: Some errors might prevent the MLoad utility from restoring primary and foreign keys, indexes, and the unique index of tables. We recommend that you examine the MLoad utility log files to determine if such an error occurred. For more information, read ***Processing Output***, next in this section.

Processing Output

The output from this processing includes:

- Various MLoad utility processing files
- MLoad utility log files from the processing
- Data files (*.mlddat* files) containing the initial-state data received by the Event Replicator Target Adapter prior to being processed by the MLoad utility.

All output files are written to the *\logs* subdirectory of your Event Replicator Target Adapter installation.

- Processing Files
- Log Files

Processing Files

The Event Replicator Target Adapter's MLoad utility processing creates three processing files per table loaded. The following processing files are stored in the *\logs* subdirectory of your Event Replicator Target Adapter installation:

- *xxxx.bat* (where *xxxx* is the table name) file: This *.bat* file is used to run the MLoad utility.
- *xxxx.mldctl* (where *xxxx* is the table name) file: This is the control file for the MLoad utility processing.
- One or more *xxxx.mlddat* (where *xxxx* is the table name) files: These files contain the data to be loaded to the database tables by the MLoad utility. These files can take a good deal of space, so make sure that there is enough space to accommodate them and be sure to manually delete them when initial-state processing has completed successfully.

Log Files

Two kinds of log files are produced by Event Replicator Target Adapter MLoad utility processing: a primary log file and multiple processing log files. All log files are stored in the *\logs* subdirectory of your Event Replicator Target Adapter installation.

1. The Teradata MLoad utility primary log files make references to the processing log files. The primary log files have names in the format *xxxx_yymmdd_hhmmss.out*, where *xxxx* is the table name, *yymmdd* is the date of the log file, and *hhmmss* is the time the log file was last updated. We recommend that you review these primary log files carefully to ensure that the MLoad utility had no problems restoring primary and foreign keys, indexes, or the unique index of tables. If MLoad utility processing fails for some reason, we recommend reviewing these primary *.out* log files first when resolving the problem, proceeding then to the processing log files, as necessary.
2. Multiple processing log files are created, one for each MLoad process executed by the Event Replicator Target Adapter. These processing log files have names in the format *Multiloadprocess-yyyy-mm-dd-hh-mm-ss-nnnn.log*, where *yyyy-mm-dd* is the processing date (year, month, and day) and *hh-mm-ss-nnnn* represents the processing time (hours, minutes, seconds, and milliseconds).

17 Requesting RDBMS Data Refreshes

Using the Adabas Event Replicator Subsystem , you can manually submit requests to the Event Replicator Target Adapter that delete the data in all relational database tables associated with the request, but leave the tables in place. This effectively clears the tables. For more information on using the Adabas Event Replicator Subsystem to do this, read *Clearing (Refreshing) the RDBMS Table Data* in *Adabas Event Replicator Subsystem User's Guide*.

18 Dropping RDBMS Tables

Using the Adabas Event Replicator Subsystem, you can manually submit requests to the Event Replicator Target Adapter that delete the data and remove the tables associated with the request from the relational database. For more information on using the Adabas Event Replicator Subsystem to do this, read *Deleting (Dropping) RDBMS Tables* in *Adabas Event Replicator Subsystem User's Guide*.

19 Replicating ADALOD Data to the RDBMS Tables

You can use the LOAD and UPDATE functions of the Adabas ADALOD utility to add and update data in the RDBMS tables. In both cases, you must specify RPLLOAD=YES when running the utility to ensure that the loaded data is automatically replicated.

The processing of ADALOD data for Event Replicator Target Adapter is the same as the processing of regularly replicated data for Event Replicator Target Adapter. In other words, the data that is loaded is processed by subscriptions; if a subscription is activated, includes a generated global format buffer field table (GFFT) for the Adabas file being loaded or updated, and includes an activated destination with a destination class type of "SAGTARG", the data is replicated and sent as an XML message to the Event Replicator Target Adapter for processing.

