



Dynamics 365 Data Export Service Deprecation Playbook

Ensure optimal transition from Data Export Service to Azure Synapse Link for Dataverse

FastTrack for Dynamics 365 Apps



Purpose and Scope

THE SITUATION

- Data Export Service (DES) will retire by end of November 2022.
- All existing implementation using DES needs to transition to either Azure Synapse Link for Dataverse or another pattern suitable for the business need.

PURPOSE

Data Export Service (DES) was originally built to support the use cases of Operational or Enterprise BI reporting, or build a staging area for data warehouse, or build machine learning models on the Analytical data by replicating data from Dataverse into Azure SQL Database/ Azure SQL VM.

With the announcement of Azure Synapse Analytics that remove the barriers between data warehousing and big data analytics, we deprecated DES and built Azure Synapse Link for Dataverse to accelerate the time-to-insights between business applications data and analytical systems. This service is built for seamless integration with the latest analytical technologies such as Azure Data Lake Storage, Azure Synapse, Power BI etc.

In this playbook, we talk about what it takes to transition from DES to Azure Synapse Link. We also discuss some of the architectural patterns and customer scenarios that help you take right decisions to transition out of DES.

Audiences:

-  Customer Solution Architects
-  Partner Solution Architects

In summary, this guide will help you with:

- ✓ Awareness of challenges for IT.
- ✓ Understanding of tools and techniques that would ensure an optimal transition.

AZURE SYNAPSE LINK FOR DATAVERSE

With a few clicks, you can...

- Link your Dataverse environment with Azure Synapse Analytics workspace.
- Visualize your Dataverse data in Azure Synapse Analytics.
- Rapidly start processing the data to discover insights using advanced analytics capabilities for serverless data lake exploration, code-free data integration, data flows for extract, transform, load (ETL) pipelines, and optimized Apache Spark for big data analytics.

INSIDE THIS PLAYBOOK:

Chapter 1: INITIALIZATION

What is ASL for Dataverse?

What is the feature set?

Chapter 2: PLANNING

Assess existing DES usage

Plan for a transition

Chapter 3: ADOPTION

Implementation

Monitoring

Appendix

Available Resources

DES Data Export Service*

Synapse Link Azure Synapse Link for Dataverse*



Chapter 1 Initialization

What is Azure Synapse Link for Dataverse?

- Azure Synapse Link for Dataverse enables you to get near real-time insights on your data in Microsoft Dataverse.

Why transition from DES to Azure Synapse Link?

- DES technology has been announced to be deprecated.
- End of support is November 2022.

What's new with Azure Synapse Link for Dataverse?

- Intuitive UI provides a frictionless experience in enabling data pipeline and democratizing Dataverse data.
- CDM Metadata updates available across all Synapse computes.
- Accelerate time to insights via built-in capabilities for data ingestion.
- Create enterprise scale Power BI reports on Dataverse data.
- Solution aware through Dataverse solution deployment.

Resources

Chapter 2 Planning

Assess Existing DES Usage Scenarios

- Understand why DES is used to effectively transition to ASL without disruption.
- Explore sample transition scenarios to effectively plan target state architecture.

Plan for a Transition

- Assess business impact
- Define Scope
- Define dependencies
- Defining ownerships
- User communications
- Testing Strategy
- Cut-over and go-live Strategy
- Pilot the transition

Resources

Chapter 3 Adoption

Implementation

- Setup your Azure Synapse Workspace and Data Lake Storage Gen2.
- Configure the Azure Synapse Link in the Power Apps Maker Portal.
- Implement integrations to sync the data from Azure Synapse Analytics to your pre-existing data warehousing solution.
- Pilot the implementation with a few tables to ensure that end-to-end scenarios are working.
- Migrate your integration scenarios to Azure Synapse Analytics.
- Test your end-to-end integration scenarios.
- Deprecate the Data Export Service Profile.

Monitoring

- Ensure all the tables have completed syncing and there are no failure notifications.
- Transfer any existing external monitoring on DES to Azure Synapse Link for Dataverse.

Resources

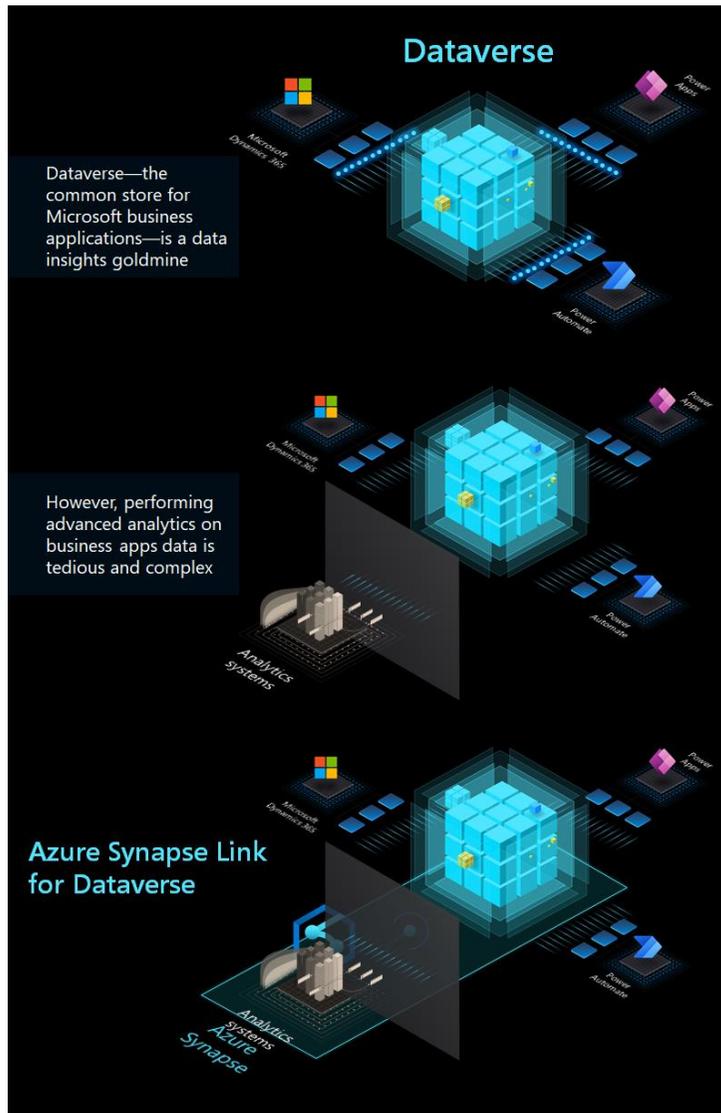


Chapter 1: Initialization



What is Azure Synapse Link for Dataverse?

With a few clicks, bring Dataverse data to Azure Synapse Analytics and enable limitless analytics.



Data is the new oil; when collated, analyzed, and connected with relevant data efficiently, you can unfold insights that drive business decisions for tomorrow. More and more customers using Dynamics 365 and modern business applications are generating massive amounts of business data that is securely stored and managed by Dataverse.

