

# DATAZEN LAB

SIMPLE AWS 2 AZURE SQL SERVER CDC REPLICATION - ~30 minutes

This document provides an overview of how to implement a secured Change Data Capture (CDC) SQL Server database replication topology from AWS into Azure using Enzo DataZen as the technology platform, using a single DataZen instance in AWS.

## Pre-Requisites

To successfully complete this lab, you will need the following:

- An AWS account
- A Microsoft Azure account

## Overview

In this lab we will configure two SQL Server databases (one in AWS and another in Azure using the Azure SQL Database service), configure Change Data Capture (CDC) in AWS, configure the target database, and observe data flowing between the two cloud platforms.




Figure 1 - Logical Replication Topology using DataZen between AWS and Azure

*NOTE: Some of the technology choices were made to simplify the lab, including the versions of the source and target databases. Additionally, DataZen can forward changes to any target system or systems; this lab uses SQL Server as the target database for simplicity. In addition, this lab uses the Synthetic Change Data Capture logic from DataZen instead of SQL Server's built-in CDC because SQL Server CDC is not supported by SQL Server Express edition used in this lab.*

## Technologies Used

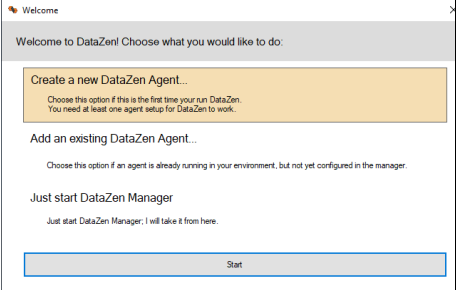
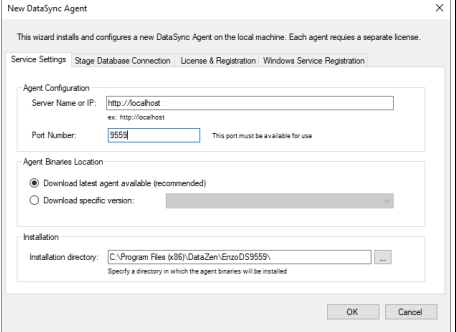
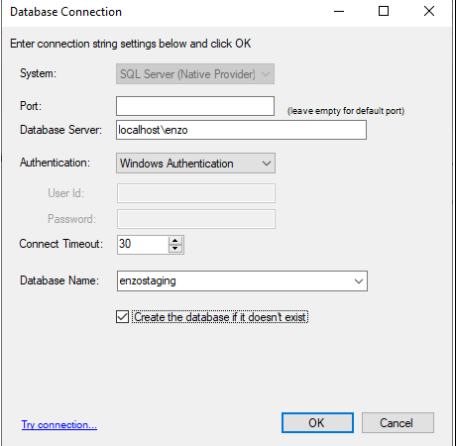
The following technologies are used in this lab:

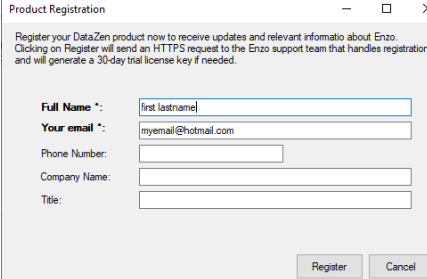
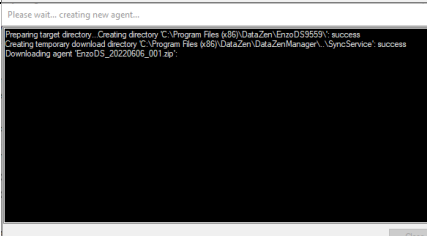
Cloud	Technology	Notes
AWS	EC2 Instance	Use the AWS Marketplace to create an Enzo Server EC2 instance that contains DataZen
	SQL Server	The EC2 image will contain a SQL Server database
	DataZen	DataZen will be installed on the same EC2 instance
Azure	Azure SQL Database	Use an Azure SQL Database instance to push changes detected in AWS

## Configure AWS

### Install DataZen on EC2 Server

Follow these steps to create a new EC2 instance from the AWS Marketplace.

