



Data Replication

Setup Notes for Replications to Amazon Redshift

Version 9.7.0.22

Software Release Date: 4/15/2020

Document History

Version	Author	Date	Reviewer	Date	Comments
1	JHLorenzin	9/27/19	ALaspertini	2/8/19	
2	JHLorenzin	4/14/20	ALaspertini		

Syniti Data Replication

Table of Contents

Connection Type	1
Additional Requirements for Amazon Redshift	1
Add Target Connection Wizard	2
Select Provider Screen	2
Set Connection String Screen	2
Connection Properties Dialog.....	4

Syniti Data Replication

Syniti Data Replication allows you to replicate data from relational database tables to **Amazon Redshift** in the following ways:

- **Refresh (Snapshot replication)**

A one-time complete replication from any major relational database source to **Amazon Redshift** as a target, according to replication settings and scripts. 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.

Continuous refresh

A regularly scheduled refresh replication as described above. The schedule is defined in the replication settings.

- **One-way mirroring (Incremental Replication)**

A continuous update of replicated tables based on changes to the source database that have been recorded in the database server log. Typically, this involves an initial refresh operation, as described above, to set up the target table. Then you can define the replication settings to check the transaction log on the source database at regular intervals. Any changes found in the log are applied to data on the **Amazon Redshift** platform.

For complete details on the setup process, check the *Syniti Data Replication User Guide* available from the Management Center **Help** menu or the *Syniti Data Replication Setup Guide*, available for download in the [Help Center](#).

Connection Type

Amazon Redshift ODBC Driver.

Download as follows:

- 32-bit Amazon Redshift (x86): <https://s3.amazonaws.com/redshift-downloads/drivers/odbc/1.4.4.1001/AmazonRedshiftODBC32-1.4.4.1001.msi>
- 64-bit Amazon Redshift (x64): <https://s3.amazonaws.com/redshift-downloads/drivers/odbc/1.4.4.1001/AmazonRedshiftODBC64-1.4.4.1001.msi>

Additional Requirements for Amazon Redshift

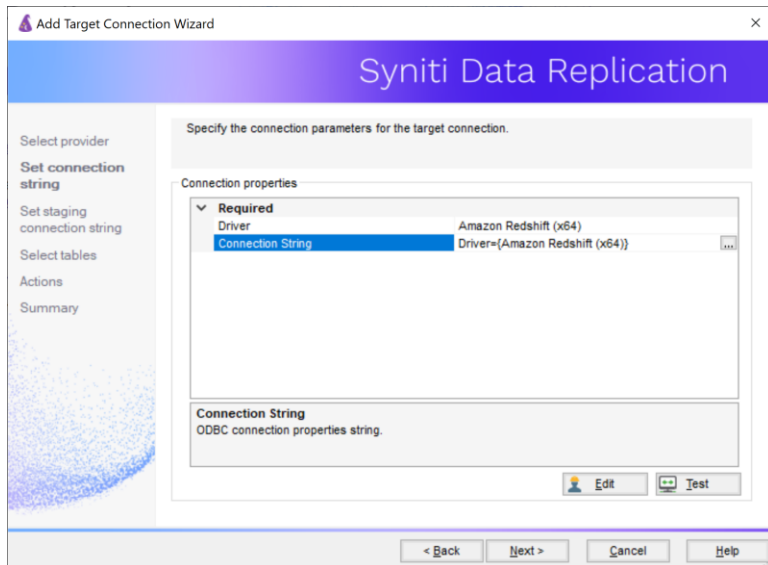
- Install the AWS Toolkit for .NET from: <https://aws.amazon.com/sdk-for-net/>
- Syniti DR will later need the pathnames to two of the DLLs in the SDK. Make sure you have the path available when you configure the target connection to Amazon Redshift.
- Set up your AWS account with an Access Key and store both the Access Key and Secret Access Key where you will be able to find them when configuring your target connection to Amazon Redshift below. The Access Key is needed to access an S3 bucket for bulk insert of data using the COPY command.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. 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.

Add Target Connection Wizard

The following field(s) require specific information for Amazon Redshift.

Select Provider Screen



Database

Choose Amazon Redshift from the drop-down list.

Provider

The default value is Amazon Redshift ODBC Driver.

Assembly

For .NET Data Provider connections only. No value is needed

Set Connection String Screen

Click in the **Connection String** value field to open the **Amazon Redshift ODBC Connection Dialog**. Values entered in this dialog should reflect those needed to make a successful connection to an Amazon Redshift cluster via an ODBC connection string, with or without Syniti DR. There are no fields specific to Syniti DR.

Syniti Data Replication

Server

Specify the Amazon Redshift cluster endpoint, as defined in the Amazon Redshift console.

Port

Type the port number for the database. The default is 5439.

Database

Type the name of the Amazon Redshift database.

Auth Type

Syniti DR has been tested using the **Standard** authentication type, with an AWS user and password. However, Syniti DR uses the values entered here as part of an ODBC connection string passed to the server. Therefore, any set of values that can connect to the server are supported. Check the [AWS documentation](#) for all the connection options.