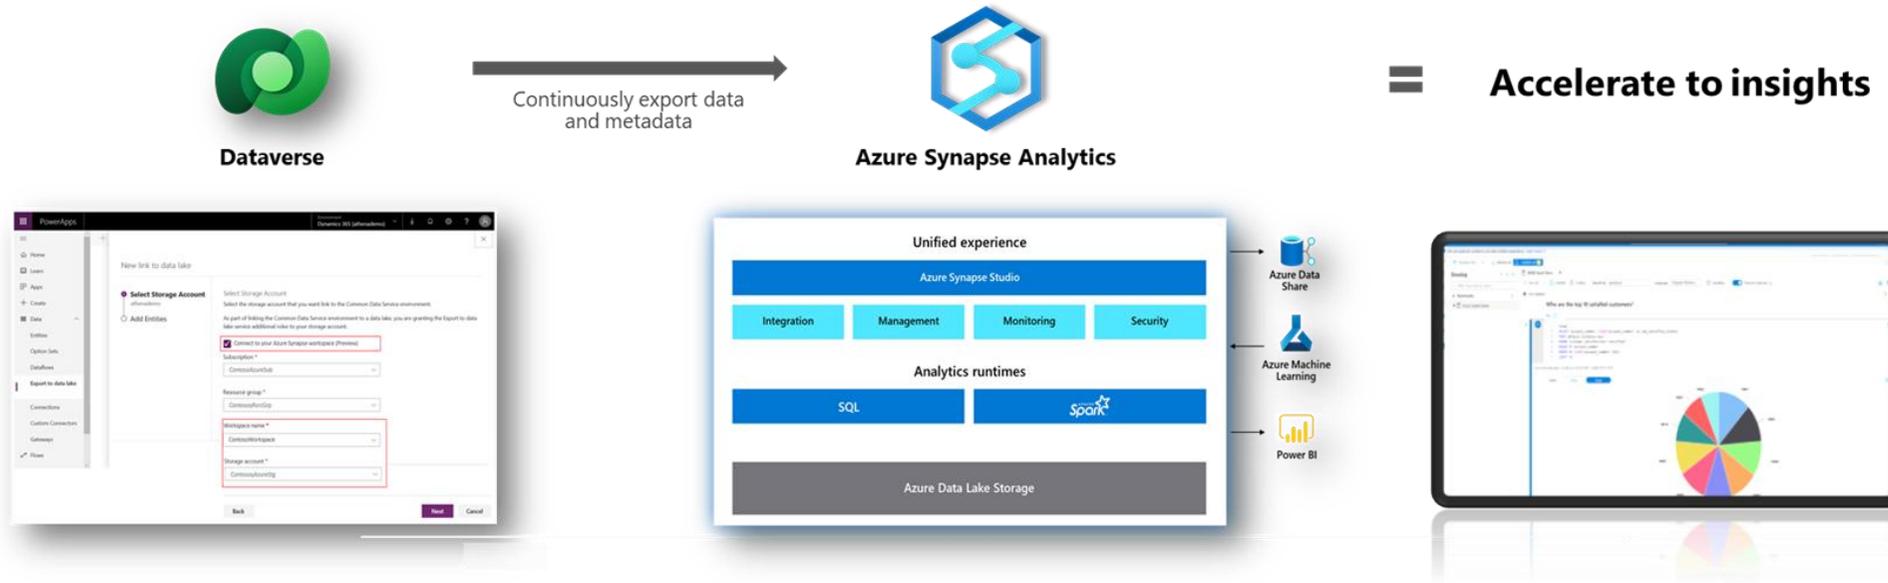
To further empower our customer's digital transformation journey, there is a need to democratize this data, provide the necessary tooling, and create end-to-end way to work with data in Dataverse that can also scale to the needs of our Enterprise customers.

[Azure Synapse Link for Dataverse](#) is the fastest path to success for our customers who want to run artificial intelligence and machine learning, integrate with external dataset, and slice and dice large volumes of Dataverse data.

Azure Synapse Link for Dataverse enables seamless integration of Dataverse with Azure Synapse and empowers customers to analyze data in the lake. Enterprise customers are now able to use the familiarity of T-SQL to analyze big data and gain insights from it, while optimizing their data transformation pipeline to leverage the deep integration of Azure Synapse with other Azure services such as Power BI Embedded, Azure Cosmos DB, Azure Machine Learning, and Azure Cognitive Services.

Our Vision

Truly democratize Dataverse data and enable advanced analytics to accelerate time to insights.



In conversations with the community, many have asked for ***a single, end-to-end way to work with data in Microsoft Dataverse, run AI and machine learning, integrate with external datasets, and slice and dice large volumes of Dataverse data.***

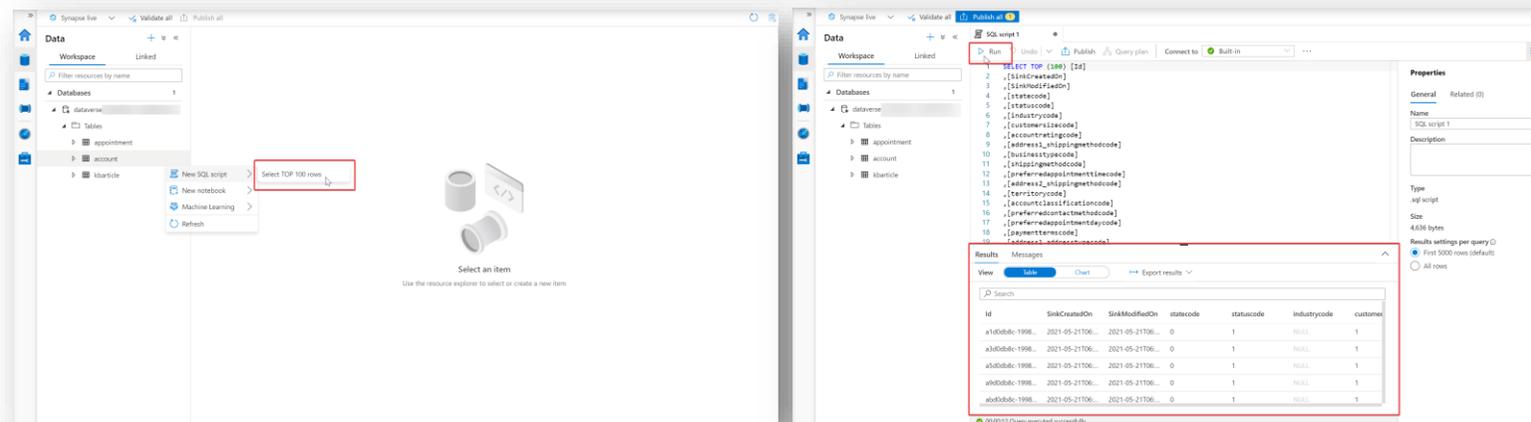
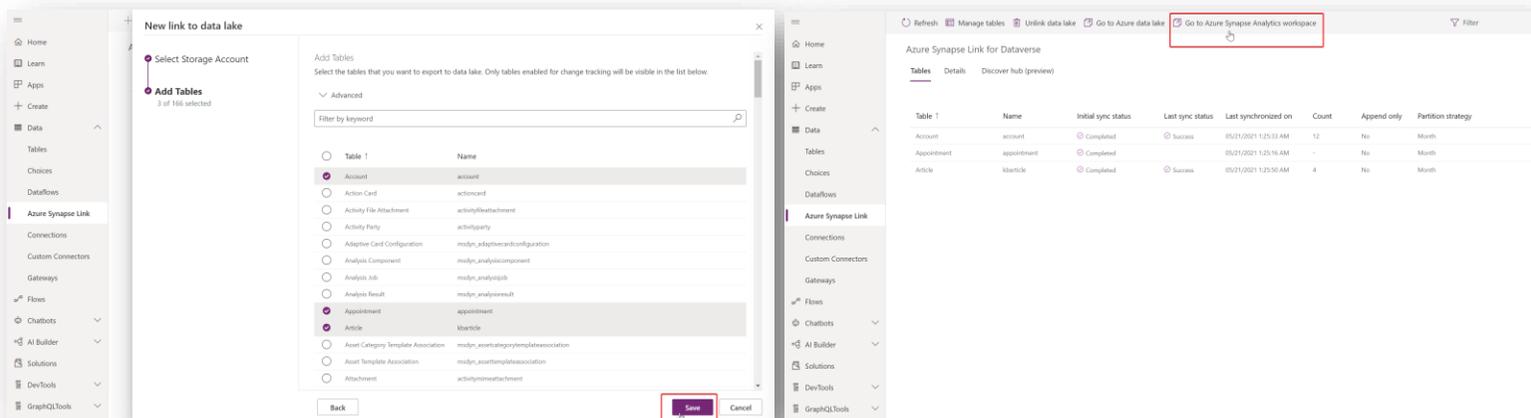
Now, instead of using multiple tools to get the job done, you can accelerate time to insight with a single comprehensive solution that can help you deliver on these goals end-to-end; something that is built-in and available out-of-the-box

