



Syniti Replicate

Kafka® Setup Guide

Version 10.1




Table of Contents

Introduction.....	1
Basic Configuration Steps	1
Connection Type.....	1
Setup Summary.....	2
Download and Install .NET Provider	2
Syniti Replicate Versions 9.8.1 and Above.....	2
Syniti Data Replication Versions 9.6.1 to 9.8.0.....	Error! Bookmark not defined.
Steps for Replicating Tables	4
1. Set Up a Source Connection to a Relational Database.....	4
2. Configure the Enable Transactional Replication Wizard	7
3. Set up a Target Connection to Kafka	9
3. Add Table Information to the Target Connection.....	13
4. Define Replications.....	15
Start Replications.....	18
Stop Replications.....	19
Using Kafka with Kerberos.....	19

Introduction

Syniti Replicate allows you to replicate data from relational database tables to Apache Kafka, an open source distributed streaming platform, using:

- **Refresh**, or **Snapshot**, replication: a one-time complete replication from any major relational database source to Kafka as a target, according to replication settings and scripts.
- **Mirroring**, or **Change Data Capture**, replication: a continuous read of changes to the source database that have been recorded in the database server log. Any changes found in the log are applied to Kafka as a target, according to replication settings and scripts.

Kafka's distributed streaming platform allows you to publish and subscribe to streams of records, similarly to a message queue system. All messages in Kafka are sent to a specific *topic*. A Kafka topic is a category or feed name to which records are published. In order to replicate data in Syniti Replicate to Kafka, the data is organized as virtual tables (with columns and data types). You can create a target table and define a replication to a Kafka "target table", applying mappings and expression logic as with other relational data sources. The Kafka message is built using a JSON, AVRO, CSV or XML serialization of the entire record, including metadata information such as the actual table name and the operation (Insert, Update or Delete), transactional information like transaction ID and timestamp, and finally the entire image (before and after) of the record. Every replication must be associated with a Kafka topic so that Syniti Replicate can determine to which Kafka topic to publish the message.

You can control the timing of the replication, identify the columns to be replicated and add scripts to transform data during replication. Source databases include Oracle, Microsoft SQL Server, IBM Db2 for i, IBM Db2 LUW, Sybase, Informix, MySQL.

Basic Configuration Steps

Use Syniti Replicate Management Center to:

- Create source connections to RDBMS tables
- Create Kafka targets
- Map RDBMS sources to Kafka targets
- Enable replication

Subsequent data management on the Kafka side depends upon your application needs.

Connection Type

Kafka version 2.1.0 and later	Syniti Replicate 9.8.1 or later	Confluent.Kafka version 1.7.0
	Syniti Replicate 9.6.3 to 9.8.0	Confluent.Kafka.StrongName 1.0.0 or 1.1.0
	Syniti Replicate 9.6.1	Confluent.Kafka.StrongName 1.0.0
Confluent Cloud	Syniti DR 9.6.3 or later	Confluent.Kafka.StrongName 1.1.0.

The provider can be downloaded via:

Syniti Replicate

Search for “Confluent Kafka”, selecting the appropriate version according to the table above. See for more detail on installing the provider.

Setup Summary

This section provides a summary of all the steps required for setting up and using Syniti Replicate. Use the link for each step for more information.

Download and Install Syniti Replicate	The Knowledge Platform Product Suites article acts as a hub to point to various resources. To download and/or register Syniti Replicate, log in to the support site, then click the relevant link in the Replicate section of the article. <ul style="list-style-type: none">• Syniti Knowledge Base• Enter a generic support ticket	
Install .NET Provider for Target Database (See Download and Install .NET Provider below)	1.	2. Go to 3. Search 4. Download 5. Install
Syniti Replicate Setup (See Steps for Replicating Tables below)		In the Syn 1. In the 2. Create field. 3. Create
Start Replicating (See Start Replications below)		In the Syn 1. Start

Download and Install .NET Provider

Syniti Replicate Versions 9.8.1 and Above

For Syniti Replicate versions 9.8.1 and above, download and install the Confluent .NET provider for Kafka as follows. Note that ALL steps below must be completed for any replication to Kafka, whether or not you are using AVRO serialization.

1. Go to and search for “Confluent.Kafka”.
2. Download the following packages:
 - Confluent.Kafka 1.7.0
 - Confluent.SchemaRegistry
 - Confluent.SchemaRegistry.Serdes.Avro

Syniti Replicate

The easiest way to do it is by clicking the Downloads menu on top of the page and get the nuget.exe command line tool. The steps below assume you are using nuget.exe.

3. From a dos prompt with admin permissions, execute the following command.

```
nuget install Confluent.Kafka -Version 1.7.0 -Framework net480 -OutputDirectory "C:\Temp\Kafka 1.7.0"
```

The command will download the entire provider into the specified output directory.

4. Open this directory to view the folders where the assemblies are located:
5. Collect the following assemblies selecting them from the shown path:
 - \Kafka 1.7.0\Confluent.Kafka.1.7.0\lib\net46\Confluent.Kafka.dll
 - \Kafka 1.7.0\librdkafka.redist.1.7.0\runtimes\win-x64\native*. * or \Kafka 1.7.0\librdkafka.redist.1.7.0\runtimes\win-x86\native*. *
 - \Kafka 1.7.0\System Buffers.4.4.0\lib\netstandard2.0\System Buffers.dll
 - \Kafka 1.7.0\System.Memory.4.5.0\lib\netstandard2.0\System.Memory.dll
 - \Kafka 1.7.0\System.Numerics.Vectors.4.4.0\lib\net46\System.Numerics.Vectors.dll
 - \Kafka 1.7.0\System.Runtime.CompilerServices.Unsafe.4.5.0\lib\netstandard2.0\System.Runtime.CompilerServices.Unsafe.dll

6. Copy these files into the Plugins/Kafka folder.
7. From a dos prompt with admin permissions, install the package Confluent.SchemaRegistry version 1.7.0:

```
nuget install Confluent.SchemaRegistry -Version 1.7.0 -Framework net480 -OutputDirectory "C:\Temp\SchemaRegistry 1.7.0"
```

8. Copy the following assembly to the Plugins/Kafka folder:
 - \SchemaRegistry 1.7.0\Confluent.SchemaRegistry.1.7.0\lib\netstandard2.0\Confluent.SchemaRegistry.dll

9. From a dos prompt with admin permissions, install the package Confluent.SchemaRegistry.Serdes.Avro version 1.7.0:

```
nuget install Confluent.SchemaRegistry.Serdes.Avro -Version 1.7.0 -Framework net480 -OutputDirectory "C:\Temp\SchemaRegistry.Serdes.Avro 1.7.0"
```

10. Collect the following assemblies, selecting them from the shown path:
 - \SchemaRegistry.Serdes.Avro 1.7.0\Apache.Avro.1.10.2\lib\netstandard2.0\Avro.dll
 - \SchemaRegistry.Serdes.Avro 1.7.0\Confluent.SchemaRegistry.Serdes.Avro.1.7.0\lib\netstandard2.0\Confluent.SchemaRegistry.Serdes.Avro.dll
 - \SchemaRegistry.Serdes.Avro 1.7.0\Newtonsoft.Json.10.0.3\lib\net45\Newtonsoft.Json.dll
 - \SchemaRegistry.Serdes.Avro 1.7.0\System.CodeDom.4.4.0\lib\net461\System.CodeDom.dll