Before you decide to add data to the RDBMS tables using the ADALOD utility, consider the following potential problems:

- All records processed by ADALOD and sent to the Event Replicator Target Adapter with action types of "Insert" are treated by the Event Replicator Target Adapter as if they were new records, unless the RESEND flag associated with the incoming data from Event Replicator for Adabas has been set to YES. If the RESEND flag is set to YES, checks are performed to determine whether a record already exists with the same primary key. If it does, the existing row is deleted and a new one is inserted for the new incoming data.

  Note: If, during Insert processing, the RESEND flag is set to NO and a record with the same primary key as an existing record is found in the RDBMS tables, Event Replicator Target Adapter processing will fail.

- Performance problems can occur due to the massive numbers of records that might be sent if you use the ADALOD utility and the Event Replicator Target Adapter to add records to the RDBMS tables. A single subscription activated for Event Replicator Target Adapter processing can update multiple RDBMS tables simultaneously (depending on the number of pertinent SFILE definitions and the number of different generated global format buffer field tables (GFFTs)

that exist in the subscription). If multiple subscriptions are activated for Event Replicator Target Adapter processing, then, the amount of traffic on your system could be quite high.

To limit the number of ADALOD records replicated through the Event Replicator Target Adapter, you can deactivate some of the subscriptions, use transaction filter definitions, or elect not to drop an RDBMS table prior to the ADALOD utility run (in this case, all insert transaction will fail if the primary key already exists for that table).

20 Managing Event Replicator Target Adapter Log Files and Console Messages

- Log File Location 94
- Reviewing the sqlrep.log File 94
- sqlrep.log File Retention 94
- Setting Log Levels 95

Log File Location

All Event Replicator Target Adapter log files are stored in the `\logs` subdirectory of your Event Replicator Target Adapter installation. The Event Replicator Target Adapter log file name is *sqlrep.log*, which contains all current log messages. If you have a problem, please keep this log file.

In addition, the following additional log files may be found at this location. If you have a problem, please keep these log files as well:

- Statistics about the Event Replicator Target Adapter session are stored in log files with names in the format: `ArtResource_yyyymmdd_hhmmss.log`, where the date and time are used to make unique log files for each session.
- Metrics up to the last transaction processed during the session are stored in log files with names in the format: `ArtMetrics_yyyymmdd_hhmmss.log`, where the date and time are used to make unique log files for each session.
- If you requested that IBM EntireX or WebSphere MQ tracing occur when you configured the source definitions for Event Replicator Target Adapter, the *etb_trace.log* file (for IBM EntireX) or the *mqs_trace.log* file (for WebSphere MQ) will be created at this location.
- If you requested that all incoming Event Replicator messages be captured for debugging purposes when you configured the source definitions for Event Replicator Target Adapter, an *etb_capture.log* file (for IBM EntireX) or an *mqs_capture.log* file (for WebSphere MQ) will be created at this location.

Reviewing the sqlrep.log File

On Windows systems, the *sqlrep.log* file can be reviewed using a text editor. On UNIX systems, it can be reviewed by running the *sqlreplog.sh* script. You can control the level of log messages written to the *sqlrep.log* file; read ***Setting Log Levels***, elsewhere in this section, for more information.

sqlrep.log File Retention

Each time you start up the Event Replicator Target Adapter, a new *sqlrep.log* file is created. The previously-existing log file will be renamed *sqlrep.log.1*. All previous log files are also retained and renamed, so that the lowest numbered file is the most recent historical log file (the *sqlrep.log* file with no number is always the current log file). In other words, if an *sqlrep.log* and an *sqlrep.log.1* file already exist when the Event Replicator Target Adapter is started, the following sequence of events occurs:

1. File *sqlrep.log.1* is renamed to *sqlrep.log.2*.

2. File *sqlrep.log* is renamed to *sqlrep.log.1*.
3. A new *sqlrep.log* file is created for the new Event Replicator Target Adapter session.

When an *sqlrep.log* file reaches 10MB in size, it is renamed *sqlrep.log.1*, and all previous log files are retained and renamed as well, as described above.

A maximum of 99,999 log files can be retained. When this maximum is reached and a new *sqlrep.log* file is created, the *sqlrep.log.99999* file is overwritten with data from the *sqlrep.log.99998* log file.

