



SQL2 Fabric Mirroring

Introduction

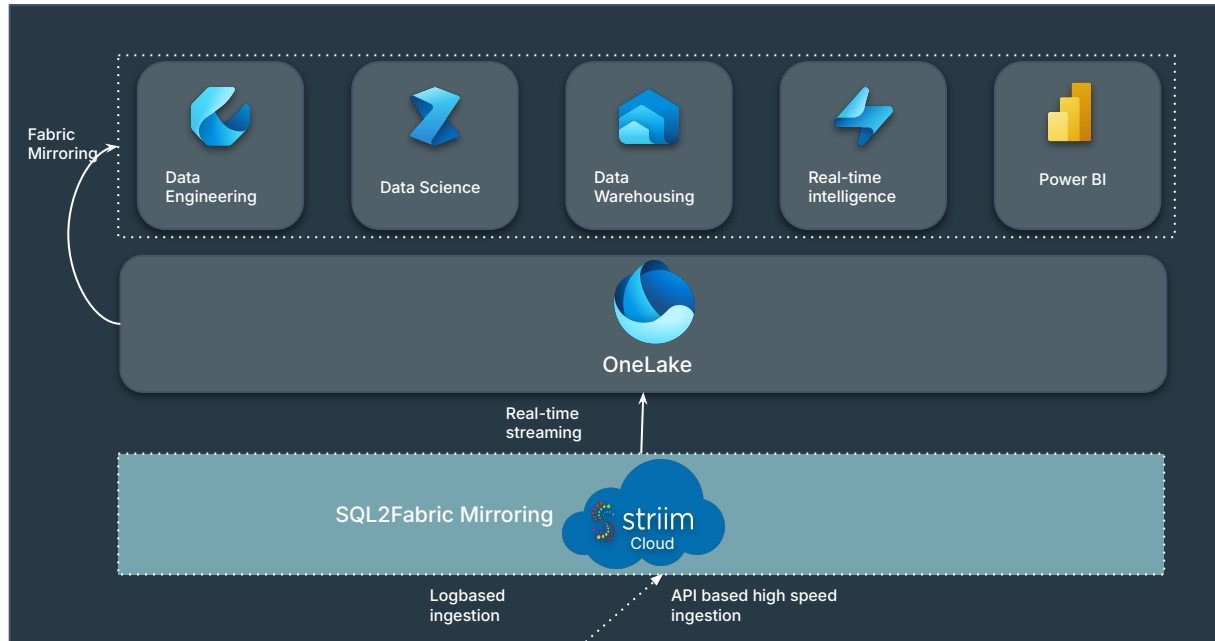
SQL2Fabric Mirroring: Seamless, Real-Time Data Integration to Microsoft Fabric

Striim's SQL2Fabric Mirroring is a fully managed, low-cost service designed to streamline the replication of on-premises SQL Server data to Microsoft Fabric. This powerful solution offers several key benefits:

- **Personalized User Experience:** Easily load your on-premises SQL Server data into Fabric with a user-friendly interface.
- **Automated Data Pipelines:** Benefit from automated initial load and Change Data Capture (CDC) processes, ensuring seamless data replication.
- **Automated Schema Conversion:** Simplify data migration with automated schema conversion.
- **High Scalability and Availability:** Handle large datasets and ensure uninterrupted data flow with cross-regional disaster recovery features.
- **Robust Security:** Protect your sensitive data with advanced security measures and compliance standards.

Architecture - Striim + Fabric Mirroring service

Striim offers high performance data ingestion from on-premise SQL server to Fabric OneLake and seamlessly integrated with Microsoft Fabric to pick up data from OneLake and write to Fabric Data warehouse or Lakehouse targets



Getting started

Simple few steps to try SQL2fabric Mirroring service by Striim

1. Subscribe to Fabric solution plan from Azure marketplace & Sign up to Striim Cloud
2. Create Striim Cloud cluster
3. Set up Private link for MSSQL & Fabric target (Mirroring, Data warehouse & Lakehouse)
4. Create fully automated smart data pipeline
5. Monitor data flowing to Fabric targets



Sign up

- Marketplace
- [Striim.com/](https://striim.com/)
- Choose appropriate plan
- Trial plan for 60 days