Currently Released Features

Transition to Azure Synapse Link for Dataverse to leverage exciting new features.

"Do more than just moving data to a data store"

Azure Synapse Link for Dataverse truly democratizes your Dataverse data and enables you to run advanced analytics to derive deep insights on your business data!



- Integrated with Power Apps maker portal
- Easily link Dataverse environment to Azure Synapse Analytics workspace
- Continuous replication of standard and custom tables to Azure Synapse Analytics
- One-click bridge to Azure Synapse Analytics workspace
- CDM Metadata updates available across all Synapse computes
- Ready to run Synapse Serverless SQL query experience
- Ready to run Synapse Apache Spark experience
- Ready to run Synapse Pipelines experience
- Solution aware which enables seamless ALM
- Cost effective

Intuitive UI provides a frictionless experience in enabling data pipeline and democratizing Dataverse data.

Integrated with Power Platform Maker Portal

Easily link Dataverse environment to Azure Synapse Analytics workspace

The screenshot displays the 'New link to data lake' wizard in the Power Apps interface. The left sidebar shows the navigation menu with 'Azure Synapse Link' highlighted. The main content area shows the 'Select Storage Account' step for the environment 'dataversesynapsestrg'. The wizard prompts the user to select a storage account in the same region as the environment (Central US). The 'Connect to your Azure Synapse Analytics workspace (Preview)' option is selected. The following fields are populated: Subscription (* DV-Synapse-sub), Resource group (* DVSynapseRsrc), Workspace name (* dataverse-synapse-workspace), and Storage account (* dataversesynapsestrg). A 'Next' button is visible at the bottom right.

Microsoft Power Apps

Environment: synapsn1testing

New link to data lake

Azure Synapse Link for Dataverse

Select Storage Account
dataversesynapsestrg

Add Tables

Select Storage Account

Select the storage account that you want link to the Dataverse environment. The storage account must be in the same region as your environment.

Your environment is located in: Central US

Please attach a storage account in one of the following location(s): Central US, Central US EUAP or South Central US