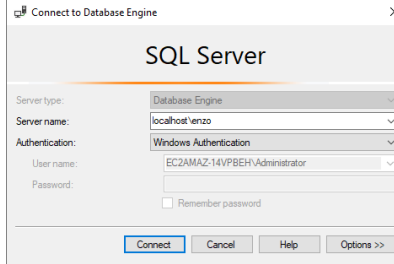
Step	Description	Details
Login to the AWS portal	Navigate to your AWS portal and login using a privilege account	
Create an EC2 image from the Enzo Server Marketplace	Navigate to <a href="https://aws.amazon.com/marketplace/pp/prodview-mdxtczpoykijk">https://aws.amazon.com/marketplace/pp/prodview-mdxtczpoykijk</a> <i>It is highly recommended to select a Medium instance size or higher.</i> <i>Make sure you enable a security rule that allows RDP connections from your computer.</i>	
Follow instructions to connect to the Instance	Connect to the EC2 Instance that was just created using an RDP connection. Once logged on, the VM may need a few minutes to download a few packages automatically.  <i>Once you see Enzo icons on the desktop, you are ready to proceed to the next step.</i>	
Starts DataZen and choose Create New Agent	To start DataZen, right-click on the DataZen icon and choose <b>Run as Administrator</b> .	
Verify the information on the General Settings tab	You can leave all the defaults provided on the General Tab. By default, a new DataZen Agent will be started on the local machine and will listen on port 9559; this is the port DataZen manager uses to control the agent.	
Click on Stage Database Connection	Enter the following information, then click OK: Database Server: <b>localhost\enzo</b> Authentication: <b>Windows Auth</b> Database Name: <b>enzostaging</b> Check the <b>Create Database if it doesn't exist</b> option	

<p><b>Click on the License tab and enter your license key</b></p>	<p>Click on the <b>Generate Free 30-day License</b> button.</p> <p>Enter your contact information in the <b>Product Registration</b> window and click <b>Register</b>.</p>	
<p><b>WAIT A FEW MINUTES</b></p>	<p>The agent installation and configuration will start automatically as a separate window.</p>	

At this point, DataZen has been installed on the AWS EC2 Server and ready to be configured for replication.

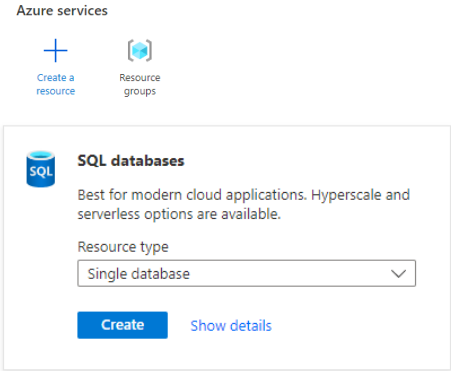
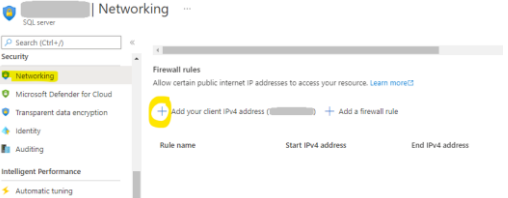
### Create a Test Database

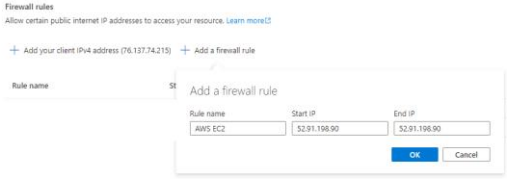
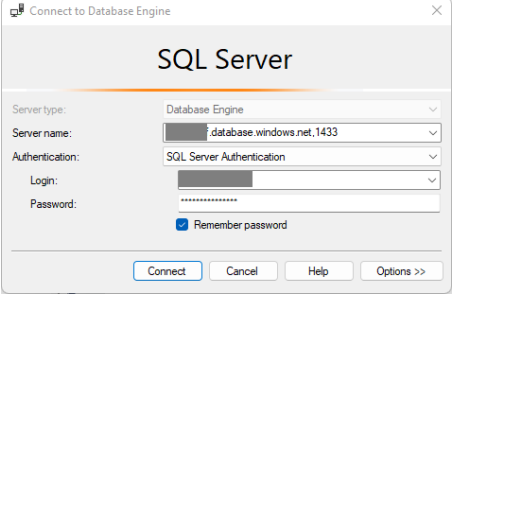

Let's create a test database that will be used for replication.