Create cluster

- Create Striim cluster
- After 60 days start charging by reporting to Azure

Setup secure connection

- SSH connection
- Azure private link connection

Create app

- Creates data pipeline

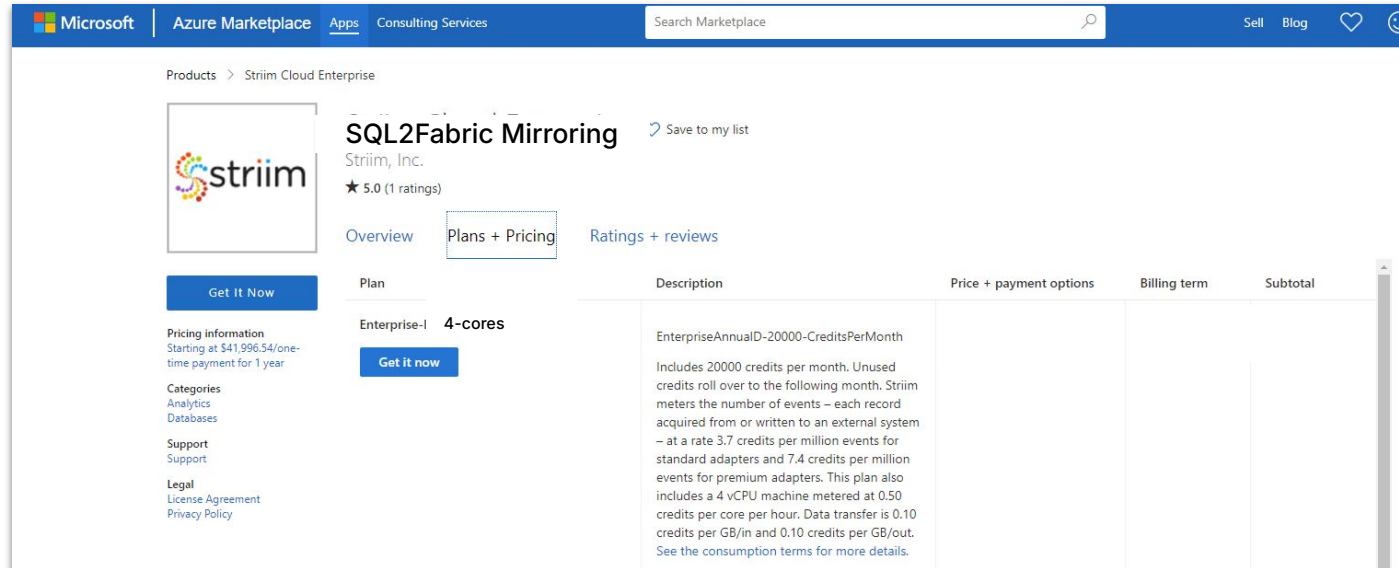
Value

- Monitor data streaming to Fabric targets

Onboarding Step 1 - Sign up

User chooses one of these options

- Sign up on Marketplace
- Sign up for trial on [Striim.com/Fabric](https://striim.com/Fabric)
- Signup from the link Microsoft shared with the customer



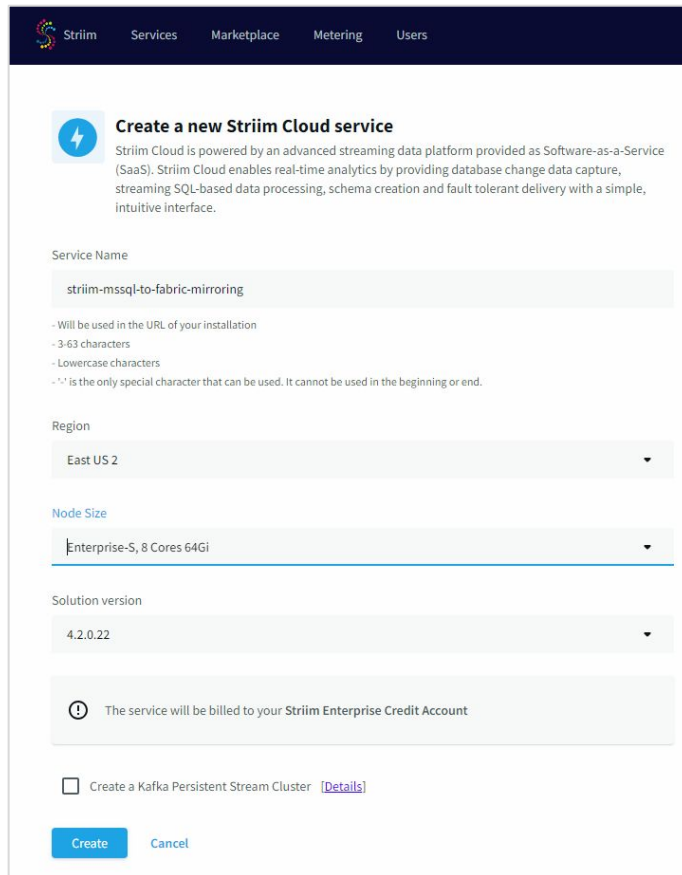
The screenshot shows the Azure Marketplace page for the product "SQL2Fabric Mirroring" by Striim, Inc. The page includes a navigation bar with "Microsoft", "Azure Marketplace", "Apps", and "Consulting Services". A search bar is present in the top right. The main content area features the Striim logo, the product name "SQL2Fabric Mirroring", and a "Save to my list" button. Below the product name are tabs for "Overview", "Plans + Pricing", and "Ratings + reviews". A "Get It Now" button is visible. The pricing information states "Starting at \$41,996.54/one-time payment for 1 year". The categories listed are "Analytics" and "Databases". The support and legal sections are also visible. The table below shows the pricing details for the "Enterprise-I 4-cores" plan.