Connect to your Azure Synapse Analytics workspace (Preview) ⓘ

Subscription *
DV-Synapse-sub

Resource group *
DVSynapseRsrc

Workspace name *
dataverse-synapse-workspace

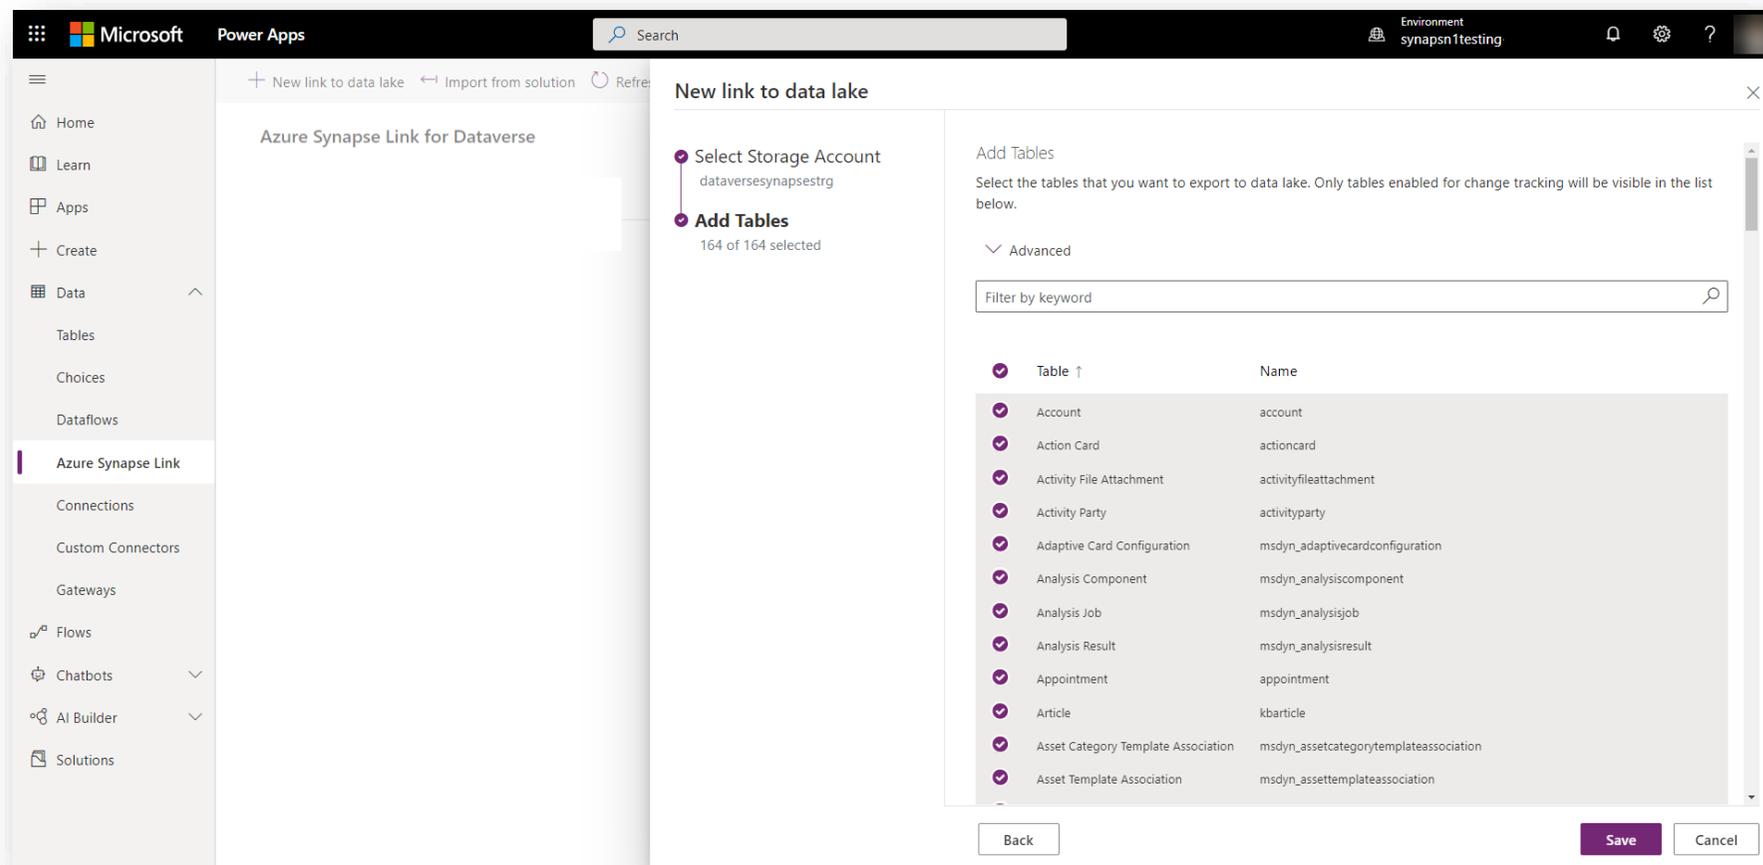
Storage account *
dataversesynapsestrg

As part of linking the Dataverse environment to a data lake, you are granting the Export to data lake service additional roles to your storage account. By using the Export to data lake service, you agree that data may go outside of Power Apps' compliance boundary. For more information, see the Privacy Notice in the technical documentation for the service.

Back Next Cancel

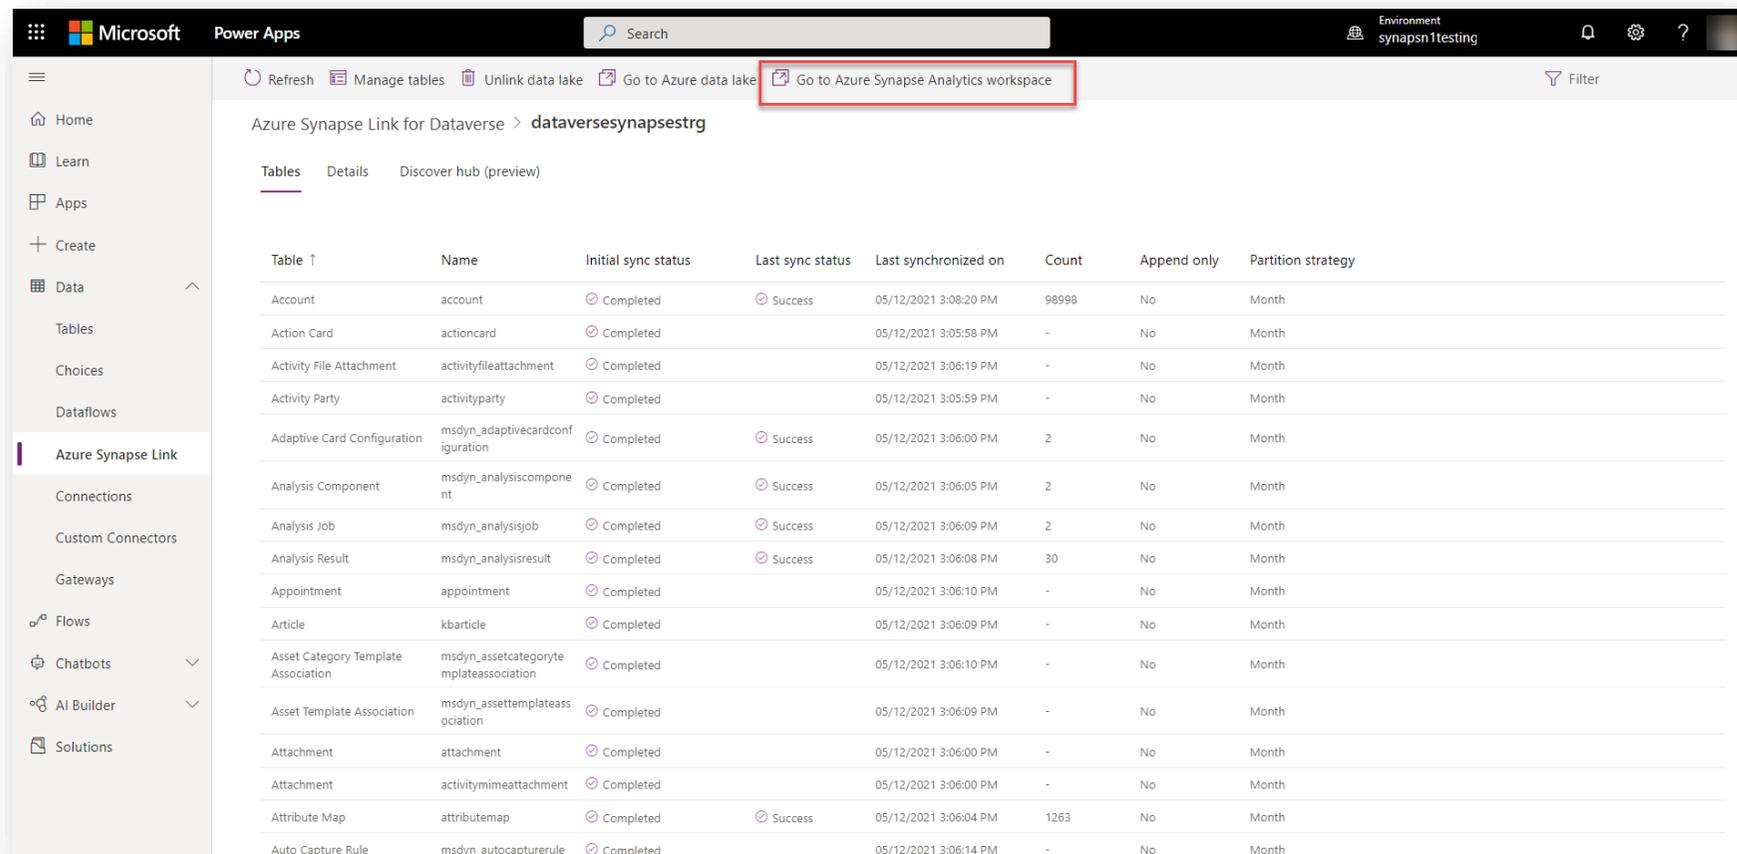
Intuitive UI provides a frictionless experience in enabling data pipeline and democratizing Dataverse data.