Syniti Replicate

11. Copy these files into the Syniti Replicate install folder Plugins/Kafka folder.

Steps for Replicating Tables

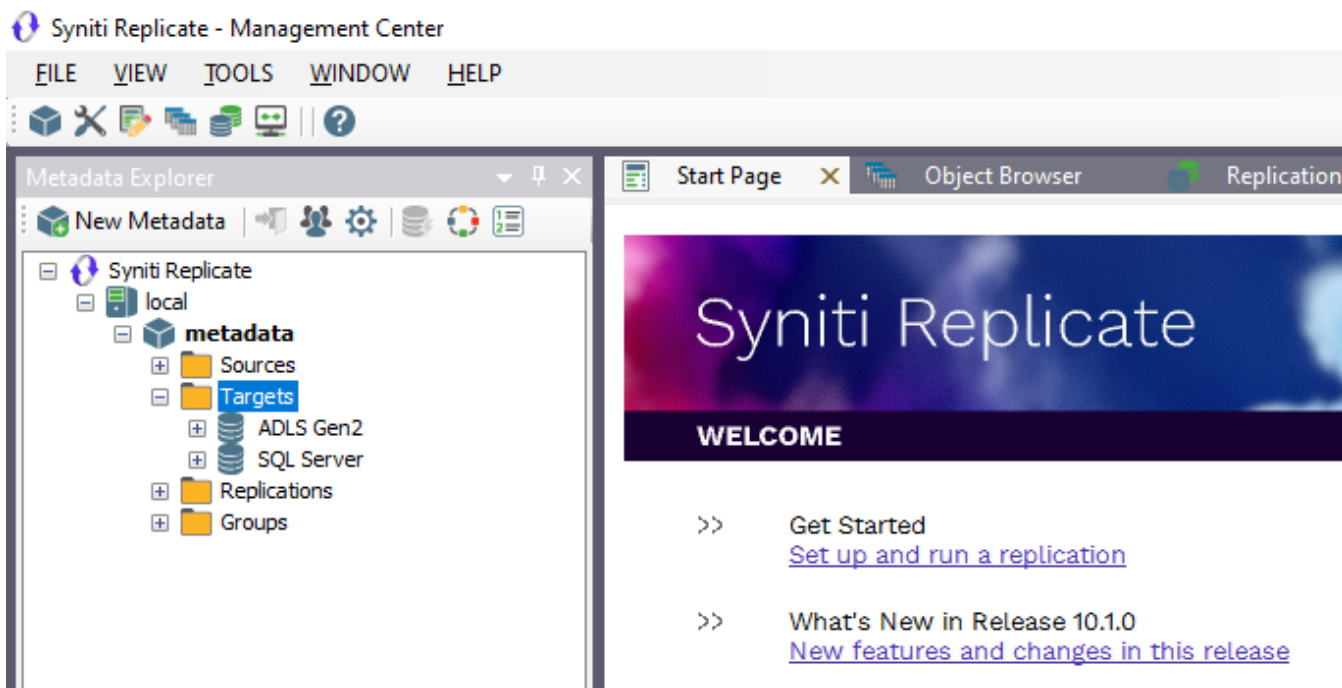
Syniti Replicate support for replicating relational data to Kafka allows you to set up replications using either **Refresh** or **Mirroring** modes.

The steps below explain how to replicate data from a relational database to the Kafka environment. Check for the latest list of supported databases.

1. Set Up a Source Connection to a Relational Database

1. Make sure you have database connections via a .NET data provider to your source database. For each database you are planning to use in your replication project:
 - Install and configure your .NET Provider.
 - From the provider, test the connection to the database.
 - Create a connection string for the data access product/database you are using. Check the documentation for the data access product for information on how to do this.
 - Check that the user ID you are planning to use has sufficient permissions to complete all operations in Syniti Replicate. Contact Syniti Technical Support for specific requirements for your database.
2. Start the Management Center.

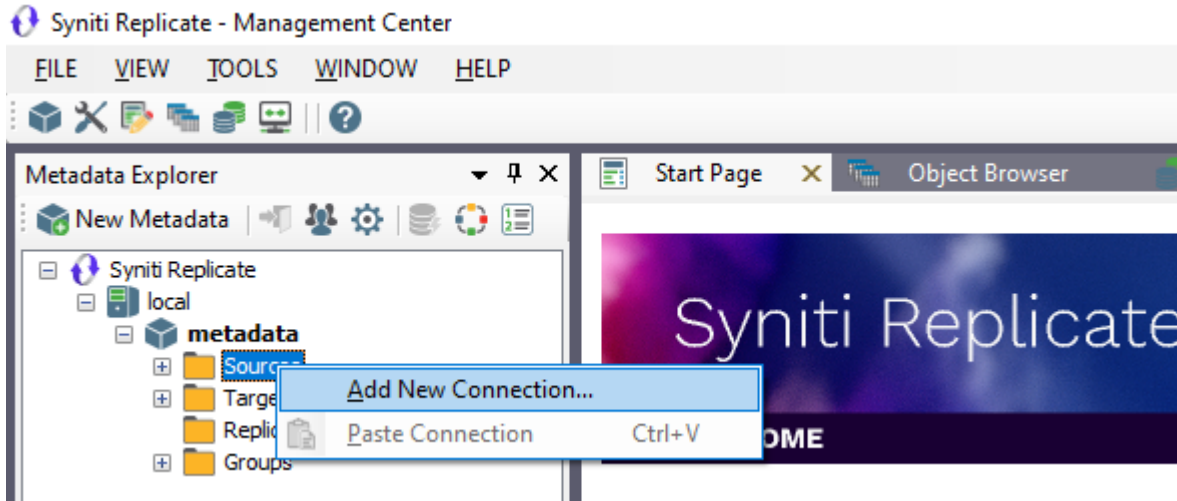
Syniti Replicate provides a default database (Microsoft SQL Server CE) for your metadata, all the information that Syniti Replicate needs to store about your replication setup.




Copyright © 2022 BackOffice Associates, LLC d/b/a Syniti. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

Syniti Replicate

3. In the Metadata Explorer, expand the metadata node to view the **Sources** and **Targets** nodes.
4. Select the **Sources** node.
5. From the right mouse button menu, choose **Add New Connection**.



6. In the Source Connection Wizard, follow steps to add a connection string and test the connection to the database.
7. Check the User's Guide **Database Access Providers and Supported Databases** page to verify what value to enter in the **Assembly** field.



8. In the **Set Connection String** screen, fill out the following fields:

Connection Properties

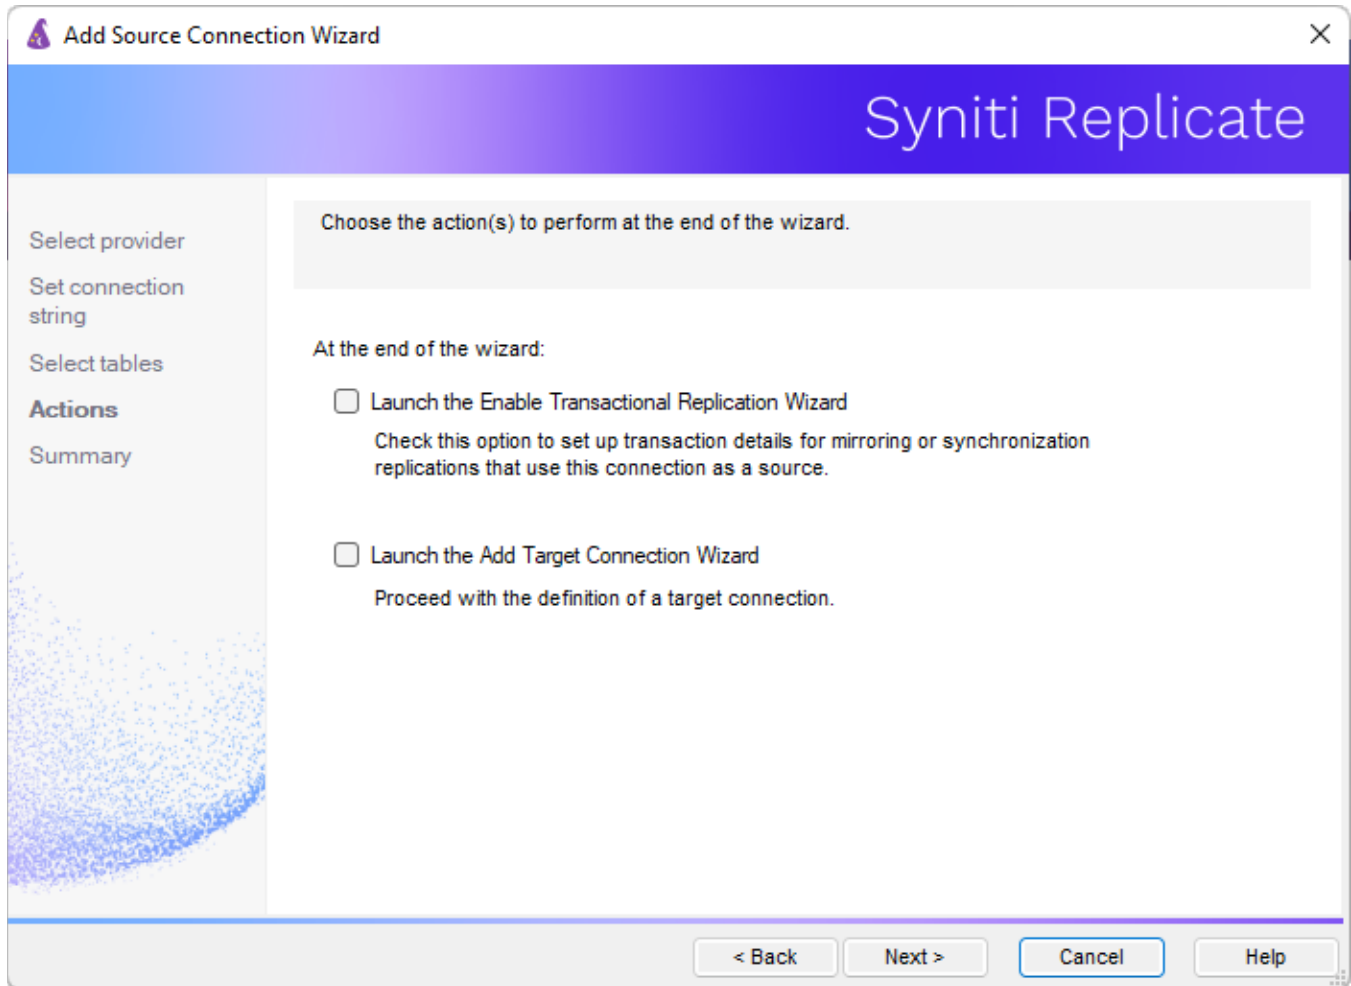
Edit at least the **Required** connection properties by clicking in the property value field and typing a new value.