Step	Description	Details
<p><b>Start SSMS</b></p>	<p>On the EC2 Instance, double-click on the SQL Server Management Studio (SSMS) shortcut and connect to the SQL Server.</p> <p>Server Name: <b>localhost\enzo</b> Authentication: <b>Windows Authentication</b></p> <p><i>It may take a few minutes for SSMS to start depending on the Instance size</i></p> <p>Once logged on, click on New Query.</p>	
<p><b>Create a Test database and table</b></p>	<p>Run the following SQL command to create the test database and table:</p> <pre>CREATE DATABASE test GO  CREATE TABLE test.dbo.source (     id int identity(1,1) PRIMARY KEY,     stateCode nvarchar(2) NOT NULL,     stateDesc nvarchar(100) NOT NULL ) GO</pre>	
<p><b>Insert a few records</b></p>	<p>Let's add three initial records to see the [source] table with initial data:</p> <pre>INSERT INTO test.dbo.source VALUES ('FL', 'Florida'), ('CA', 'California'), ('NY', 'New York')</pre>	

### Create an Azure SQL Database

Next, let's create a new Azure SQL Database in Microsoft Azure and configure its firewall settings.

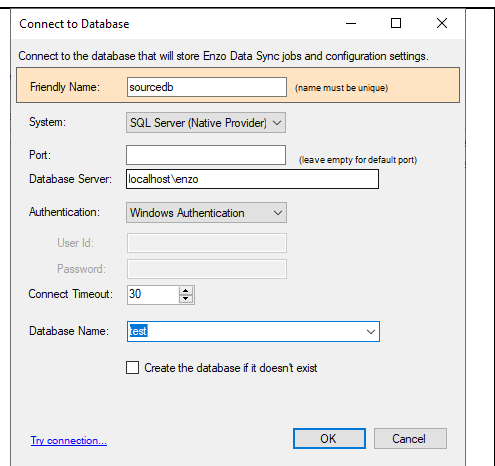
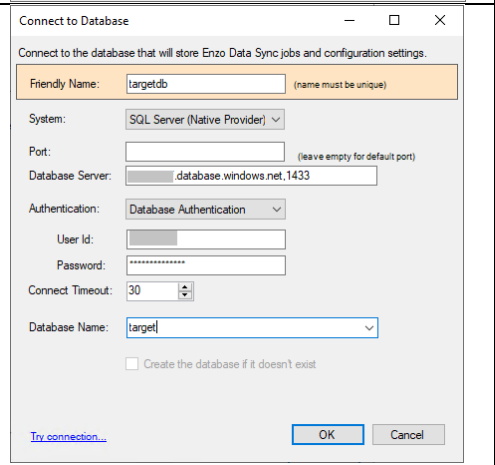
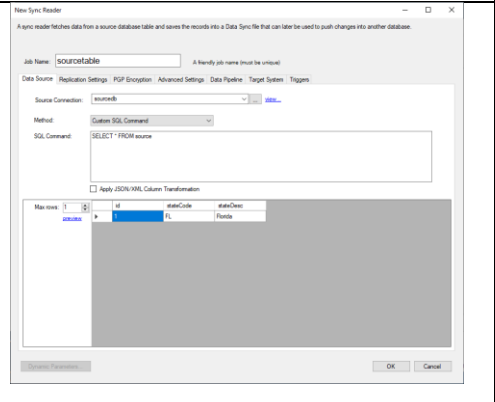
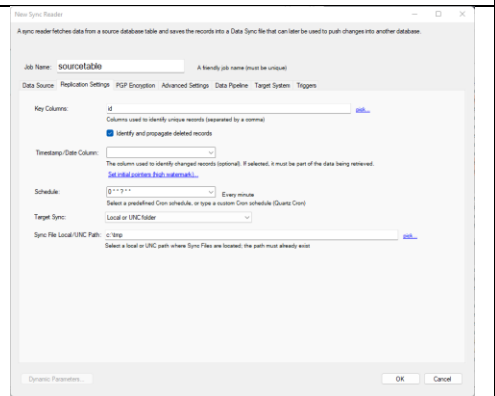
Step	Description	Details
<b>Login to Azure</b>	Open a browser and navigate to the Azure portal, then login.  <a href="https://portal.azure.com">https://portal.azure.com</a>	
<b>Create an Azure SQL Database</b>	<p>From the Azure Portal, click on Create a Resource. For complete instructions on how to create a Single Azure SQL Database, follow Microsoft's documentation here:</p> <p><a href="https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart?view=azuresql&amp;tabs=azure-portal">https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart?view=azuresql&amp;tabs=azure-portal</a></p> <p>Search for Azure SQL, and click on it, then choose <b>Create</b>. Then, choose <b>SQL Database / Single Database</b> and click <b>Create</b>.</p> <p><b>Enter the following information:</b></p> <p><b>Resource Group:</b> (leave default)  <b>Database Name:</b> <b>target</b>  <b>Server:</b> (create new if needed)  <b>Workload Env:</b> development</p> <p><b>On the Networking Tab:</b></p> <p>Add current IP Address: <b>Yes</b></p> <p><b>On the Security Tab:</b></p> <p>Enable Microsoft Defender: <b>Not now</b></p> <p><b>Click on Create</b></p>	
<b>Find you Database Server connection string</b>	<p>Once the database has been created, navigate to the Overview tab, and click on the server name. The server name should look something like this:</p> <p><i>abcdefghijklm.database.windows.net,1433</i></p> <p>You are now on the Database Server page.</p>	
<b>Locate Server Firewall Rules</b>	<p>Find the Fire Rule options, located under the Networking tab.</p> <p>If you already see a rule that allows connection from the EC2 server's IP Address, skip the next step.</p>	