Setting Log Levels

You can set the log level of messages written to the *sqlrep.log* log file, the Windows and UNIX event logs, and to the console. Three log levels are available (listed in order of settings that result in the most logging to those resulting in the least logging):

Log Level	Description
debug	Only debugging, informational, warning, error, and fatal messages are logged or displayed. **Note:** Setting this log level can impact performance.
info	Only informational, warning, error, and fatal messages are logged or displayed.
error	Only error and fatal messages are logged or displayed.

To change the log level of messages sent to the console:

1. Edit the *log4j2.xml* file. This file is located in the *\webapps\sqlrep\WEB-INF\classes* directory wherever you installed Event Replicator Target Adapter.
2. To produce a debug log, edit the following line as indicated:

```
<Loggers>
  <Logger name="com.softwareag.adabas.jas" level="error" additivity="false">
    <AppenderRef ref="File"/>
  </Logger>
  <Logger name="com.softwareag.adabas.xts" level="error" additivity="false">
    <AppenderRef ref="File"/>
  </Logger>
 <Logger name="com.softwareag.ada.reptor" level="debug" additivity="false"> ↩
 ↩
```

Tip: Change “info” to “debug” on the last line to activate debug trace.

```
<AppenderRef ref="File"/>
      <AppenderRef ref="Console" level="info"/>  ↩
```

Tip: Leave this at “info” to avoid debug information in the console. Set this to “debug” only if you want debug trace in the console.

```
</Logger>
    <Root level="info">
      <AppenderRef ref="Console"/>
    </Root>
  </Loggers>
```

Note: Setting the log level to DEBUG can impact performance.

21 Controlling Event Replicator for Adabas Destination Definitions

Event Replicator Target Adapter provides two commands (batch utilities) you can use to activate, deactivate, open, or close a destination definition in the Event Replicator for Adabas.

These commands are located in the *bin* directory, wherever the Event Replicator Target Adapter source code and program files are installed.

- Use the `etbdest` command to control your IBM EntireX destination definitions.
- Use the `mqsdest` command to control your WebSphere MQ destination definitions.

From the Target Adapter Command Prompt, enter the appropriate command, using one of the following syntaxes:

```
etbdest -bbrokerID -sclass/server/service -ooperation -ddestination-name -ttracelevel ↩
(-mscm -uuserid -kusertoken -ppassword)
```

or

```
mqsdest  -hhost -pport -mqueuemgr -cchannel -qqueuename -ooperation -ddestinationname ↩
-ttracelevel
```

The parameters in this syntax can be specified in any order because each parameter is identified by a parameter identifier. Note that there are no spaces between the parameter identifier and the value of the parameter, although a single space separates each parameter. Each parameter identifier is described in the following table:

Parameter Identifier	Specify
-b	The IBM EntireX Broker ID.
-c	The case-sensitive WebSphere MQ channel, defined as an Event Replicator Input Queue. The mqdest message will sent to this queue.
-d	The name of the Event Replicator for Adabas Destination you want to control.
-h	The WebSphere MQ host name.
-k	User Token if required for Single Conversation Mode communication with Broker.
-m	(etbdest) Mode. Follow with 'scm' (no quotes) if the input queue is configured for Single Conversation Mode.
-m	(mqsdest) The case-sensitive WebSphere MQ queue manager name.
-o	The controlling operation you want to issue for the destination. Valid operation are OPND (open the destination) and CLSD (close the destination).
-p	(etbdest) Password if Broker is running with security.
-p	The WebSphere MQ port number.
-q	The case-sensitive WebSphere MQ queue name.
-s	The case-sensitive class, server, and service name (separated by slashes) of the IBM EntireX Broker, defined as an Event Replicator Input Queue. The etbdest message will sent to this queue.
-t	The trace level that should be used for the attempt. Valid values range from "0" (no trace) through "5". The default is "0".
-u	Broker userid if required for Single Conversation Mode.

If errors occur when you run *etbdest* or *mqsdest*, rerun them using a trace level setting of "3". Then submit the trace log to your IBM EntireX or WebSphere MQ administrators to identify the problem. For more information about Event Replicator Target Adapter log files, read ***Managing Event Replicator Target Adapter Log Files and Console Messages***, in the *Event Replicator Target Adapter User Guide*.