Note: If using an Oracle 11 client, in the Data Source field, be sure to add the **Oracle Service Name** after the IP address.

9. In the **Setup Info** screen, click **Next** without completing any fields. This screen is used specifically for transactional replications, not refresh or snapshot replications.
10. Choose the tables that you plan to replicate.
11. If using Refresh mode to replicate data to Kafka, in the **Actions** screen, check the option **Launch the Add Target Connection Wizard**.

-OR-

If using Mirroring mode to replicate data to Kafka, in the **Actions** screen, check the option **Launch the Enable Transactional Replication Wizard**.



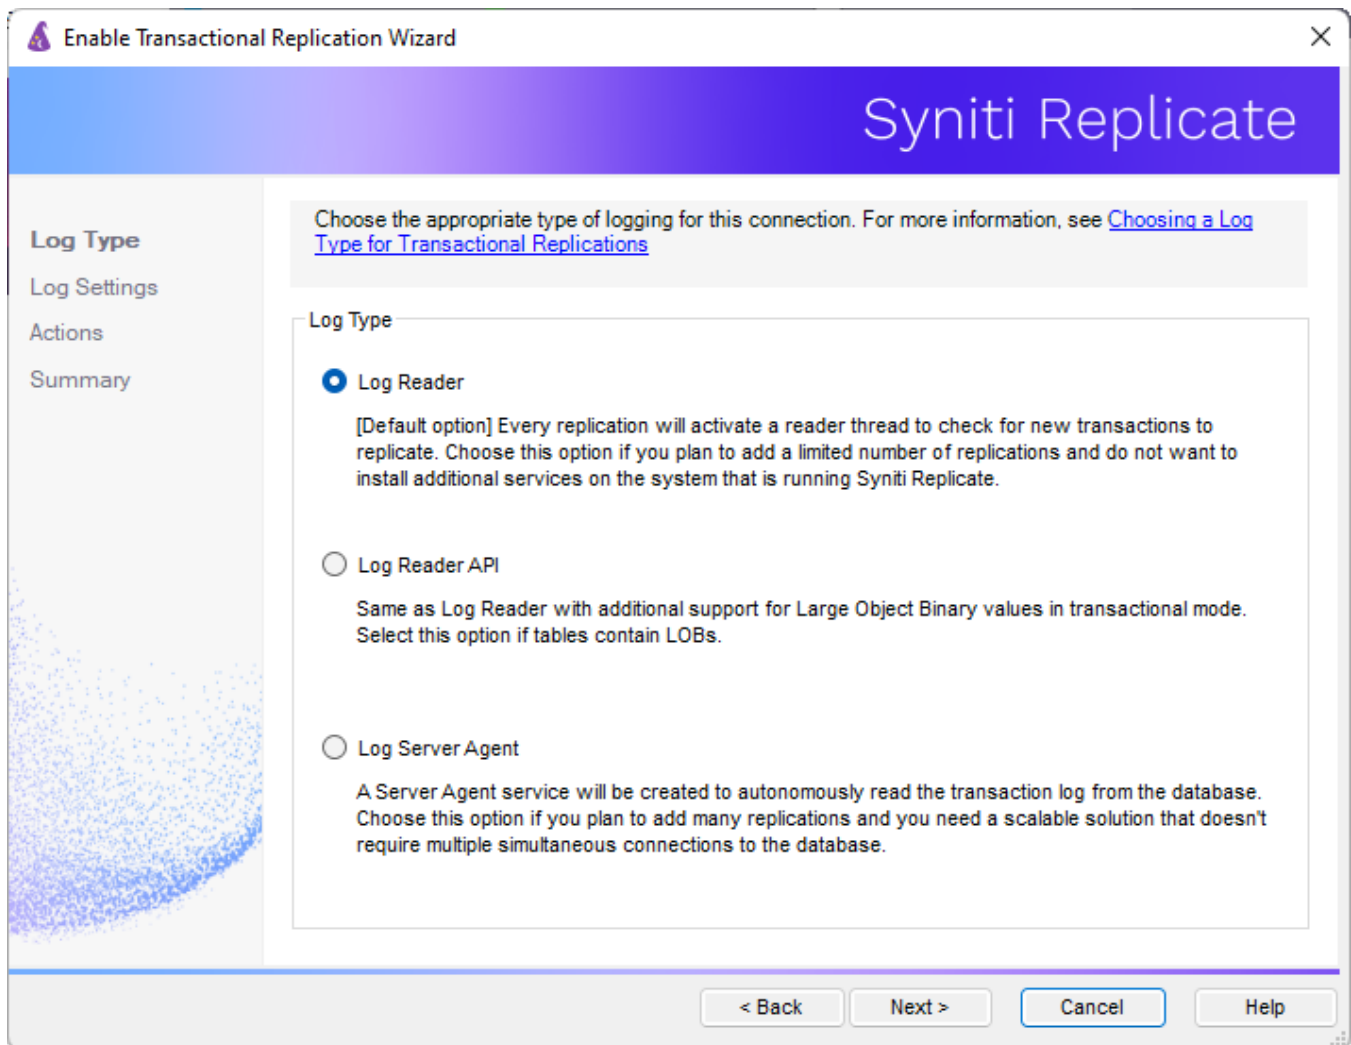
2. Configure the Enable Transactional Replication Wizard

Skip this section if you are setting up a Refresh replication. Continue here if you are setting up a Mirroring replication.

This section assumes you have checked the Source Connection wizard option to launch the Enable Transactional Replication wizard. To open the wizard from the Management Center, choose the connection in the Metadata Explorer, then right-click to choose Transactional Setup > Enable...

In the **Enable Transactional Replication** wizard:

1. Select the type of transactional replication to use. The options depend on the source database and can include Log Reader, Log Server Agent, Triggers, plus Log Reader API (for IBM Db2 for i only.)



2. Click **Next** to enter the log settings. The fields and appropriate values depend on the database and log type. Submit a request in the to obtain a Setup Guide for your database.

All Log Server Agent setups require the following Agent Settings to establish details for the Windows service and log files:

Log Files Folder

An existing folder where all the intermediate binary logs files are written. The folder is also used to contain trace files.

Prefix

A prefix for all files created in the folder. This provides easy identification and management of files associated with your connection. The primary purpose for the prefix is to support the case when you configure more than one connection to use the same folder. In general, however, it is advisable to use different folders for different connections.

Log File Size

Syniti Replicate

The maximum size of each binary log file in megabytes.

Keep Max Files

The maximum number of binary log files to keep. Combined with the log file size, this number needs to be large enough to make sure that all the files are read by Syniti Replicate before they get deleted.

Trace

Unchecked by default. When checked, enables tracing for diagnostics and problem reporting. Trace files are saved in the Log Server folder.