Continuous replication of Standard and Custom Dataverse tables to Azure Synapse Analytics



Intuitive UI provides a frictionless experience in enabling data pipeline and democratizing Dataverse data.

One-click bridge to
Azure Synapse Analytics
workspace



The screenshot displays the Microsoft Power Apps interface for the 'dataversesynapsestrg' workspace. The left sidebar shows navigation options, with 'Azure Synapse Link' selected. The main content area shows a table of data tables with columns for Name, Initial sync status, Last sync status, Last synchronized on, Count, Append only, and Partition strategy. The 'Go to Azure Synapse Analytics workspace' button is highlighted with a red box.

Table ↑	Name	Initial sync status	Last sync status	Last synchronized on	Count	Append only	Partition strategy
Account	account	Completed	Success	05/12/2021 3:08:20 PM	98998	No	Month
Action Card	actioncard	Completed		05/12/2021 3:05:58 PM	-	No	Month
Activity File Attachment	activityfileattachment	Completed		05/12/2021 3:06:19 PM	-	No	Month
Activity Party	activityparty	Completed		05/12/2021 3:05:59 PM	-	No	Month
Adaptive Card Configuration	msdyn_adaptivecardconfiguration	Completed	Success	05/12/2021 3:06:00 PM	2	No	Month
Analysis Component	msdyn_analysiscomponent	Completed	Success	05/12/2021 3:06:05 PM	2	No	Month
Analysis Job	msdyn_analysisjob	Completed	Success	05/12/2021 3:06:09 PM	2	No	Month
Analysis Result	msdyn_analysisresult	Completed	Success	05/12/2021 3:06:08 PM	30	No	Month
Appointment	appointment	Completed		05/12/2021 3:06:10 PM	-	No	Month
Article	kbarticle	Completed		05/12/2021 3:06:09 PM	-	No	Month
Asset Category Template Association	msdyn_assetcategorytemplateassociation	Completed		05/12/2021 3:06:10 PM	-	No	Month
Asset Template Association	msdyn_assettemplateassociation	Completed		05/12/2021 3:06:09 PM	-	No	Month
Attachment	attachment	Completed		05/12/2021 3:06:00 PM	-	No	Month
Attachment	activymimeattachment	Completed		05/12/2021 3:06:00 PM	-	No	Month
Attribute Map	attributemap	Completed	Success	05/12/2021 3:06:04 PM	1263	No	Month
Auto Capture Rule	msdyn_autocapturerule	Completed		05/12/2021 3:06:14 PM	-	No	Month

Accelerate time to insights via built-in capabilities for data ingestion, data preparation, and machine learning.

Ready-to-run Synapse Serverless SQL query experience

The screenshot displays the Microsoft Azure Synapse Analytics interface. The left sidebar shows navigation options: Home, Data (highlighted with a red box), Develop, Integrate, Monitor, and Manage. The main area is divided into a 'Data' pane on the left and a 'SQL script 1' editor on the right. The 'Data' pane shows a tree view of databases and tables, with 'account' highlighted by a red box. The SQL editor contains the following query:

```
1 SELECT TOP (100) [name]
2 , [SinkCreatedOn]
3 , [SinkModifiedOn]
4 , [statecode]
5 , [statuscode]
6 , [preferredcontactmethodcode]
7 , [accountratingcode]
8 , [businessstypecode]
9 , [shippingmethodcode]
10 , [preferredappointmenttimecode]
11 , [address_shippingmethodcode]
```

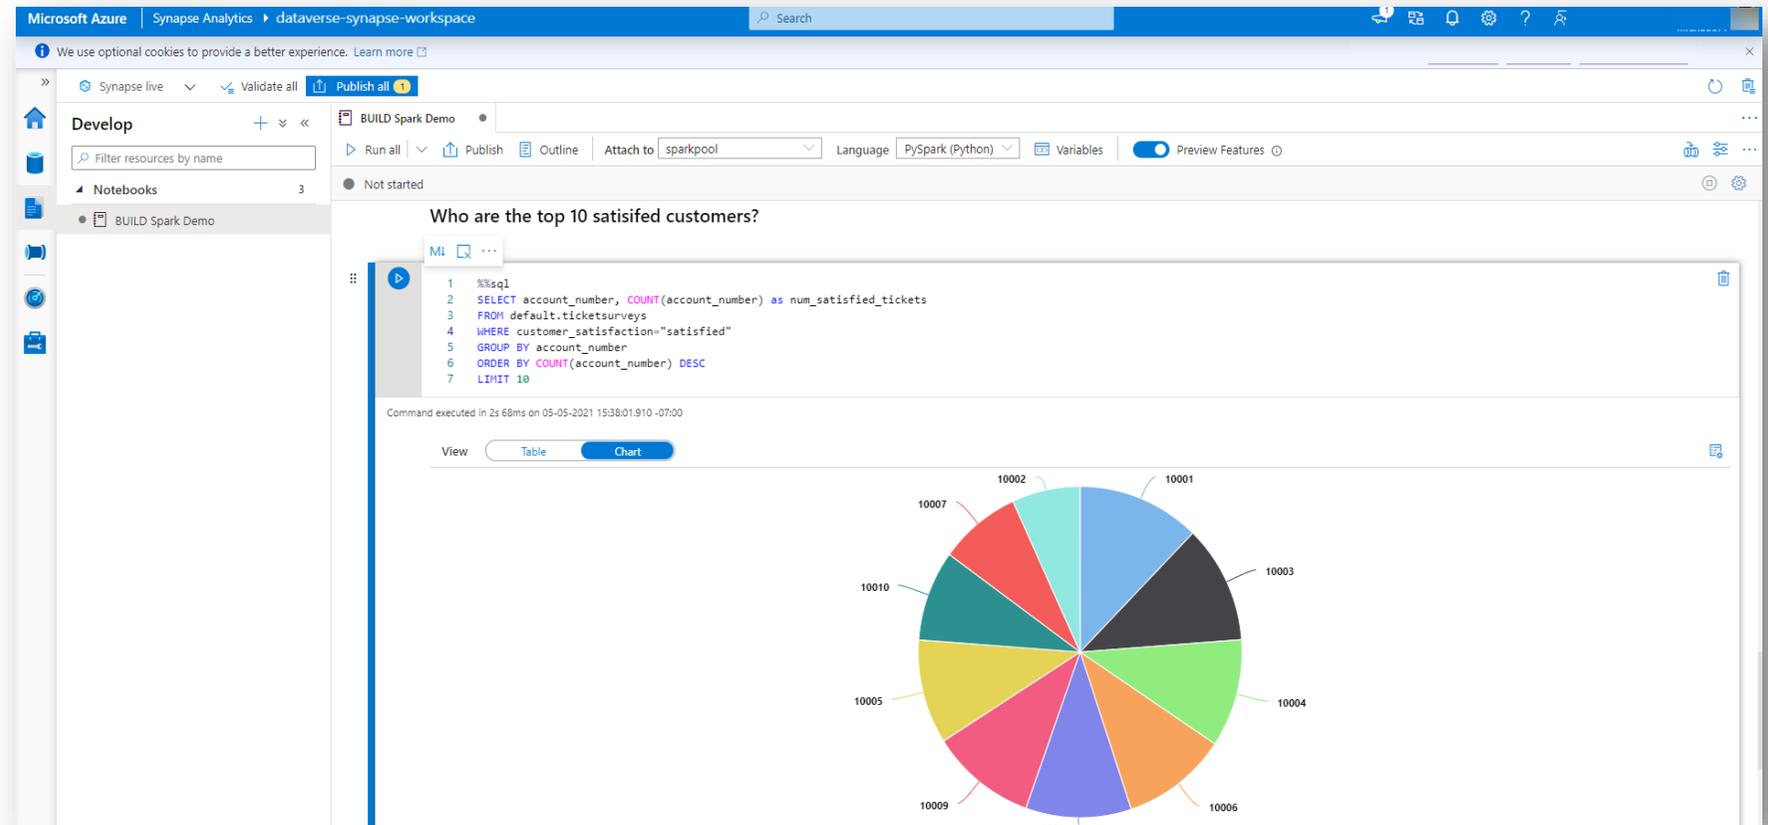
The 'Results' pane shows a table with the following data:

name	SinkCreatedOn	SinkModifiedOn	statecode	statuscode	preferredconta...	accountra
Fourth Coffee (sample)	2021-05-12T22:...	2021-05-12T22:...	0	1	1	1
Litware, Inc. (sample)	2021-05-12T22:...	2021-05-12T22:...	0	1	1	1
Adventure Works (sample)	2021-05-12T22:...	2021-05-12T22:...	0	1	1	1
Blue Yonder Airlines (sample)	2021-05-12T22:...	2021-05-12T22:...	0	1	1	1
City Power & Light (sample)	2021-05-12T22:...	2021-05-12T22:...	0	1	1	1

The status bar at the bottom indicates: 00:00:05 Query executed successfully.

Accelerate time to insights via built-in capabilities for data ingestion, data preparation, and machine learning.

Ready-to-run Synapse
Apache Spark experience



The screenshot displays the Microsoft Azure Synapse Analytics interface. The top navigation bar shows "Microsoft Azure | Synapse Analytics | dataverse-synapse-workspace". The left sidebar contains navigation icons and a "Develop" section with a search bar and a list of notebooks, including "BUILD Spark Demo". The main workspace shows a PySpark notebook titled "BUILD Spark Demo" with a SQL query:

```
1 %%sql
2 SELECT account_number, COUNT(account_number) as num_satisfied_tickets
3 FROM default.ticketsurveys
4 WHERE customer_satisfaction="satisfied"
5 GROUP BY account_number
6 ORDER BY COUNT(account_number) DESC
7 LIMIT 10
```

The query was executed in 2s 68ms on 05-05-2021 15:38:01.910 -07:00. Below the code, there are "View" options for "Table" and "Chart". The "Chart" view is selected, displaying a pie chart with 10 segments, each labeled with an account number: 10001, 10003, 10004, 10006, 10008, 10009, 10005, 10010, 10007, and 10002.

Accelerate time to insights via built-in capabilities for data ingestion, data preparation, and machine learning.

Ready-to-run Synapse Pipelines experience

The screenshot displays the Microsoft Azure Synapse Analytics interface. The top navigation bar includes "Microsoft Azure", "Synapse Analytics", and "contoso". A search bar and several utility icons are also present. Below the navigation bar, there is a cookie consent banner and a toolbar with "Synapse live", "Validate all", and "Publish all" buttons. The left sidebar shows a "Data" section with "Workspace" and "Linked" tabs, a search box for resources, and a tree view of databases and tables. The main area shows a "Dataflow1" configuration. The "Source settings" tab is active, displaying the following fields:

- Output stream name ***: source1
- Source type ***: Workspace DB (highlighted with a red box)
- Database ***: dataverse_contoso_unq2b40367a1fe... (highlighted with a red box)
- Table ***: account
- Sampling ***: Disable (selected)

The "Source settings" tab also includes a "Refresh" button and a "Learn more" link. The "Source options" tab is also visible, showing a "Data flow debug" toggle.

Limitless analytics at incredible value.

380% faster and costs up to 49% less than other cloud providers

Azure Synapse Analytics pricing

[Request a pricing quote](#) [Try Azure for free](#)

[Overview](#) [Pricing table](#) [Purchase options](#) [Resources](#) [FAQs](#) [Azure Synapse Analytics >](#)