Plan	Description	Price + payment options	Billing term	Subtotal
Enterprise-I 4-cores Get it now	EnterpriseAnnualD-20000-CreditsPerMonth Includes 20000 credits per month. Unused credits roll over to the following month. Striim meters the number of events – each record acquired from or written to an external system – at a rate 3.7 credits per million events for standard adapters and 7.4 credits per million events for premium adapters. This plan also includes a 4 vCPU machine metered at 0.50 credits per core per hour. Data transfer is 0.10 credits per GB/in and 0.10 credits per GB/out. See the consumption terms for more details.			

Onboarding Step 2 - Create cluster

Create Striim cluster

1. Create new service
2. Select the region
3. Select the size of the cluster



The screenshot shows the Striim console interface for creating a new service. The navigation bar at the top includes 'Striim', 'Services', 'Marketplace', 'Metering', and 'Users'. The main content area is titled 'Create a new Striim Cloud service' and includes a description of Striim Cloud as a SaaS platform. Below the description are several configuration fields: 'Service Name' (text input with value 'striim-mssql-to-fabric-mirroring'), 'Region' (dropdown menu with value 'East US 2'), 'Node Size' (dropdown menu with value 'Enterprise-S, 8 Cores 64Gi'), and 'Solution version' (dropdown menu with value '4.2.0.22'). A warning message states 'The service will be billed to your Striim Enterprise Credit Account'. At the bottom, there is a checkbox for 'Create a Kafka Persistent Stream Cluster' with a link to 'Details', and two buttons: 'Create' and 'Cancel'.

Create a new Striim Cloud service

Striim Cloud is powered by an advanced streaming data platform provided as Software-as-a-Service (SaaS). Striim Cloud enables real-time analytics by providing database change data capture, streaming SQL-based data processing, schema creation and fault tolerant delivery with a simple, intuitive interface.

Service Name

striim-mssql-to-fabric-mirroring

- Will be used in the URL of your installation
- 3-63 characters
- Lowercase characters
- '-' is the only special character that can be used. It cannot be used in the beginning or end.

Region

East US 2

Node Size

Enterprise-S, 8 Cores 64Gi

Solution version

4.2.0.22

⚠ The service will be billed to your Striim Enterprise Credit Account

Create a Kafka Persistent Stream Cluster [\[Details\]](#)