Windows Service Name

Prefix:

This value is provided and cannot be changed. It allows you to easily identify the service in the Microsoft Windows Services tool.

Name:

Specify a unique name for the Windows service. Each connection runs its own instance of the Log Server as a Windows service.

Start service after completing the wizard:

When checked, starts the service automatically after you click **Finish** to complete the wizard.

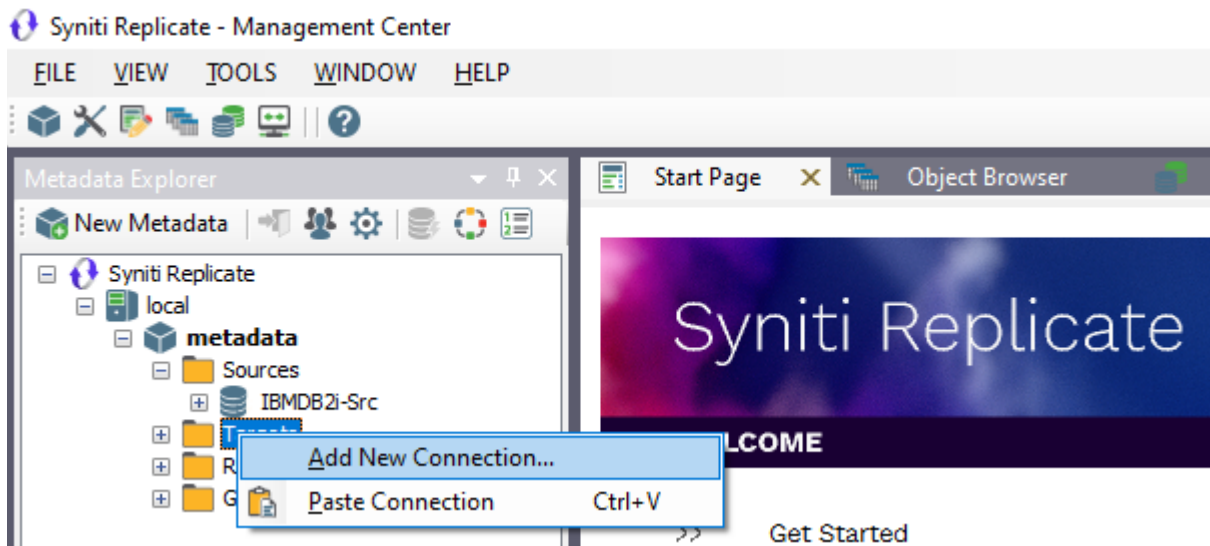
3. Click **Next** to verify your settings against the source connection to the database. If any information is missing, you will not be able to proceed.
4. In the **Actions** screen, check the option to launch the Add Target Connection wizard.
5. Click **Next** to review your changes.
6. Click **Finish** to complete the wizard.

The source connection is now set up for transactional replications.