<p><b>Add EC2 Firewall Rule</b></p>	<p>Add a firewall rule manually on the SQL Database Server itself so that your EC2 instance can connect to Azure SQL.</p>	
<p><b>Connect using SSMS</b></p>	<p>Now that your Azure SQL Database is running and configured, connect to it using SSMS from your EC2 Instance.</p> <p>Click on New Query, then right-click inside the query window, and select <b>Connection -&gt; Change Connection</b>.</p> <p><b>Server Name:</b> The Azure server name  <b>Authentication:</b> SQL Server Authentication  <b>Login:</b> Your Azure SQL Server login name  <b>Password:</b> Your Azure SQL password</p> <p><i>IF YOUR EC2 Instance PUBLIC IP ADDRESS WAS NOT ADDED TO THE FIREWALL RULES IN PRIOR STEPS YOU WILL GET AN ERROR.</i></p>	
<p><b>Select the target database</b></p>	<p>Once connected, select the <b>target</b> database from the database dropdown.</p> <p>You are now connected to the target database from the EC2 Instance.</p>	

### Configure DataZen for Replication

To finalize the source configuration, the next step is to configure DataZen to capture the changes identified in the source table, and create a Change Log file that will be stored in an Azure Blob (note other target stores are possible, including AWS Bucket, and FTP locations).

Step	Description	Details
<p><b>Configure Connections in DataZen</b></p>	<p>Next, we need to configure DataZen's database connections. On the AWS EC2 instance, using DataZen Manager, click on the DataZen agent on the left and choose <b>Configuration -&gt; Central Connection Strings</b>.</p>	