Create Cancel

Onboarding Step 3 - Create secure connections

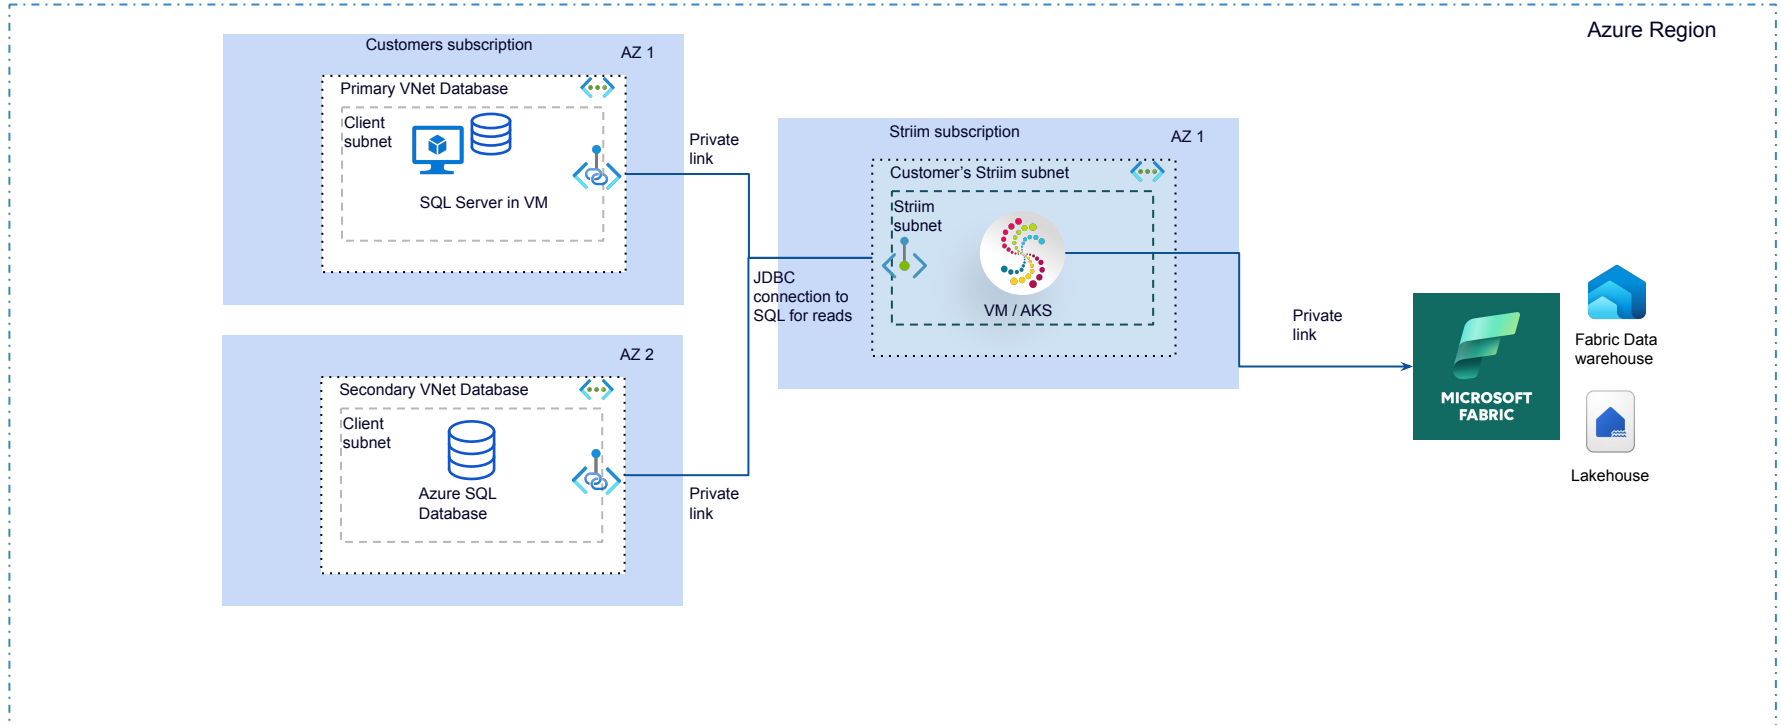
Striim supports below connection methods to source and targets

1. SSH tunnel
2. Azure Private link (recommended)

There are three possible topologies to connect the sources depending on where source is located.