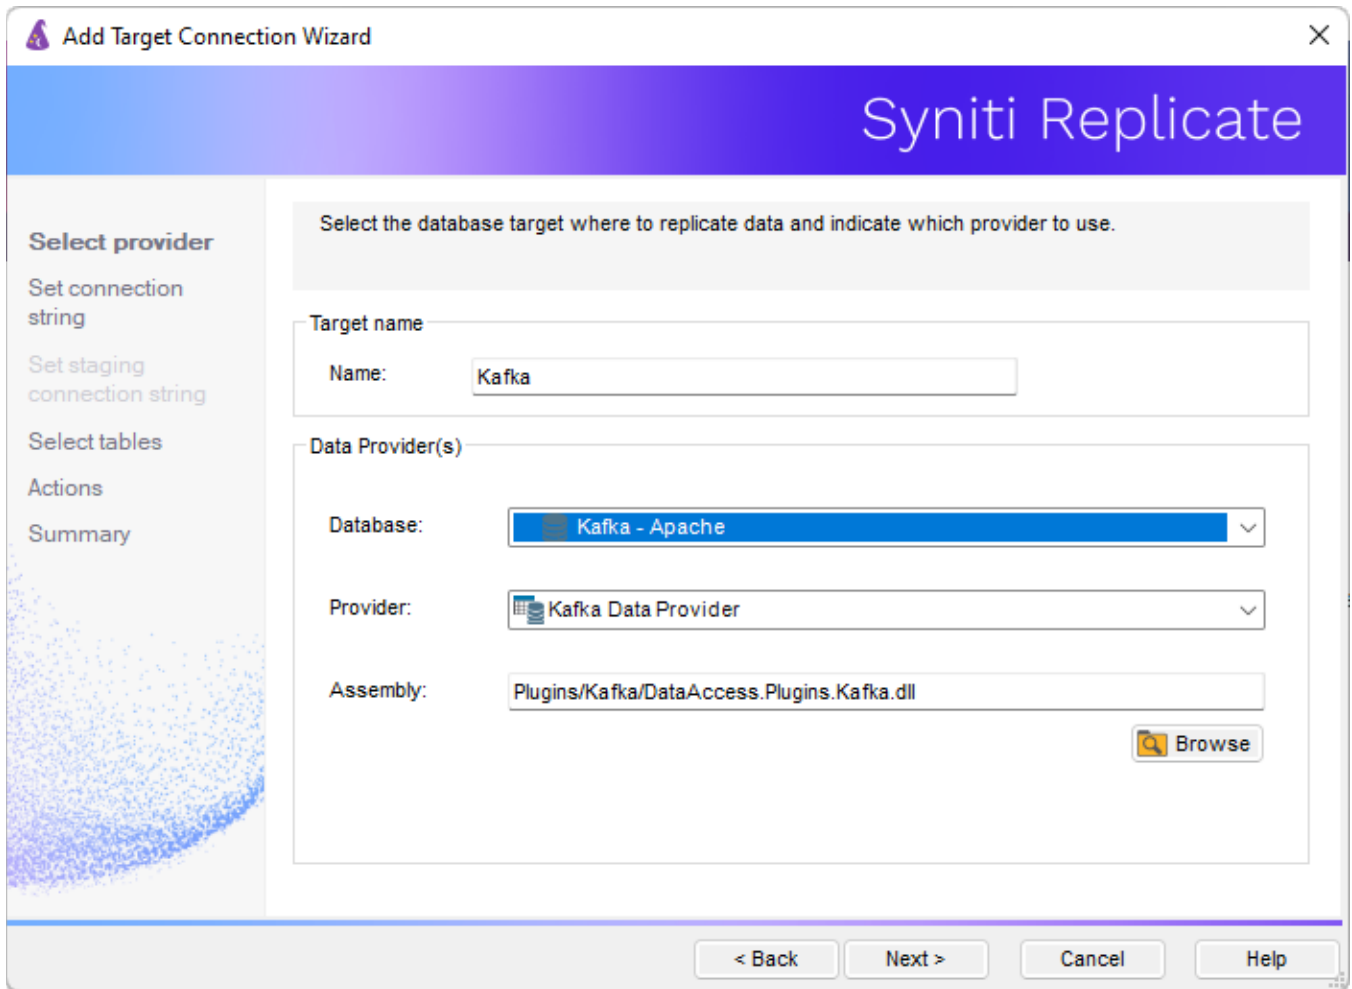
3. Set up a Target Connection to Kafka

1. Select the **Targets** node.
2. From the right mouse button menu, choose **Add New Connection**.

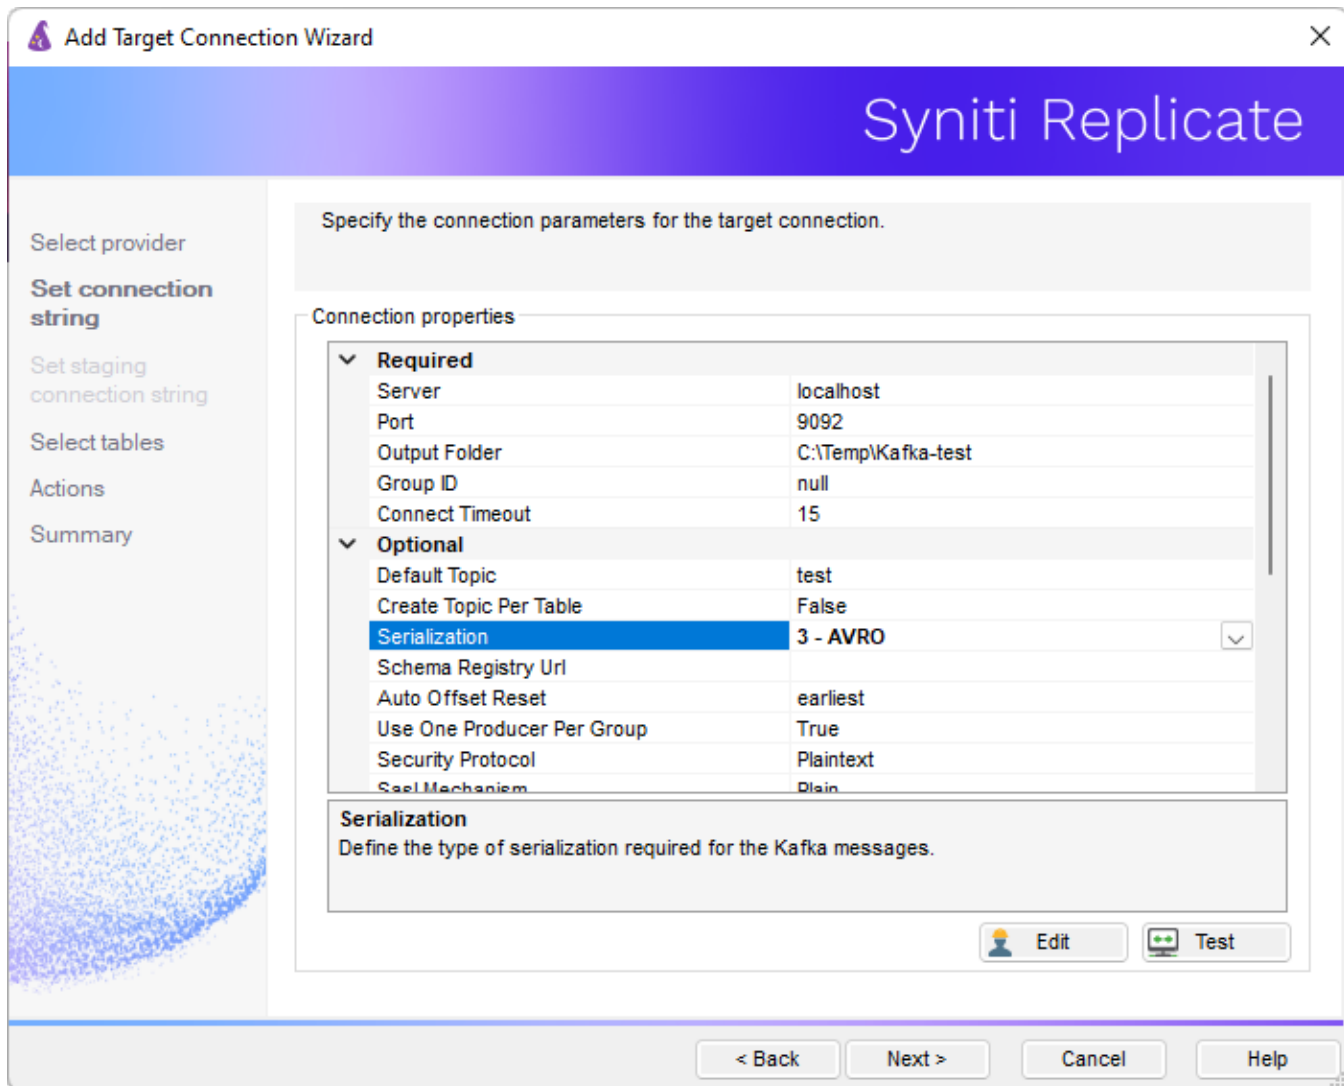
Syniti Replicate



3. In the Add Target Connection Wizard **Database** field, select the Apache Kafka option. The **Provider** and **Assembly** fields are automatically filled out for you.



4. In the **Set Connection String** page, set properties as described in the table below.



Server	The server name or IP address of the system running Kafka.
Port	The port number of the server running Kafka.
Output Folder	The schema name and location to hold config files for the Kafka objects.
Group ID	Kafka property not used by Syniti Replicate. A string that uniquely identifies the group of consumer processes to which this consumer belongs. By setting the same group id multiple processes indicate that they are all part of the same consumer group.

Connect Timeout	Amount of time before the connection will throw an exception.
Default Topic	All Kafka streams need to be targeted to a specific topic. In Syniti Replicate you can indicate a default topic for all replications in that connection. Alternatively, if the Default Topic is left blank, every replication will need to specify a topic name.
Serialization	Possible types of message serialization. Currently supported: 0 – JSON 1 – CSV 2 – XML 3 – AVRO
Schema Registry URL	Used only when AVRO serialization is selected. Because Avro does not include any schema information (column names, types, sizes and format) within the message that is being sent, Kafka producers and consumers use an implicit “contract” to determine how the message is serialized. Schema Registry helps ensure that this contract is met with compatibility checks and it provides centralized schema management and compatibility checks as schemas evolve.
Auto Offset Reset	Kafka property not used by Syniti Replicate. Determines how to proceed when there is no initial offset in ZooKeeper or if an offset is out of range: * smallest : automatically reset the offset to the smallest offset; * largest : automatically reset the offset to the largest offset; * anything else: throw exception to the consumer.
Use One Producer Per Group	True by default. If True, Syniti Replicate uses the same Producer for all streams within the group (this has a better impact on performance).

5. If using Kafka with Kerberos security, there are additional connection values to set. See below.
6. Click **Next** to display the **Select tables** page.
At this point, there is no text output structure available to display. You can add the information after completing the Target Connection wizard.
7. Click **Next** to display the **Summary** page.
8. Click **Finish** to complete the wizard.

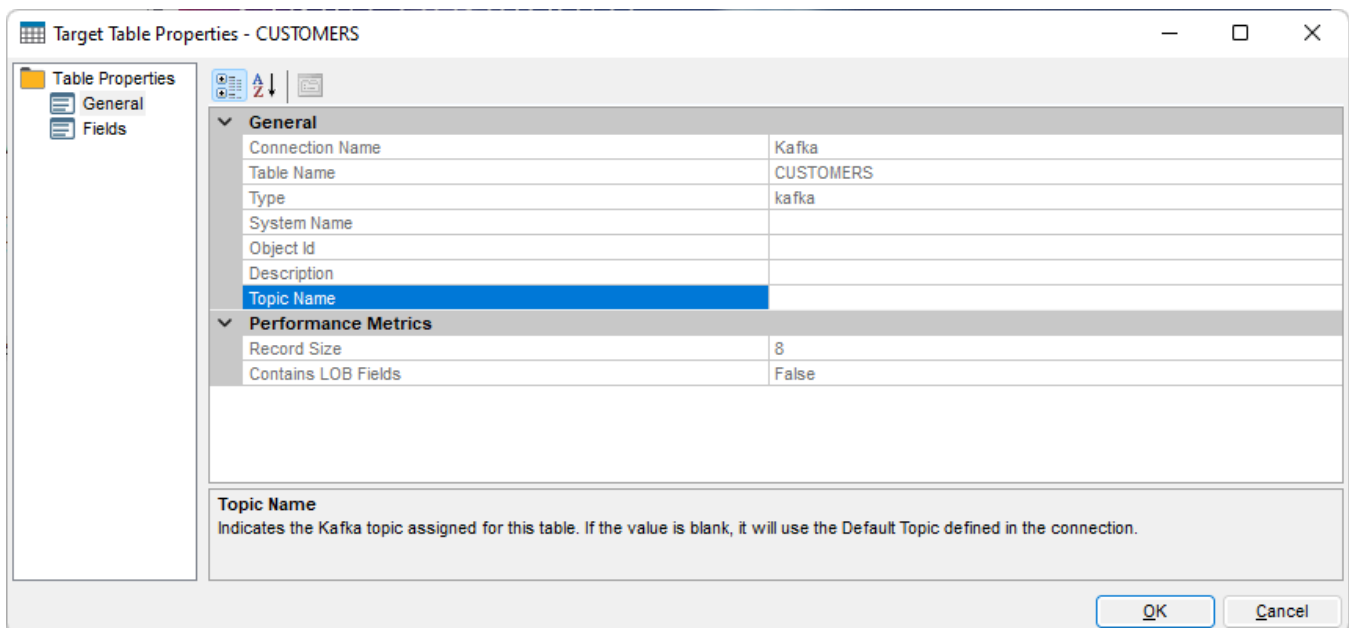
3. Add Table Information to the Target Connection

The target connection is displayed in the Metadata Explorer, but you still need to add the file representation for source table data so that when you create replications below, you can specify a source table and target “file.”