22 Samples

- Generated Field Table (GFFT) 100
- Schema and Replicated Data 101

The sample in this chapter shows a generated field table (GFFT) and its corresponding XML schema.

Generated Field Table (GFFT)

The field table is created when you generate a global format buffer for a subscription in Event Replicator for Adabas. It is generated from Predict file definitions. To generate a global format buffer and field table (GFFT) using the Adabas Event Replicator Subsystem, read *Generating a GFB* in *Adabas Event Replicator Subsystem User's Guide*.

```
I T L DB Name                              F       Leng S D Remark
- - - -- -------------- top ------------- - ---------- - - -----------

    1 AA PERSONNEL-ID                      A          8   D
  G 1 AB FULL-NAME
    2 AC FIRST-NAME                        A         20 N
    2 AD MIDDLE-I                          A          1 N
    2 AE NAME                              A         20   D
    1 AF MAR-STAT                          A          1 F
  *      M=MARRIED
  *      S=SINGLE
  *      D=DIVORCED
  *      W=WIDOWED
    1 AG SEX                               A          1 F
  G 1 A1 FULL-ADDRESS
    2 AJ CITY                              A         20 N D
    2 AK ZIP                               A         10 N
    2 AL COUNTRY                           A          3 N
  G 1 A2 TELEPHONE
    2 AN AREA-CODE                         A          6 N
    2 AM PHONE                             A         15 N
    1 AO DEPT                              A          6   D
    1 AP JOB-TITLE                         A         25 N D
  G 1 A3 LEAVE-DATA
    2 AU LEAVE-DUE                         N        2.0
    2 AV LEAVE-TAKEN                       N        2.0 N
```

Schema and Replicated Data

The XML schema for the data is created from the generated field table (GFFT) for the subscription that processed the replicated data. It is provided in the Create operation in the sample below. Insert, Update, and Delete operations for the replicated data are also shown in the sample. The Create operation will create the SQL table in the RDBMS, the Insert operation will insert replicated data into the created SQL table, the Update operation will update replicated data in the SQL table, and the Delete operation will delete replicated data in the SQL table. The name of the table created is determined by combining the name of the subscription (PAYROLL in this sample) with the name of the schema file (EMPLOYEES in this sample): PAYROLL_EMPLOYEES. The names of the columns in the table are taken from the long names in the field table.