<p><b>Create the source SQL Server Connection</b></p>	<p>From the Central Connection Strings window, choose <b>New -&gt; Database Connection</b>.</p> <p>Enter a name for this connection (ex: <b>sourcedb</b>) and fill out the rest of the information needed to connect to SQL Server:</p> <p>Source System: <b>SQL Server</b>                  Database Server: <b>localhost\enzo</b>                  Authentication: <b>Windows Authentication</b>                  Database Name: <b>test</b></p> <p>Click OK.</p>	
<p><b>Create the target SQL Server Connection</b></p>	<p>Choose <b>New -&gt; Database Connection</b> again.</p> <p>Enter a name for this connection (ex: <b>targetdb</b>) and fill out the rest of the information needed to connect to SQL Server:</p> <p>Source System: <b>SQL Server</b>                  Database Server: <b>the Azure SQL Server</b>                  Authentication: <b>Database Authentication</b>                  User Id: <b>You Azure SQL login id</b>                  Password: <b>Your Azure SQL Login password</b>                  Database Name: <b>target</b></p> <p>Click OK, and close the Central Connection Strings window.</p>	
<p><b>Create a new Sync Job</b></p>	<p>Click on <b>New -&gt; New Data Sync</b>.</p> <p>Enter the following information:</p> <p>Job Name: <b>sourcetable</b>                  Source Connection: <b>sourcedb</b>                  Method: <b>Custom SQL Command</b>                  SQL Command: <b>SELECT * FROM test.dbo.source</b></p> <p>Click Preview...</p>	
<p><b>Select the Replication Settings Tab</b></p>	<p>Enter the following information:</p> <p>Key Columns: <b>id</b>                  Schedule: <b>0 * * ? * *</b>                  Identify &amp; Propagate Deleted Records: <b>Yes</b>                  Target Sync: <b>Local or UNC Folder</b>                  Sync File: <b>c:\tmp</b></p> <p><i>NOTE: In production environments, when large volumes of data changes are expected, it is best to use Change Data Capture (CDC) or Change Tracking for improved performance.</i></p>	


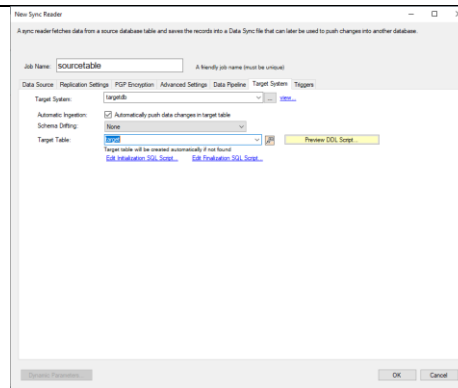
**Select the Target System tab**

Enter the following information:

Target System: **targetdb**  
 Schema Drifting: **Auto add and update columns**  
 Target Table: **target**

Click OK. The job will start immediately and will run every minute to automatically detect new/updated records.


Note: A warning may come up indicating the target table doesn't exist; click Yes. The table will be created automatically.



**Verify Sync Job completed**

Once the job is completed, within a few seconds, you can inspect the output window of the job itself. Click on the job, and expand the output window at the bottom.

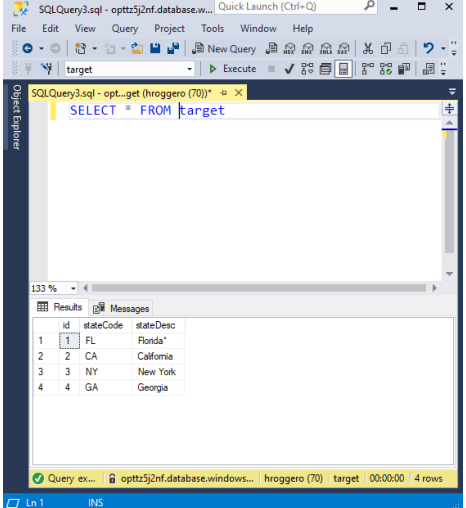
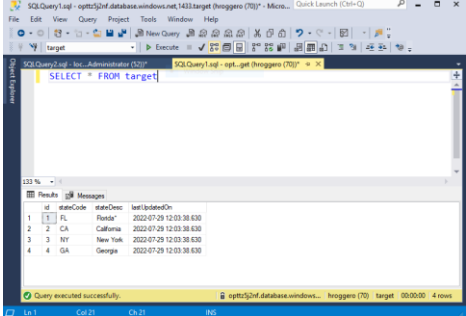
At this point, the target table contains a copy of the source table, along with its records.