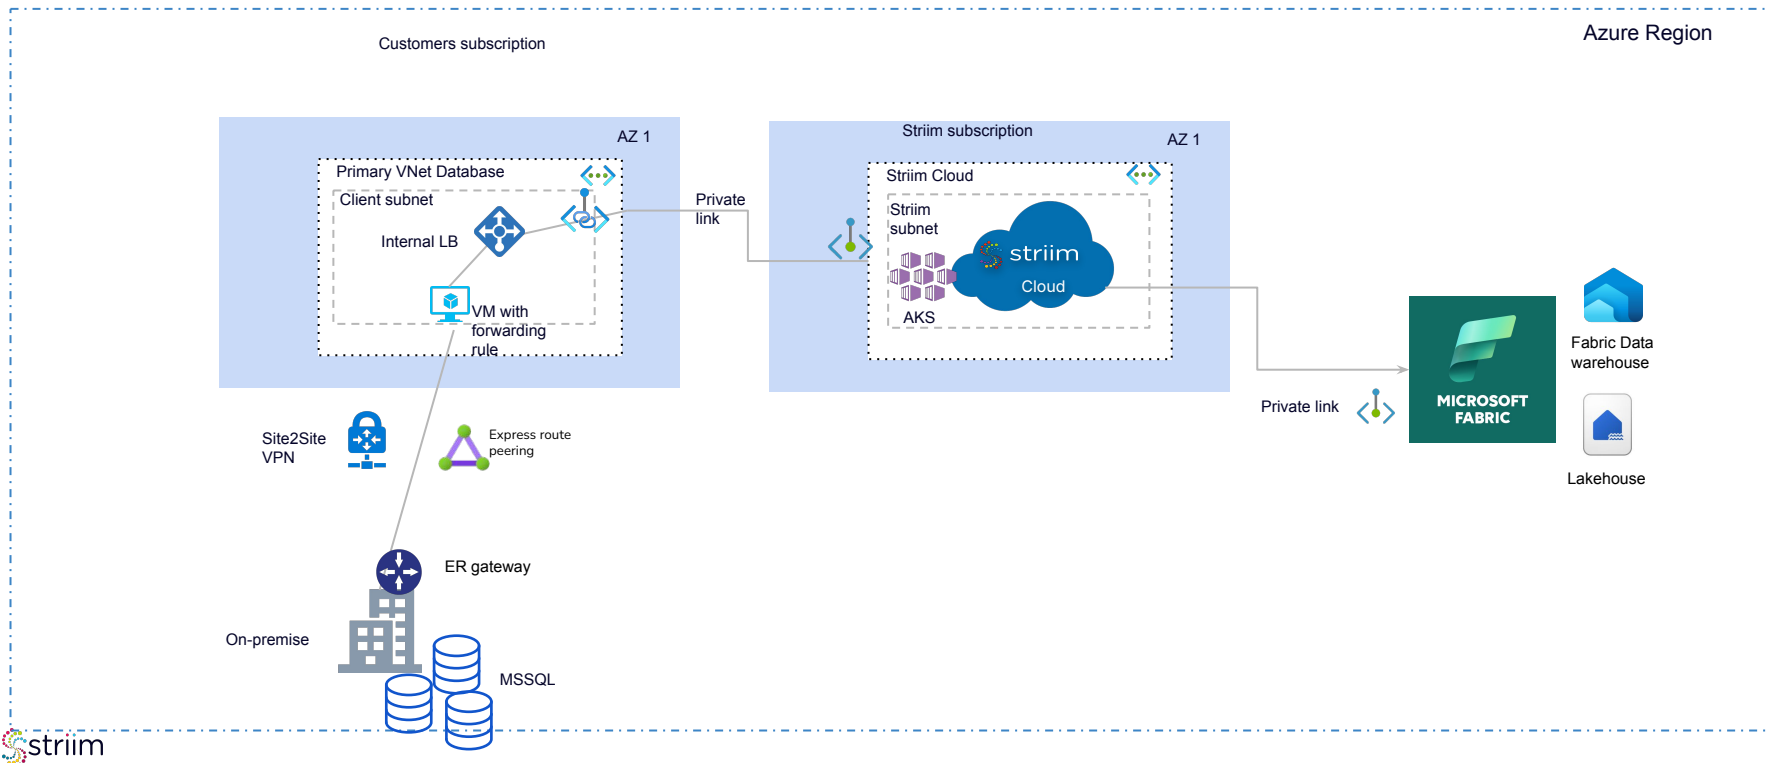
Topology 1 (source DB in Cloud)

Striim is deployed into AKS service in Azure cloud in Striim subscription with a dedicated VPC for customers