Note: If no primary key specified, the ISN will be used as the primary key. The primary key must be defined as an Adabas Unique Descriptor.

```
EXAMPLE "Create":

<TransactionPart Sub='SUB1' Origin='I' Prefix='N' Part='1' ↩
Id='C0CF513C08C1F20400000000' En='4091'>
 <Event Row='1' File='EMPLOYEES-FILE' Op='Create' Time='2007/06/27-15:20:42.444665' ↩
FNR='10'>
  <Field name='PERSONNEL-ID' ei='1' fmt='string' prc='0' len='8' key='uk'/>
  <Field name='FIRST-NAME' ei='2' fmt='string' prc='0' len='20'/>
  <Field name='NAME' ei='3' fmt='string' prc='0' len='20' key='ky'/>
  <Field name='MIDDLE-NAME' ei='4' fmt='string' prc='0' len='20'/>
  <Field name='MAR-STAT' ei='5' fmt='string' prc='0' len='1'/>
  <Field name='SEX' ei='6' fmt='string' prc='0' len='1'/>
  <Field name='BIRTH' ei='8' fmt='date' prc='0' len='10' key='ky'/>
  <Field name='ADDRESS-LINE' ei='9' fmt='string' prc='0' len='20' att='mu'/>
  <Field name='CITY' ei='10' fmt='string' prc='0' len='20' key='ky'/>
  <Field name='POST-CODE' ei='11' fmt='string' prc='0' len='10'/>
  <Field name='COUNTRY' ei='12' fmt='string' prc='0' len='3'/>
  <Field name='AREA-CODE' ei='13' fmt='string' prc='0' len='6'/>
  <Field name='PHONE' ei='14' fmt='string' prc='0' len='15'/>
  <Field name='DEPT' ei='15' fmt='string' prc='0' len='6' key='ky'/>
  <Field name='JOB-TITLE' ei='16' fmt='string' prc='0' len='25' key='ky'/>
  <Field name='CURR-CODE' ei='18' fmt='string' prc='0' len='3' att='pe' ↩
gname='INCOME'/>
  <Field name='SALARY' ei='19' fmt='decimal' prc='0' len='10' att='pe' ↩
gname='INCOME'/>
  <Field name='BONUS' ei='20' fmt='decimal' prc='0' len='10' att='mu' gname='INCOME'/>

  <Field name='LEAVE-DUE' ei='21' fmt='decimal' prc='0' len='2'/>
  <Field name='LEAVE-TAKEN' ei='22' fmt='decimal' prc='0' len='2'/>
  <Field name='LEAVE-START' ei='24' fmt='decimal' prc='0' len='8' att='pe' ↩
gname='LEAVE-BOOKED'/>
  <Field name='LEAVE-END' ei='25' fmt='decimal' prc='0' len='8' att='pe' ↩
```

```
gname='LEAVE-BOOKED'/>
  <Field name='LANG' ei='26' fmt='string' prc='0' len='3' att='mu'/>
    <Alias name='INCOME' ei='17'/>
    <Alias name='LEAVE-BOOKED' ei='23'/>
 </Event>
 <TransactionEnd RowCount='1' DBID='44444' Sent='2007/06/27-15:20:58.782005' ↩
Committed='2007/06/27-15:20:58.434591' Version='1.0'/>
</TransactionPart>

EXAMPLE "Insert":

<TransactionPart Sub='SUB1' Origin='I' Prefix='N' Part='1' ↩
Id='C0D5D7140B13D3C900000000' En='4091'>
 <E Row='1' File='EMPLOYEES-FILE' ISN='1' Op='Insert' FNR='10'>
  <F ei='1' ai='50005800' key='uk'/>
  <F ei='2' ai='SIMONE'/>
  <F ei='3' ai='ADAM' key='ky'/>
  <F ei='5' ai='M'/>
  <F ei='6' ai='F'/>
  <F ei='8' ai='1952/01/30' key='ky'/>
  <F ei='9' ai='26 AVENUE RHIN ET DA' att='mu' ix='1'/>
  <F ei='10' ai='JOIGNY' key='ky'/>
  <F ei='11' ai='89300'/>
  <F ei='12' ai='F'/>
  <F ei='13' ai='1033'/>
  <F ei='14' ai='44864858'/>
  <F ei='15' ai='VENT59' key='ky'/>
  <F ei='16' ai='CHEF DE SERVICE' key='ky'/>
  <F ei='18' ai='FRA' att='pe' ix='1' gei='17'/>
  <F ei='19' ai='159980' att='pe' ix='1' gei='17'/>
  <F ei='20' ai='23000' ix='1' ix2='1' att='mu' gei='17'/>
  <F ei='21' ai='19'/>
  <F ei='22' ai='5'/>
  <F ei='24' ai='19990801' att='pe' ix='1' gei='23'/>
  <F ei='25' ai='19990831' att='pe' ix='1' gei='23'/>
  <F ei='26' ai='FRE' att='mu' ix='1'/>
  <F ei='26' ai='ENG' att='mu' ix='2'/>
 </E>
 <TransactionEnd RowCount='10' DBID='44444' Sent='2007/07/02-19:51:44.051797' ↩
Committed='2007/07/02-19:51:43.712573' Version='1.0'/>
</TransactionPart>

EXAMPLE "Update":

<TransactionPart Sub='SUB1' Origin='I' Prefix='N' Part='1' ↩
Id='C0D5D7140B13D3C900000000' En='4091'>
 <E Row='1' File='EMPLOYEES-FILE' ISN='10' Op='Update' FNR='10'>
  <F ei='1' bi='50004000' ai='11111111' key='uk'/>
  <F ei='20' bi='8880' ai='0000' ix='1' att='mu' ix2='1' gei='17'/>
  <F ei='18' bi='USD' ai='TWN' att='pe' ix='31' gei='17'/>
  <F ei='26' bi='ENG' ai='CHN' att='mu' ix='2'/>
 </E>
```

```
 <TransactionEnd RowCount='1' DBID='44444' Sent='2007/07/08-19:51:44.051797' ↩
Committed='2007/07/08-19:51:43.712573' Version='1.0'/>
</TransactionPart>

EXAMPLE "Delete":

<TransactionPart Sub='SUB1' Origin='I' Prefix='N' Part='1' ↩
Id='C0D5D7140B13D3C900000000' En='4091'>
 <E Row='1' File='EMPLOYEES-FILE' ISN='10' Op='Delete' FNR='10'>
 </E>
 <TransactionEnd RowCount='1' DBID='44444' Sent='2007/07/09-19:51:44.051797' ↩
Committed='2007/07/09-19:51:43.712573' Version='1.0'/>
</TransactionPart>
```

