

Open Systems Event Replicator Frequently Asked Questions

- **General Definitions**

Q: What is the Replication Controller?

A: The Replication Controller receives its work orders from the Replication Administrator. Once it gets a replication configuration it can pass the configuration on to the Event Producers on the Adabas nucleus and spawn Replication Engines to handle the work.

Q: What is the Replication Engine (EC)?

A: The Replication Engine takes transactions that have been completed on the source table and replicates the changes to the target database.

Q: What is the Event Producer (EP)?

A: The Event Producer is part of the source database; it sends information about which data has changed on the source table to the target table in a given replication.

Q: What is the latency of changes?

A: The latency of changes is the time it takes for an update/insert/delete on the source database to be reflected in the target database.

- **Requirements**

Q: Does the Open Systems Event Replicator require a special version of Adabas / ACE?

A: Yes, the Open Systems Event Replicator requires both Adabas version 6.1.4 and ACE version 6.6.3, or later for Adabas to Relational replication. For Adabas to Adabas replication, Adabas 6.3.1 or later is required and ACE 11.5 or later is required.

- **Initial State**

Q: What is an initial state?

A: An initial state is the snapshot of the data on the source table taken when the initial replication starts.

Q: What events require an initial state?

A: A new deployment or an Adabas nucleus crash will require an initial state. A normal shutdown of the Adabas nucleus will not require an initial state.

Q: Do I have to initiate an initial state manually or will it happen automatically?

A: When you deploy a data dictionary for replication, an initial state will automatically happen. There are also some error cases (such as an Adabas nucleus crash or an improper shutdown) that cause an automatic initial state. This restores data integrity on the target database.

Q: How long does an initial state take once it has started?

A: The length of time an initial state takes depends on the system and table size. For example, initial states take longer on larger tables than on smaller tables.

Q: What happens if I shut down the Adabas database while replicating? Do I have to take any special action, such as performing an initial state again?

A: If you shut down Adabas, the event replicator does not need any outside intervention. When an Adabas nucleus is shutdown during replication it stops events from being sent to the target database. When the nucleus is started, events will automatically be sent to the target database.

- **Databases**

Q: Is it possible for me to replicate from more than one source databases?

A: Yes, just import all the source databases to the data dictionary before setting up the replications. When the data dictionary is opened by the Replication Administrator, all the databases will be available for replication.

- **Tables**

Q: Will the Open Systems Event Replicator create the tables for me on the target database, or do I need to do that before creating the replications?

A: Sometimes. If the target is an Adabas database, the target table must exist on the Adabas database and be imported into the CDD. Currently, the Open Systems Event Replicator can not create tables on an Adabas database. For other databases, the Replication Administrator will create new tables on the target database. Existing tables can be used in the replications, but that is not a requirement.

Q: Can I control the shape of the tables on the target database?

A: Yes, when replicating to a new table that does not exist, the columns created and the order they are created in can be specified in the Replication Administrator column mapping form.

Q: Can I split a single source table into two separate target tables?

A: Yes, one source table can be sent to multiple target tables. Each of the target tables do not have to contain the same source columns.

Q: Can I use a table as a source table in one replication and a target table in another replication?

A: No, if a table is being used as a target table in one replication, it cannot be used as a source table in another replication. The original source table needs to be replicated to both target tables.

Q: Can I take two or more source tables and merge them into a single target table?

A: No, the current version of the Open Systems Event Replicator can not send data from multiple source tables to one target table.

Q: I am replicating from an Adabas source table to an Adabas target table. My Adabas source table has Mu's, PE's or Mu-PE's. Do I have to do anything special to replicate this Adabas table?

A: Yes, in order to replicate an Adabas source tables that contain either Mu's, PE's or Mu-PE's to an Adabas target table, the Open Systems Event Replicator must use the CONNX generated logical flat table instead of the Adabas source table. When an Adabas source table containing Mu's, PE's or Mu-PE's is imported into CONNX, multiple logical tables are created so the table appears relational to ODBC. Most of these tables are for ODBC access only; they are not appropriate for event replication. One of the generated source table tables is a logical flat table; use this table to replicate data from the Adabas source table with Mu's, PE's or Mu-PE's to the Adabas target table. Both the source and target tables must have the same design.

Q: What is the maximum number of Mu plus PE plus Mu-PE occurrences an Adabase Open Systems Event Replication source table can have?

A: The maximum number of Mu plus PE plus Mu-PE occurrences for an Open Systems Event Replicator Adabas source table is 191. All occurrences must be predefined.

Q: When I replicate from an Adabas table to an Adabas table, why can't I use the ISN as the replication index on the target table?

A: The target table must have a unique index that the replication tool can use to propagate changes. The target table ISN is an automatically generated number that cannot map to any column on the source table.

- **Data Transformations**

Q: Can I perform data transformation during the replication?

A: SQL column functions can be created in the column mapping form and used during replication. The Open Systems Event Replicator does not currently include aggregation functions.

Q: I am replicating from an Adabas Source table to a non-Adabas target table. My tables contain non-character data. Is there anything special I should be aware of?

A: Yes, if your non-character data has a large number of significant digits, you could lose some data. If the number of significant digits in your Adabas source table is greater the maximum number of allowed significant digits in the non-Adabas target table (this depends on the relational database type), the Open Systems Event Replicator will truncate the target data during replication and will not generate an error condition.

- **Replication Engine**

Q: What happens if the Replication Engine is shutdown or taken offline? Will we lose data integrity on the target database?

A: No. If the Replication Engine is shutdown during replication the Event Producer stores transactions that have been completed on the source table. When the Replication engine is restarted the stored transactions will then be processed and sent to the target table. Since the transactions are stored, no data will be lost if the Replication Engine shuts down.

- **Log File**

Q: Does the Open Systems Event Replicator have a log file?

A: Yes, the Open Systems Event Replicator has a log file. It can normally be found in

- **C:\connx32\replication\log** on Windows
- **\$CNXDIR/OSER/log** on UNIX

Q: Can I get notified if an error occurs that causes replication to stop?

A: Yes. The Open Systems Event Replicator has a user-specified job exit. The Event Replicator will run your specified job if event replication stops because of an Nx action code error. The Event Replication Error Messages topic in the User Guide has a list of these errors and a description of how to use the notification exit.

- **Performance**

Q: How many computers are required to use Open Systems Event Replicator? What is the recommended configuration?

A: The Open Systems Event Replicator can run on one computer if both the source and target databases are located on the same computer. The configuration required to successfully deploy and replicate your databases depends on many factors including hardware, software and network limitations.

Q: Will the Open Systems Event Replicator negatively affect the performance of my Adabas database?

A: No, the Event Producer was designed to minimize the interaction with the Adabas nucleus. The Adabas database will continue on after it passes the transaction to the Event Producer. If the Replication Controller, the Replication Engine or the target becomes disabled, the new transactions are stored and Adabas continues to work normally.