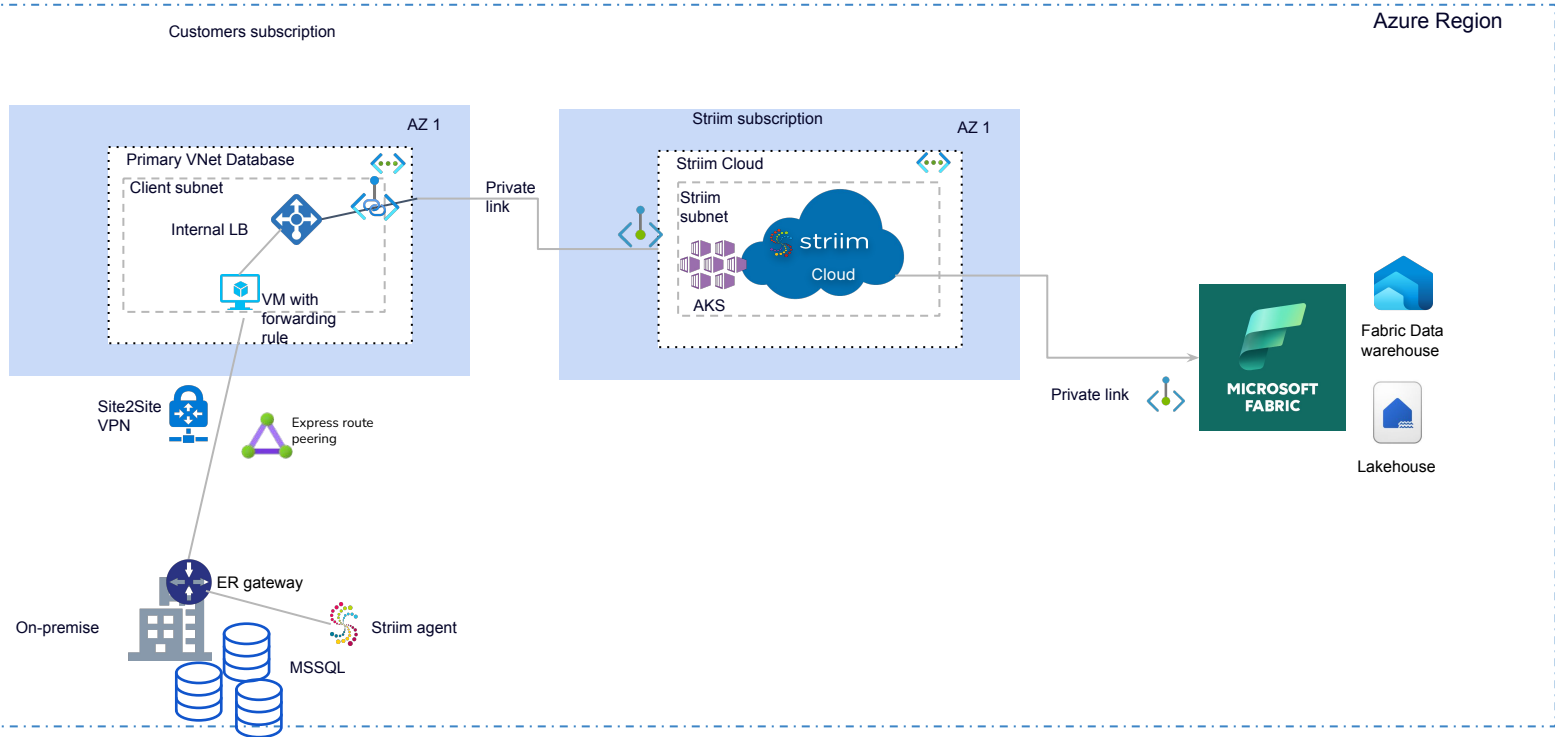
Topology 2 (source DB in on-premise)

Striim will require customer to setup VPN between on-premise & customers subscription to securely connect to on-premise via Private link



Topology 3 (source DB in on-premise with agent)

Striim will require customer to setup VPN between on-premise & customers subscription to securely connect to on-premise via Private link and install Striim MSJet agent



Creating Private link

Create a private link with an endpoint for each source & target.

To create PE, go to Striim Cloud service → Security and follow the prompts

Striim Cloud | Azure | East US

Launch Stop More

You can securely connect the service to your databases by [configuring SSH Tunnels](#).

[Learn more](#) [Dismiss](#)

You can now securely connect the service to your databases by [creating Private Endpoints](#) without exposing the data stream into public networks.

[Learn more](#) [Dismiss](#)

Overview Secure connection Monitoring Users API Persistent Streams **Private Endpoint**

Connect Securely To Striim

You can securely connect your applications to Striim Cloud by creating Private Endpoint.