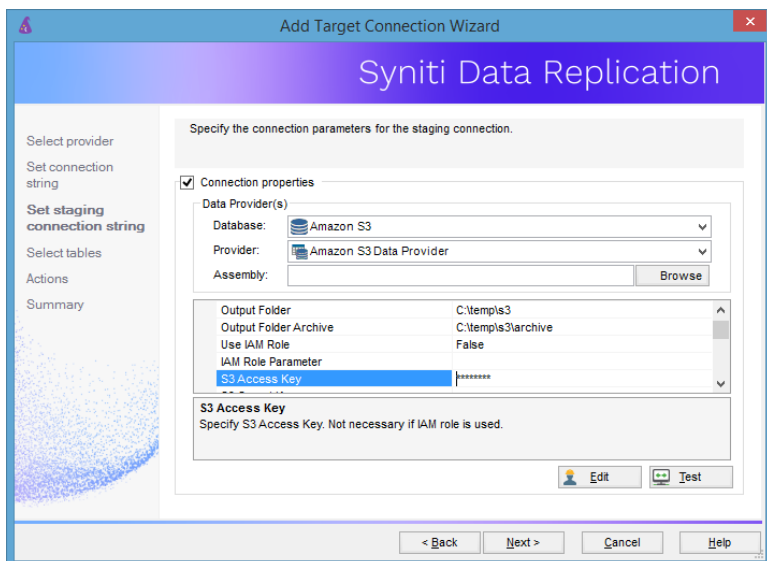
Complete the fields in the Target Connection wizard. Use the *Syniti Data Replication User Guide* available from the Management Center **Help** menu or the *Syniti Data Replication Setup Guide*, available for download in the [Help Center](#) as needed.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. 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 Data Replication

Set Staging Connection String

Redshift setup requires a staging connection in S3.



Add values for the properties as follows:

Output Folder	Local path to temporarily store replication output files.
Output folder Archive	(Optional) Set a path to archive files that are generated and uploaded.
Use IAM Role	Set to False by default. Set to True to use an IAM role for connections instead of the Access Key/Secret Key approach. For more information about using IAM roles, see https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html . Note that IAM roles work only for access between AWS services (such as when using S3 as a staging connection for bulk uploads to Redshift.)
IAM Role Parameter	(Optional – for use with Redshift only) The IAM role parameter should be configured as follows to access S3: <code>arn:aws:iam:<aws-account-id>:role/<role-name></code> The parameter information is required when using S3 in conjunction with Redshift in the COPY from statement to access files you imported from S3 to Redshift. For more information about using IAM roles with Redshift, see https://docs.aws.amazon.com/redshift/latest/mgmt/authorizing-redshift-service.html
S3 Access Key	The Access Key from your AWS login ID.
S3 Secret Key	The Secret Access Key from your AWS login ID.
S3 Bucket Name	The bucket must already exist. It can either be a bucket that you are already using, or one you have created for this purpose.
AWS SDK S3 Assembly Path	The pathname to AWSSDK.S3.dll installed with the AWS .NET SDK.
AWS SDK Core Assembly Path	The pathname to the AWSSDK.Core.dll installed with the AWS .NET SDK.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. 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.

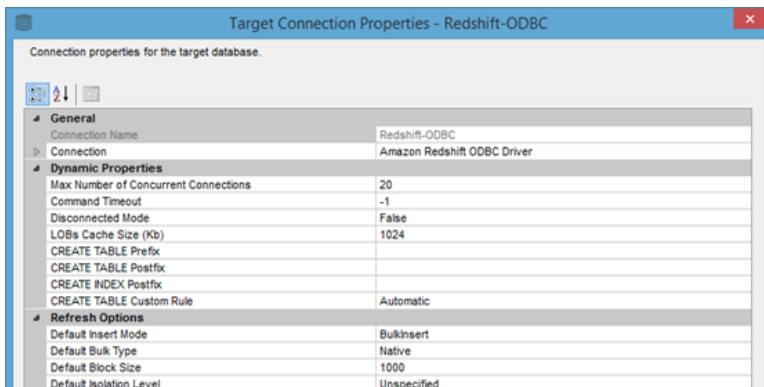
Connection Properties Dialog

The default insert mode for replicated data has been set as follows:

- Refresh replications - BulkInsert
- Mirroring replications – Single Insert

To change these settings, and to complete setup for Redshift, you need to edit the connection properties for the connection created using the Create Target Connection wizard (above).

1. In the Metadata Explorer, select the Amazon Redshift target connection you have created. From the right mouse button menu, choose **Connection Properties**.



2. Scroll to the **Refresh Options** section to change the Bulk Insert settings, as needed. Bulk Insert is the default and offers better performance.
3. Scroll to the **Mirroring Options** section to change the Bulk Insert settings, as needed. Bulk Insert is NOT set as the default but it offers better performance, so you may want to change the settings.
4. Click **OK** to save your changes.

Last Updated on 4/14/20