## Test Cross-Cloud Replication

Now that the replication topology is in place, let's test ongoing database replication between AWS and Azure.

Step	Description	Details
Add/Modify source data	<p>Bring up SSMS, and make sure to select the Query window that is connected to the local SQL Server database.</p> <p>To add records to the <b>source</b> table, run this SQL batch:</p> <pre>INSERT INTO test.dbo.source VALUES ('GA', 'Georgia')  UPDATE test.dbo.source SET stateDesc = 'Florida*' WHERE stateCode = 'FL'</pre> <p>Within a minute the changes will be detected and replicated directly to the Azure SQL Database table.</p>	

<p><b>Verify the data in the target table</b></p>	<p>Switch the query window to the target database, and run this SQL command to visualize the records (note, it may take up to 1 minute before the changes take place).</p> <pre>SELECT * FROM target</pre>	
<p><b>Propagate Schema Changes</b></p>	<p>On the source database, run the following command; a new datetime field will be added:</p> <pre>ALTER TABLE test.dbo.source ADD lastUpdatedOn DATETIME NOT NULL DEFAULT(GETUTCDATE())</pre>	
<p><b>Inspect target table</b></p>	<p>A minute later, the new column will be added to the target table. Run this command on the target Azure SQL database:</p> <pre>SELECT * FROM target</pre> <p>NOTE: We explicitly selected the Schema Drifting option to 'Auto add and update columns'. This option automatically propagates new columns and changes the target data types.</p>	
<p><b>Delete a record</b></p>	<p>On the source database, run the following command to delete a record:</p> <pre>DELETE FROM test.dbo.source WHERE id=1</pre>	
<p><b>Inspect target table</b></p>	<p>A minute later, run this command on the target Azure SQL database:</p> <pre>SELECT * FROM target</pre>	

## How to Restart the DataZen Agent

Under production deployments, the DataZen Agent runs as a Windows Service and restarts automatically upon reboots. However, in this lab, the agent runs as a console application. As a result, if the agent stops, or after a reboot, you will need to restart the agent manually.

To start the agent manually from the EC2 Instance, open Windows Explorer, navigate to **C:\Program Files (x86)\DataZen\EnzoDS9559**, and start **EnzoDS.exe** in administrative mode.

## Further Considerations

This lab was designed to introduce you to the capabilities of DataZen replication and its automated Synthetic Change Data Capture capabilities. Several choices were made to simplify



this lab so that the core capabilities could be easily demonstrated, including the ability to establish a fully decoupled replication topology between cloud providers and database systems.

Additionally, the following considerations apply for more complex scenarios and production environments:

- **Cross Database Replication**  
The target system can be any SQL Server edition (including Azure SQL Database), any other database server (such as Oracle) or platform (including Azure Event Hub or ADLS Parquet files).
- **Native SQL Server CDC and Change Tracking Support**  
In this lab, we selected DataZen's automated Synthetic CDC capabilities; however, you can also select SQL Server CDC or Change Tracking. CDC and Change Tracking allow you to leverage SQL Server's built-in data capture to detect inserts, updates, and delete operations. This is the preferred option when dealing with tables that experience large data changes.
- **Decoupled Replication**  
We chose to install and configure a single VM (in AWS) to quickly demonstrate how to replicate data from a single server. However, in some instances where security is important, it may be necessary to install a second VM in Azure, and use a cloud drive to push the change log between AWS and Azure.
- **Multicast**  
Because DataZen uses Change Logs to capture changes, you can replay the same logs to any number of target systems, or stand up a new target system at any time in the future.
- **PGP Encryption**  
Change Logs can also be encrypted using PGP; this can be useful if you do not trust the intermediate storage system.
- **Any Source**  
We chose SQL Server Express as the source system for this lab; however, virtually any source system could be used, including other database platforms, messaging systems, files or HTTP data sources.

---

*For more information about DataZen and its capabilities, please contact [info@enzounified.com](mailto:info@enzounified.com) or visit <https://www.enzounified.com>*

---