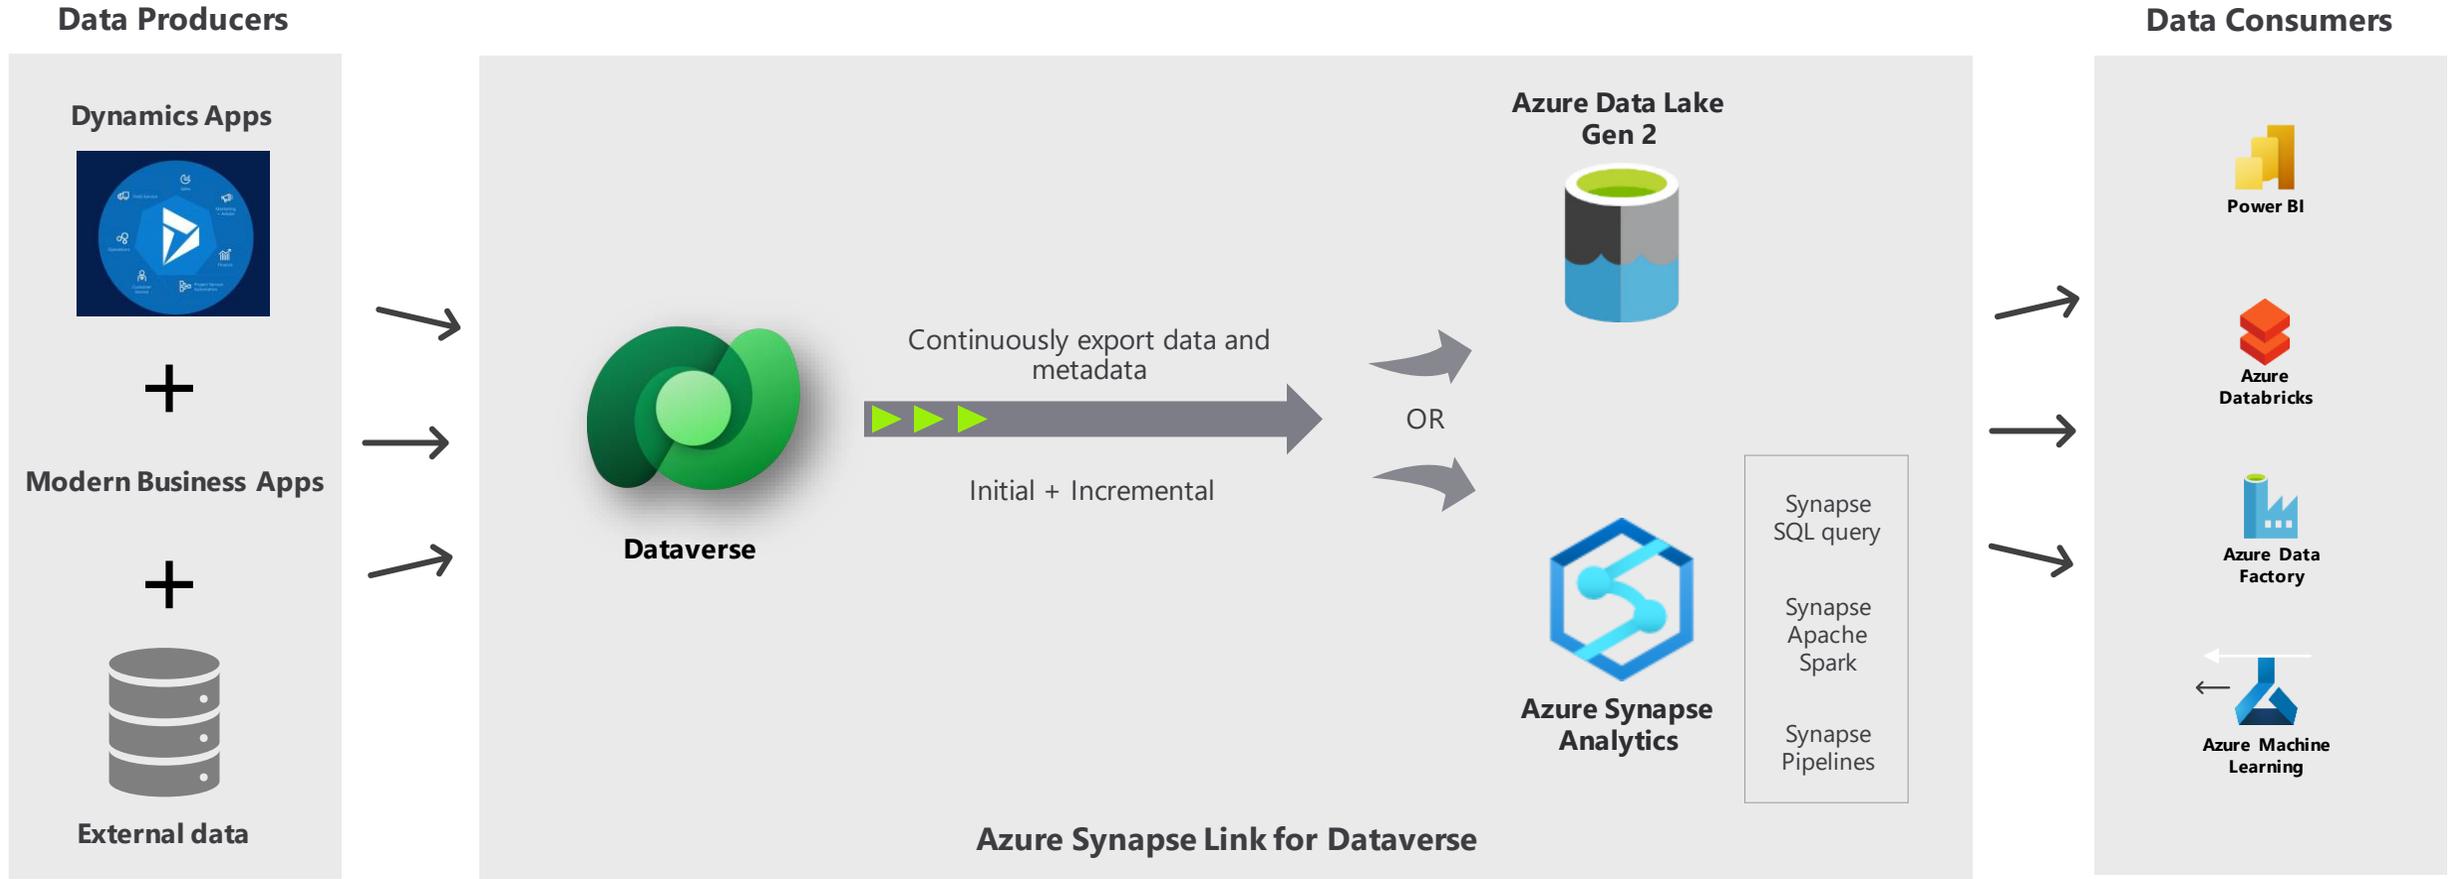
Limitless analytics service with unmatched time to insight

Azure Synapse Analytics is a limitless analytics service that brings together data integration, enterprise data warehousing, and big data analytics. It gives you the freedom to query data on your terms, using either serverless or dedicated resources—at scale. Azure Synapse brings these worlds together with a unified experience to ingest, explore, prepare, manage, and serve data for immediate BI and machine learning needs.