Syniti Replicate

1. In the Metadata Explorer, expand the source connection you created above.
2. Select and drag a source table to the target file connection.
3. The Create Target Table wizard is displayed.
Although you are not actually creating tables, you can use this wizard to create a representation of the data.
4. In the **Source Connection** screen, you should see the source table you selected above.
5. Click **Next** to display the **Target Connection** screen.
6. Verify that the screen displays the correct target connection name, and table name.
7. Click **Next** to review the table structure.
At this point, you can modify data types, null values and so on, if you want to modify the data eventually sent to a file.
8. Click **Next** to display the **SQL Script** screen.
The contents of this screen are inactive because there is no editable SQL script to create a table. Instead, Syniti Replicate outputs the table information to a file.
9. Click **Next** to display the **Summary** screen.
10. Click **Finish** to create the table representation in the Metadata Explorer.
11. Click **Yes** to add the table name to the target connection entry in the Metadata Explorer
12. Repeat steps 2 through 11 for each source table that you want to replicate to a file.

At the end of this process, you should have a list of table representations under the target connection in the Metadata Explorer. The only difference between these and relational database tables in the display is that, if you select the table and edit the Properties, there is a new property called **Topic Name**:



Target Table Properties - CUSTOMERS

Table Properties

- General
- Fields

General	
Connection Name	Kafka
Table Name	CUSTOMERS
Type	kafka
System Name	
Object Id	
Description	
Topic Name	

Performance Metrics	
Record Size	8
Contains LOB Fields	False

Topic Name
Indicates the Kafka topic assigned for this table. If the value is blank, it will use the Default Topic defined in the connection.

OK Cancel

Syniti Replicate

Topic Name can be edited. If blank, Syniti Replicate uses the **Default Topic** property of the connection to determine to which topic to deliver the messages. Otherwise, it uses the topic indicated in this field.

Note that if you replicate a table that has a blank topic both here and in the connection, the Data Replicator generates an error and disables the replication.

4. Define Replications

1. Expand the Metadata Explorer tree to display the table that contains the data you want to replicate.
2. Select the table.
3. From the right mouse button menu, choose Replication then Create New Replicaton....
4. In the Define Replication Type screen, type a name to identify the replication.
5. Optionally provide a description of the replication.
6. In the Replication Mode area, choose Refresh or Continuous Mirroring.

Create Replication Wizard

Syniti Replicate

Mirroring mode will define a one-way transactional replication, from the source server to the target server.

Replication Name

Replication Name: CUSTOMER

Description:

Use Group: <undefined> Create

Replication Type

Refresh

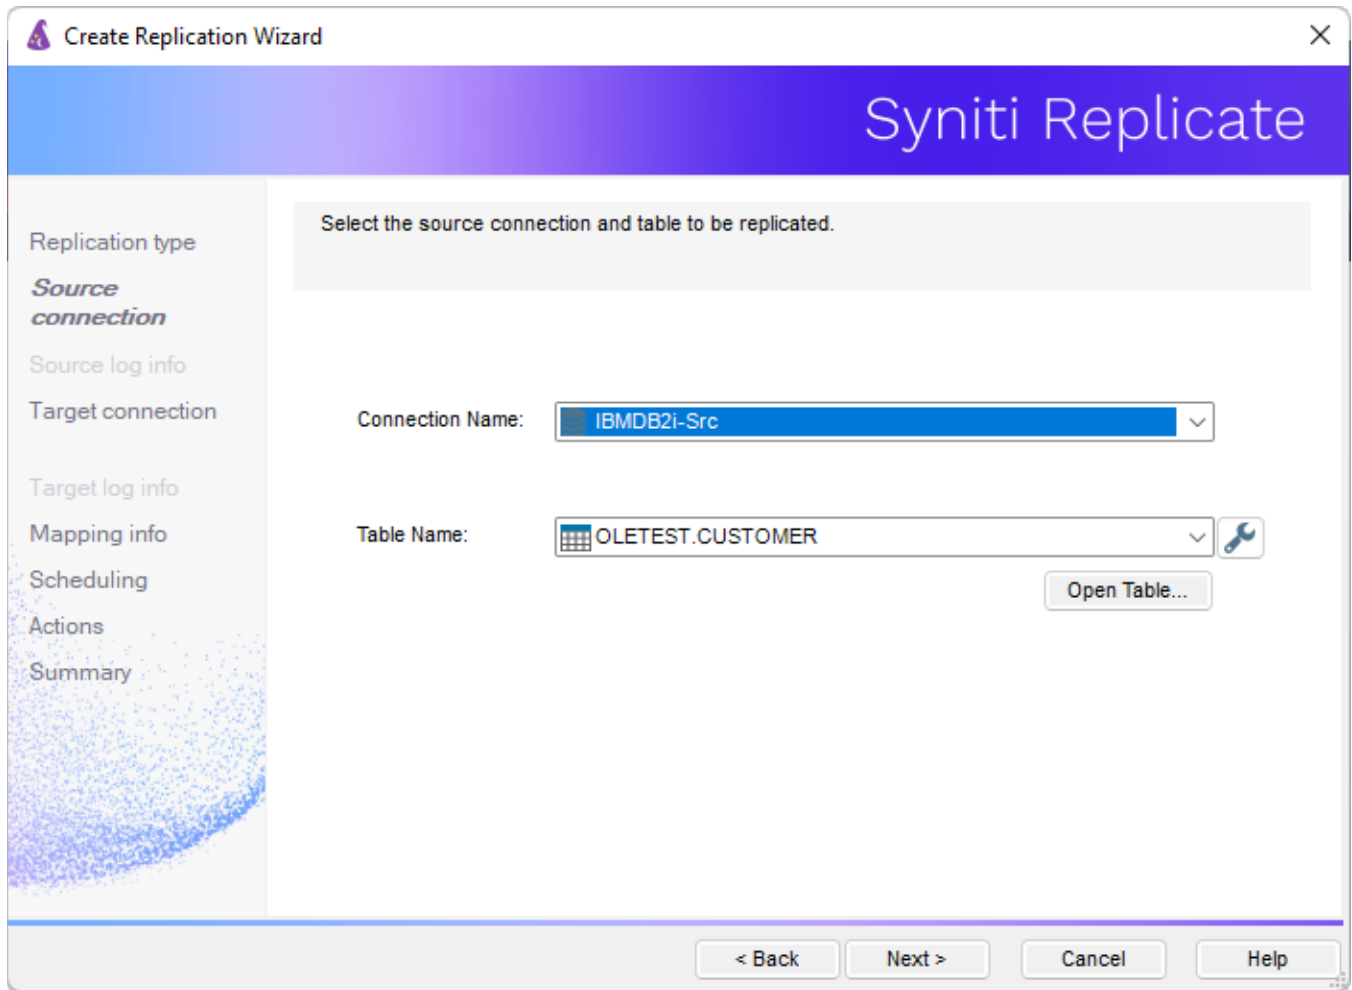
Continuous Mirroring

Synchronization

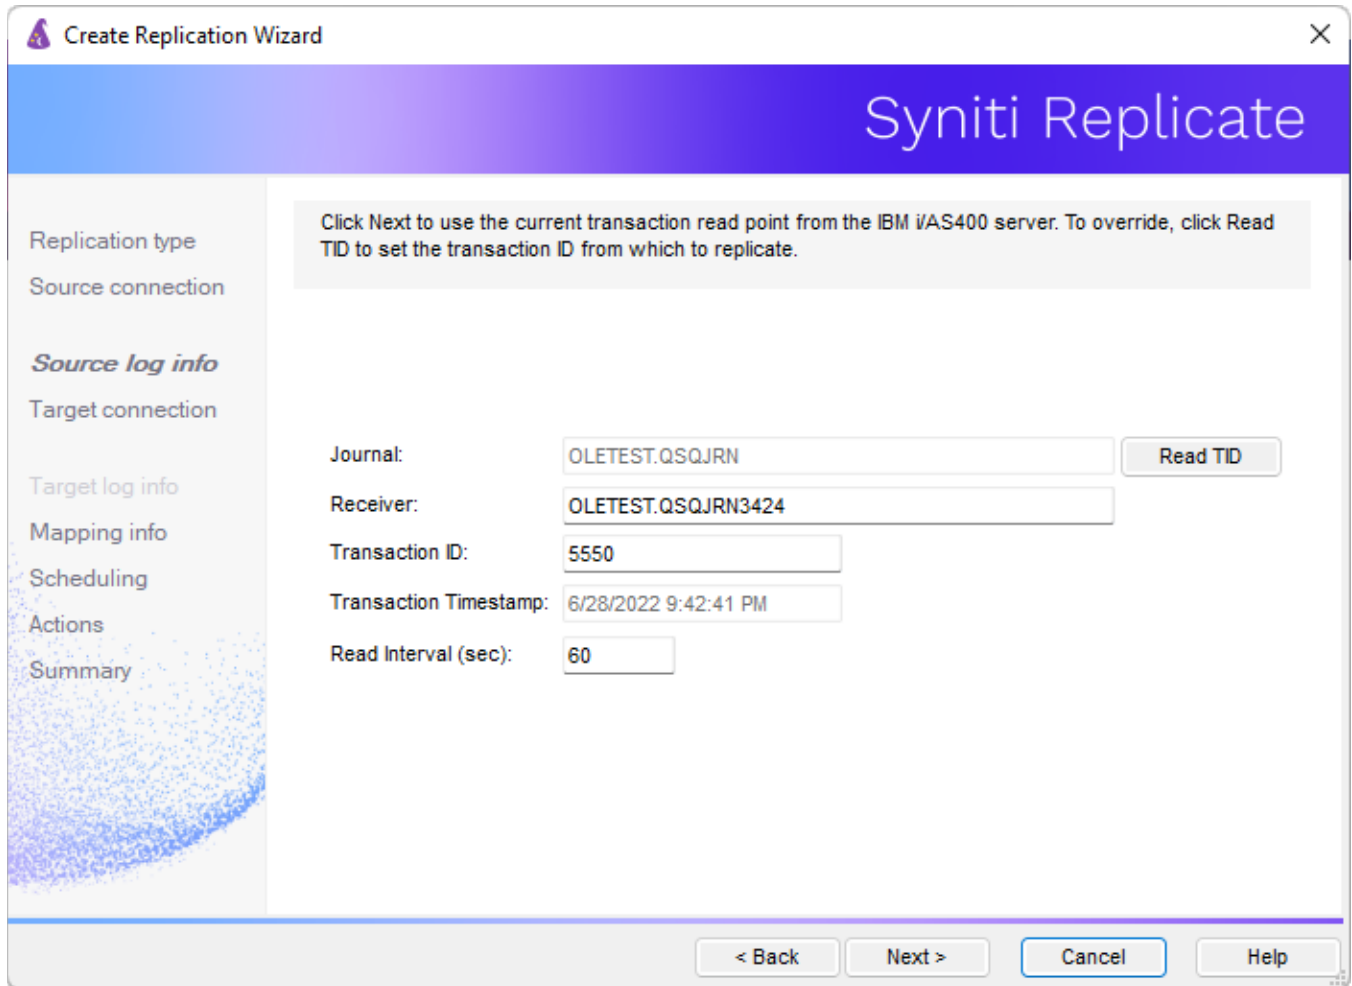
< Back Next > Cancel Help

Syniti Replicate

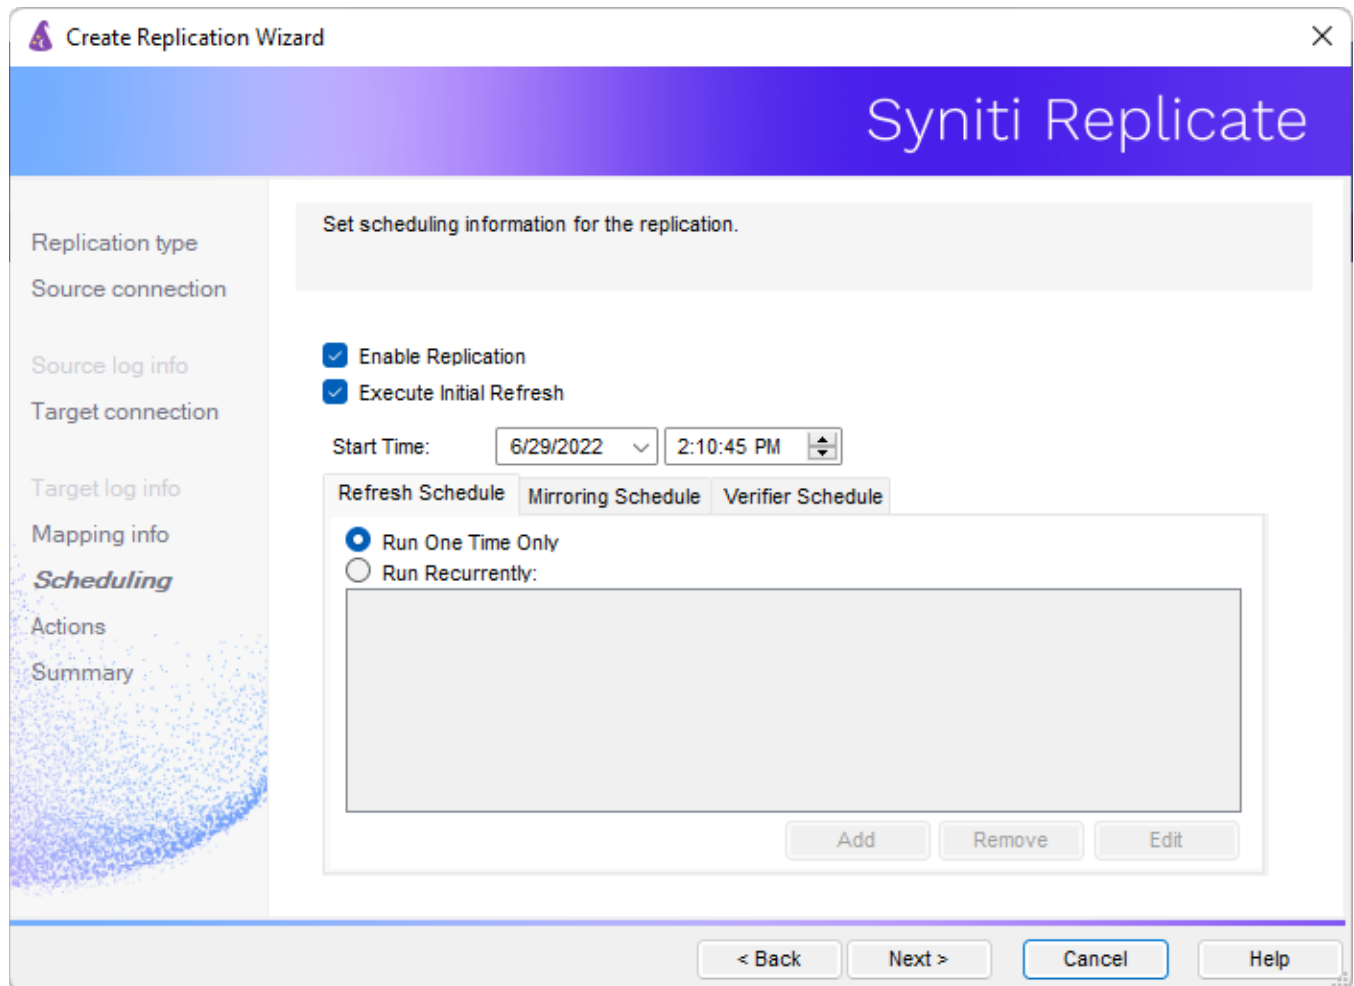
- Click **Next** to go to the **Select Source Connection** screen.



- Choose the source connection name from the drop-down list that includes all the source connections you have created in Syniti Replicate.
- Choose the table that you want to replicate from the drop-down list.
- If you want more information about the table before proceeding, click **Open Table...**
- Click **Next** to go to the **Source Log Info** screen.
Complete the fields in this screen only if you are setting up a mirroring replication. The fields displayed depend on the source database log type.