Index

A

activating processing
  Event Replicator Target Adapter, 9
ADALOD
  replicating data to RDBMS tables, 91
Administration tool
  restarting Event Replicator Target Adapter, 6
  shutting down Event Replicator Target Adapter, 16
  starting Event Replicator Target Adapter, 6

B

batch utilities, 97
bcp utility
  loading initial-state data, 67

D

DAT files, 76
DB2 CONNECT IMPORT command
  loading initial-state data, 61
DB2 database load processing
  files, 65
  output, 64
  steps, 62
DB2 databases
  loading with initial-state data, 61
deactivating
  Event Replicator Target Adapter processing, 13
destination definitions
  controlling from the Event Replicator Target Adapter, 97
dropping RDBMS tables, 89

E

etbdest command, 97
Event Replicator Target Adapter
  activating processing, 9
  controlling Event Replicator for Adabas destination definitions, 97
  controlling webMethods EntireX destination definitions, 97
  controlling WebSphere MQ destination definitions, 97
  deactivating processing, 13
  documentation,
  dropping tables, 89
  loading initial-state data with bcp utility, 67
  loading initial-state data with DB2 CONNECT IMPORT command, 61
  loading initial-state data with Oracle SQL*Loader, 73
  log files, 93
  populating tables with initial-state data, 19
  replicating data to the tables using ADALOD, 91
  requesting RDBMS data refreshes, 87
  restarting, 6
  samples, 99
  shutting down, 15
  shutting down in UNIX, 17
  shutting down in Windows, 16
  starting, 5
  starting in UNIX, 7
  starting in Windows, 7

G

generated field table
  sample, 100

I

initial-state data
  enabling Oracle SQL*Loader processing, 74
  loading data with Oracle SQL*Loader, 73
  loading DB2 databases using the DB2 CONNECT IMPORT command, 61
  loading Microsoft SQL Server databases using the bcp utility, 67
  Oracle SQL*Loader processing DAT files, 76
  Oracle SQL*Loader processing log files, 76
  Oracle SQL*Loader processing output, 75
  populating RDBMS tables, 19

J

JMS support
  description, 21

L

log files
  location, 94
  managing, 93
  Oracle SQL*Loader processing, 76
  retention, 94
  setting log levels, 95
log levels, 95

M

Microsoft SQL Server databases
    loading with initial-state data, 67
mqsdest command, 97

O

Oracle databases
    loading with initial-state data, 73
Oracle SQL*Loader
    enabling initial-state processing, 74
    initial-state processing DAT files, 76
    initial-state processing log files, 76
    initial-state processing output, 75
    loading initial-state data, 73
    prerequisites, 74

R

RDBMS data, 87
RDBMS tables
    ADALOD data as input, 91
    dropping, 89
    populating with initial-state data, 19
    refreshing, 87
restarting
    Event Replicator Target Adapter, 6

S

samples
    Event Replicator Target Adapter, 99
    generated field table, 100
    schema and replicated data, 101
schema
    sample, 101
shutting down
    Event Replicator Target Adapter, 15
sqlrep.log file, 94
starting
    Event Replicator Target Adapter, 5

U

UNIX
    shutting down Event Replicator Target Adapter, 17
    starting Event Replicator Target Adapter, 7

W

webMethods EntireX
    controlling destination definitions, 97
WebSphere MQ
    controlling destination definitions, 97
Windows
    shutting down Event Replicator Target Adapter, 16
    starting Event Replicator Target Adapter, 7