Create Private Endpoint [Create Private Endpoint](#)

NAME	FQDN	SERVICE ALIAS	STATUS	ACTIONS
------	------	---------------	--------	---------

No Private Endpoints created yet

Once created, you can view and manage your Private Endpoints from here.

Create Private Endpoint

Name

pe2

Private Link Alias

privatelinksvc.ed7b438a-7b95-484c-8196-8f8dfb20e63b.eastus.azure.privatelinkservice

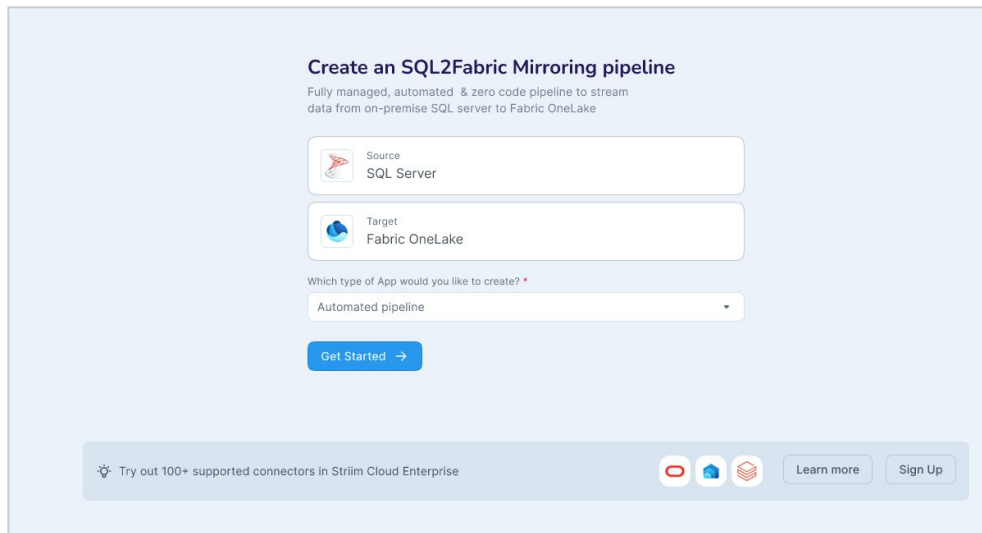
This service will be billed to your Striim Credit Account. Reach out your admin for Private Link Alias details.

[Learn more](#)

Cancel [Create Private Endpoint](#)

Onboarding Step 4 - Create data pipeline

1. Click on Launch to build apps
2. Striim by default selects the source SQL Server and target for you
3. Simply provide the required inputs with Private link details created earlier
4. Data Pipeline starts moving data



Create an SQL2Fabric Mirroring pipeline
Fully managed, automated & zero code pipeline to stream data from on-premise SQL server to Fabric OneLake

Source
SQL Server

Target
Fabric OneLake

Which type of App would you like to create? *

Automated pipeline

[Get Started →](#)

Try out 100+ supported connectors in Striim Cloud Enterprise

[Learn more](#) [Sign Up](#)

Source configuration

1. Select the SQL server and Fabric Mirror template
2. Enter SQL server credentials
3. Enter the Private Endpoint details to connect to the SQL
4. If MSSQL running on-premise, make sure private end point is able to reach source database via VPN
5. Select schema and tables

Create an App using a Template
An App is a sequence of actions that moves data from a source to a target.

SQL Server Change

Fabric Mirror Change

What type of App would you like to create?*

Automated pipeline

Get Started →

Build a Custom Flow

Start from scratch Import a TQL

[View all templates](#)

Configure SQLServer Source
Let's get started by configuring your Source. Please fill in the connection details so we can check if Striim is able to reach your SQLServer Database.

Hostname*
10.5.8.189

Port*
1433

Username*
salesuser

Private Endpoint to connect
.....

Database Name*
sales

Where is the database located?*

OnPrem

Connect using SSL
Use SSL to connect

Cancel Next

Select the Tables

dbo

Please select the list of Tables to load from dbo

Select All

pricing

stores

systranschemas

Fabric target configuration

1. Enter Fabric credentials
2. Choose user/pwd or OAuth connection types
3. Review and finish configuration

Configure Fabric Mirror as Target

Please fill in the connection details so we can check if Striim is able to reach your Fabric Mirror target

Connection details

Mirror Database Name*
SalesDatabase

Please choose the flow to create a new connection profile to Fabric *

Select a connection profile

+ New Connection Profile

Advanced Settings

Batch Policy
eventcount:10000,interval:5m

Connection Profile Name*
ConnectToFabric

Namespace*
admin

Authentication Type*
Resource Owner Password Credential O Auth

Username*
Kirupashankaran.nataraj@azureteststriim.onmicrosoft.c

Password*
.....