12. Click **Next** to go to the **Select Target Connection** screen.
13. Choose the target connection for text output from the drop-down list that includes all the target connections you have created in Syniti Replicate.
14. Choose the data set you want to replicate from the drop-down list.
If the drop-down list is empty, exit the wizard and add or create a target data set.
15. Click **Next** to go to the **Set Mapping Info** screen.
Source columns and target data with the same name are automatically mapped.
16. Click **Next** to go to the **Scheduling** screen.



17. Make sure the **Enable Replication** option is checked. This is required for the replication to run.
18. Set a start time for the replication. The **Start Time** field indicates the time at which the Data Replicator will begin considering the replication for execution.
19. Check the option to **Execute Initial Refresh**.
A full replication will be performed from the source table to the data file.
20. Click **Next** to go to the **Summary** screen.
21. Click **Finish** to complete the wizard.



Start Replications

If you installed the Replication Agent as a service during Syniti Replicate setup, you just need to start the service using the ServiceMonitor program  in the Windows Notification Area.


- The replication that you have scheduled should start at the specified time.
- Use the Replication Monitor tab in the Management Center to track the progress of the replication.

Syniti Replicate

If you would like to set up the Replication Agent as a service:

- From the ServiceMonitor program  in the Windows Notification Area, choose **Launch Service Installer**.
- Manage the service from Service Monitor program (located in the Windows Notification Area ).
- Use the Replication Monitor tab in the Management Center to track the progress of the replication.

To run the Replication Agent interactively:

- In the Windows Notification Area, select the Service Monitor icon .
- From the right mouse button menu, choose **Replication Agent**, then **Start** then **Application**. The replication that you have scheduled should start at the specified time.
- Use the Replication Monitor tab in the Management Center to track the progress of the replication.

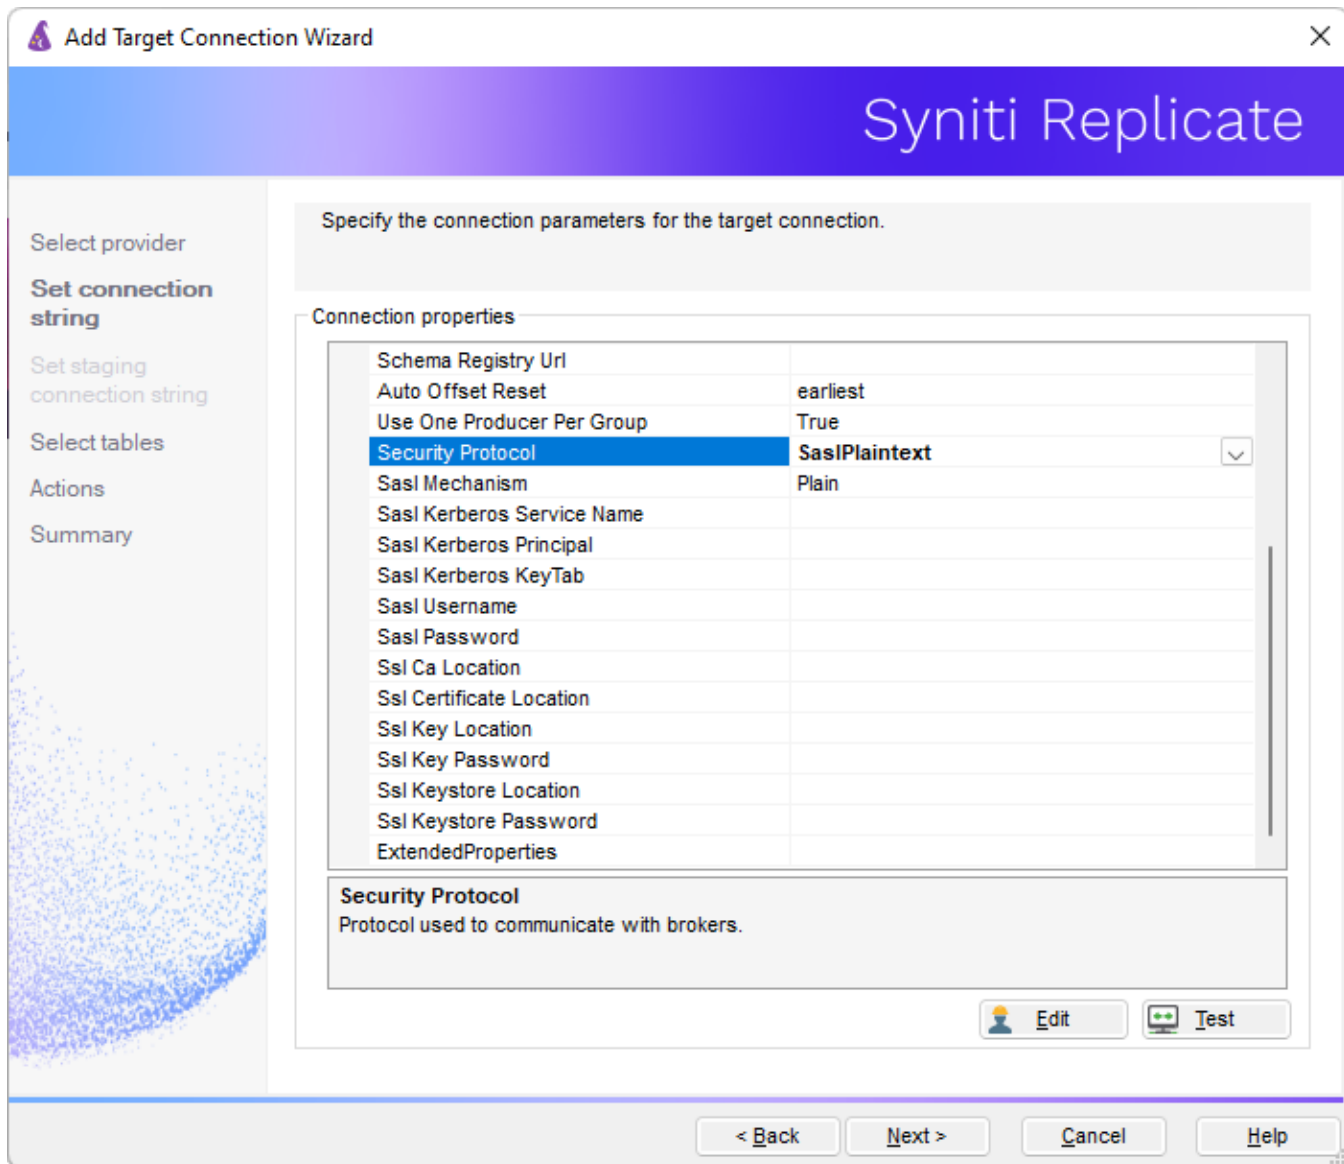
Stop Replications

Stop the Replication Agent from the Service Monitor in the Windows Notification Area.

Using Kafka with Kerberos

When using the Kerberos network authentication protocol to secure Kafka streams, you need to add certain properties to the connection string used by Syniti Replicate, either when creating the connection to Kafka using the Create Target Connection wizard, or by later editing the connection properties. The steps below assume that you are editing the connection properties for an existing target connection to Kafka.

1. In the Metadata Explorer, select the Kafka target connection.
2. From the right mouse button menu, choose Connection Properties
3. In the Connection Properties dialog, scroll to the Optional section.
4. Set the following properties as needed for your environment. Note that these properties are Kerberos properties, not Syniti Replicate properties. Any settings are passed directly to Kerberos as part of the connection string.



Security Protocol	<p>Options are:</p> <ul style="list-style-type: none"> • Plaintext - the listener does not use authentication and data is not encrypted • Ssl - an SSL channel is used • SaslPlaintext - uses SASL authentication, non-encrypted channel • SaslSsl - uses SASL authentication, SSL encrypted channel <p>For Kerberos, choose either SaslPlaintext or SaslSsl. If SaslSsl is used, properties starting with Ssl also need to be set for the SSL encryption.</p>
Sasl Mechanism	Select GssAPI

Syniti Replicate

Sasl Kerberos Service Name	Kerberos principal name that Kafka runs as, not including /hostname@REALM
Sasl Kerberos Principal	Sasl Kerberos principal name (not supported in MS Windows)
Sasl Kerberos Key Tab	Path to the Sasl Kerberos key tab file (not supported in MS Windows)