Workspace*
StoresWorkspace

Workspace Id
a63e507c-18fd-48ce-8672-6cad0b46b0db

Test

Review

Automated Pipeline Configuration

3. Target Configuration (Fabric Mirror)

CONNECTION DETAILS

Mirror Database Name SalesDatabase
Connection Profile Name admin.ConnectToFabric
Workspace Name StoresWorkspace

IL Table mapping

"sales"."dbo"."%",dbo.%"

CDC Table mapping

"dbo"."%",dbo.%"

ADVANCED SETTINGS

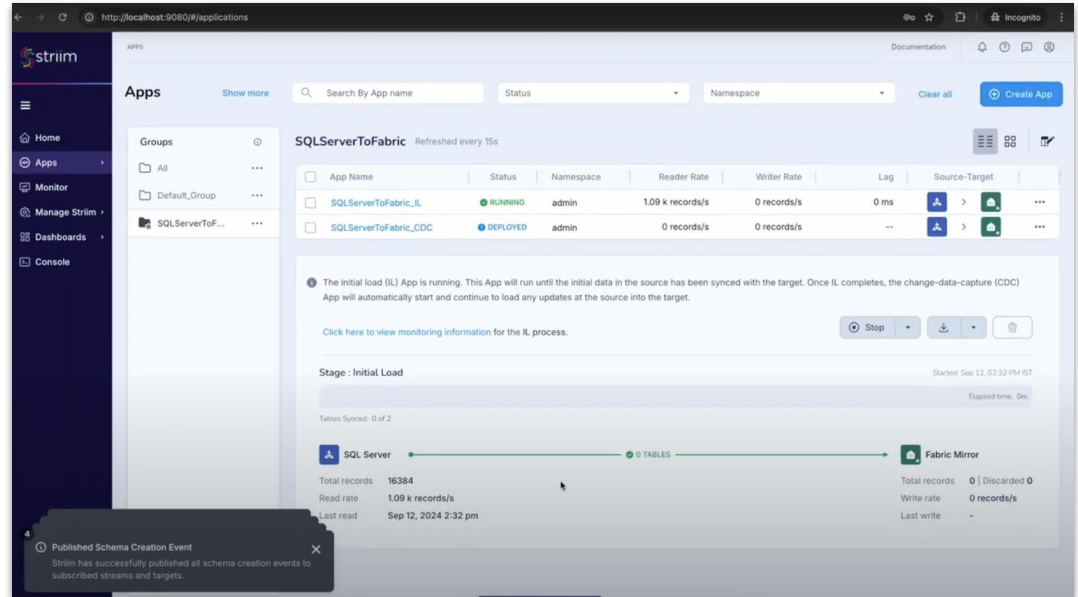
Batch Policy eventcount:10000,interval:5m

Deploy and start the automated pipeline

Cancel Save & Exit Save & Start

Monitor your automated Fabric data pipeline

1. Monitor data streaming of IL & CDC. Striim will automate setting up data pipeline from on-premise SQL server to Fabric Mirroring pair creation to your target
2. Validate the data on Fabric target replicated by Fabric mirroring service seamlessly



The screenshot displays the Striim monitoring console for the 'SQLServerToFabric' application. The interface includes a navigation sidebar on the left with options like Home, Apps, Monitor, Manage Striim, Dashboards, and Console. The main area shows a table of application components:

App Name	Status	Namespace	Reader Rate	Writer Rate	Lag	Source-Target
SQLServerToFabric_IL	RUNNING	admin	1.09 k records/s	0 records/s	0 ms	[Icons]
SQLServerToFabric_CDC	DEPLOYED	admin	0 records/s	0 records/s	--	[Icons]

Below the table, a detailed view of the 'Initial Load' stage is shown, indicating that 0 tables have been synced out of 2. A progress bar shows the 'SQL Server' source and 'Fabric Mirror' target. A notification at the bottom states: 'Published Schema Creation Event: Striim has successfully published all schema creation events to subscribed streams and targets.'