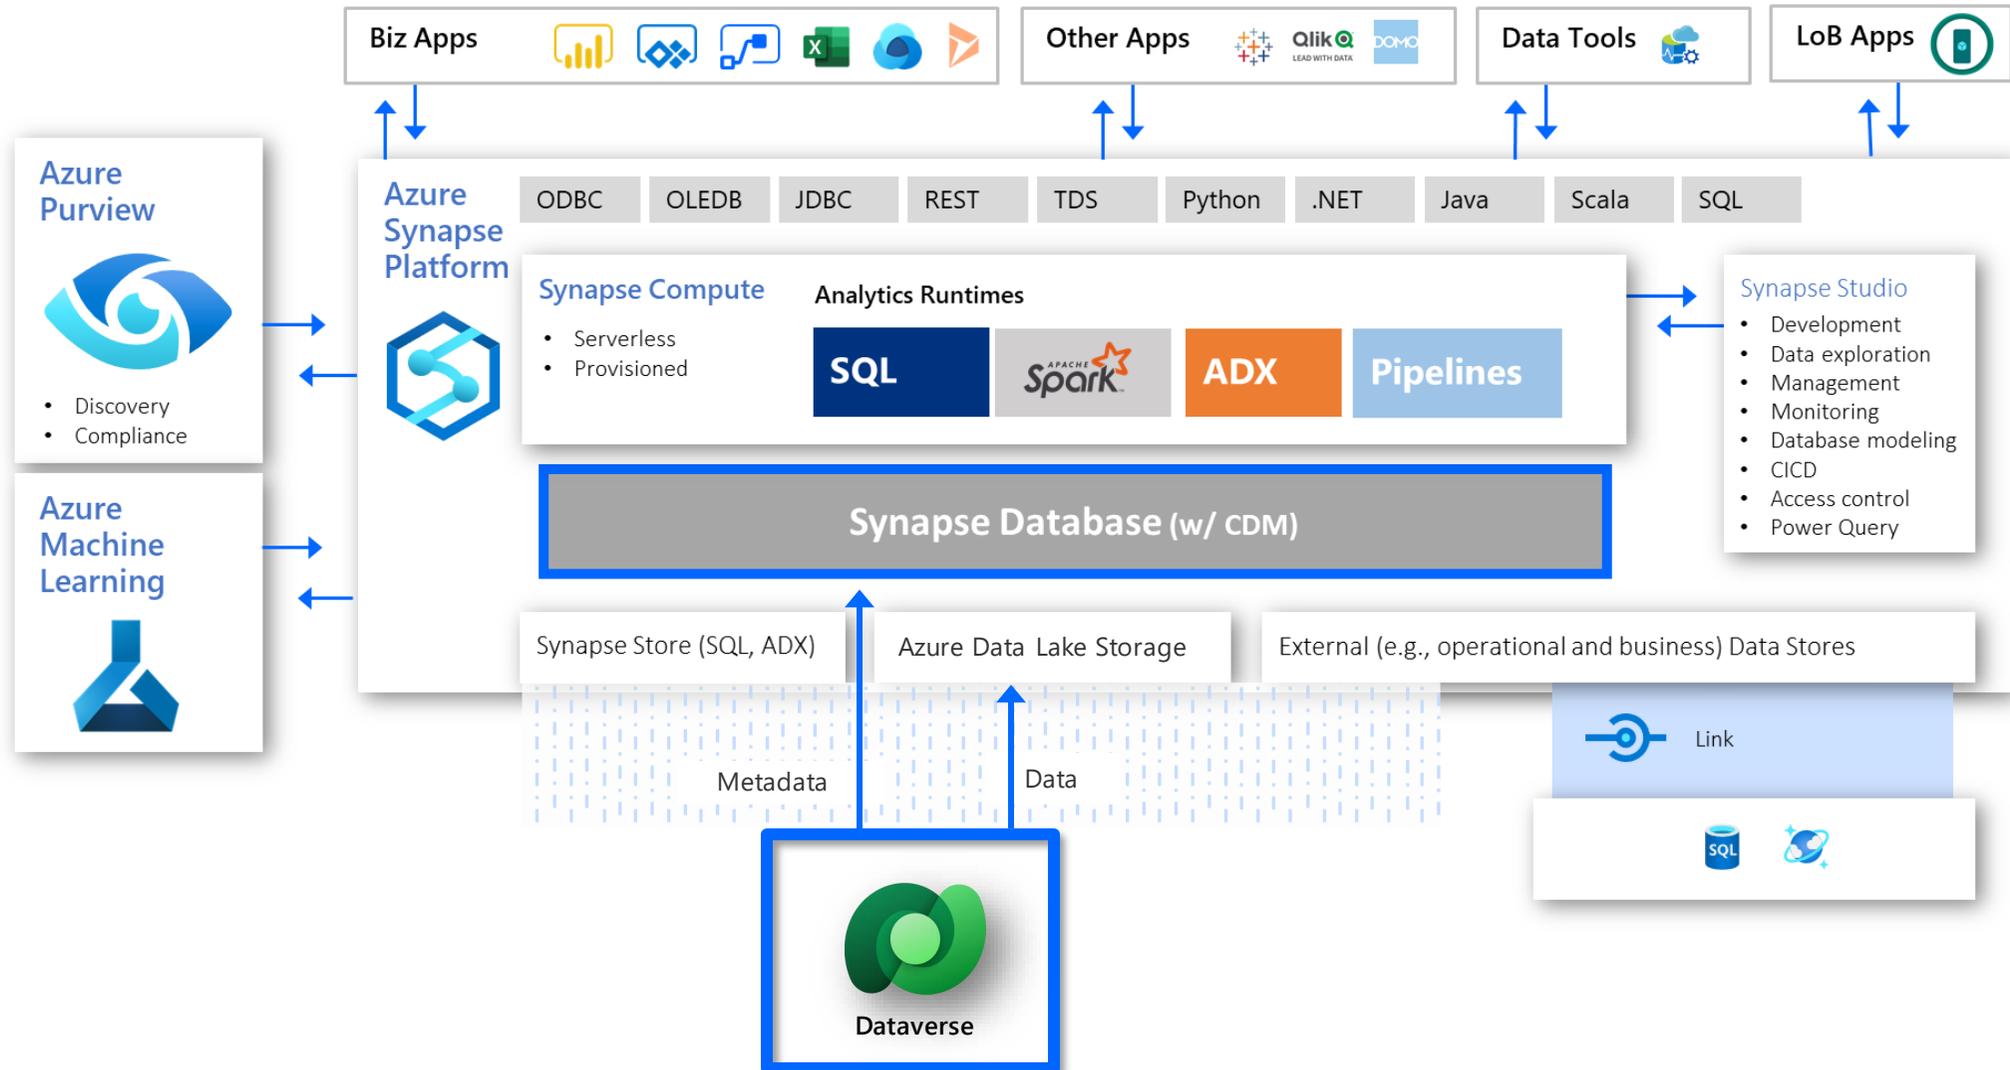
When using Azure Synapse, you will only pay for the capabilities you opt in to use. Beginning January 1, 2022 there will be a cost when enabling additional workspace features, such as a Managed Virtual Network. For customers that are using Azure Synapse workspace before January 1, 2022 and have enabled t

[Chat with Sales](#)

Azure Synapse Link for Dataverse



Azure Synapse Link for Dataverse Architecture Diagram



Understand the availability and reliability of Synapse Link.

Using Azure Synapse Link, we currently support two options to export data from Dataverse:

1. Export Dataverse data to Azure Data Lake Storage Gen2 (General Availability since February 2020)

- Applicable to customers who are using ADLS Gen2 as a staging area.
- Collate data from multiple sources to ADLS Gen2 for downstream integration.
- Reader reads data from read-only snapshot folder(s) updated every hour.

Note: Since GA, 4B rows of CDM data is pushed to ADLS Gen2 each month on average.

2. Export Dataverse data to Azure Synapse Analytics (Generally Available since November 2021)

- Near real-time replication to Azure Synapse (soft SLA of 15 mins for delta syncs when service operates under specific conditions).
- Azure Synapse Analytics provides the necessary compute needed to run Enterprise scale queries and create Enterprise scale Power BI reports on Dataverse data.
- Tight integration with Azure Synapse Analytics provides several advantages that were not available in DES:
 - Limitless analytics service that brings together data integration, enterprise data warehousing, and big data analytics.
 - Ingest, explore, prepare, manage, and serve data for immediate business intelligence and machine learning needs all from a single service.
 - Freedom to query data on your terms, using either serverless or dedicated resources—at scale. It also allows you to build analytics solutions on top of the Apache Spark engine.

While exporting Dataverse data to ADLS Gen2 provides flexibility to stage data and later move to the desired destination, exporting to Azure Synapse Analytics is highly recommended since it provides an end-to-end way to work with data in Microsoft Dataverse. After the export to Azure Synapse Analytics flow is Generally Available, customers will reap more benefits as we continue to make significant investments in this flow to further streamline consumption and analytical patterns.

High level overview of cost for using Azure Synapse Link for Dataverse.

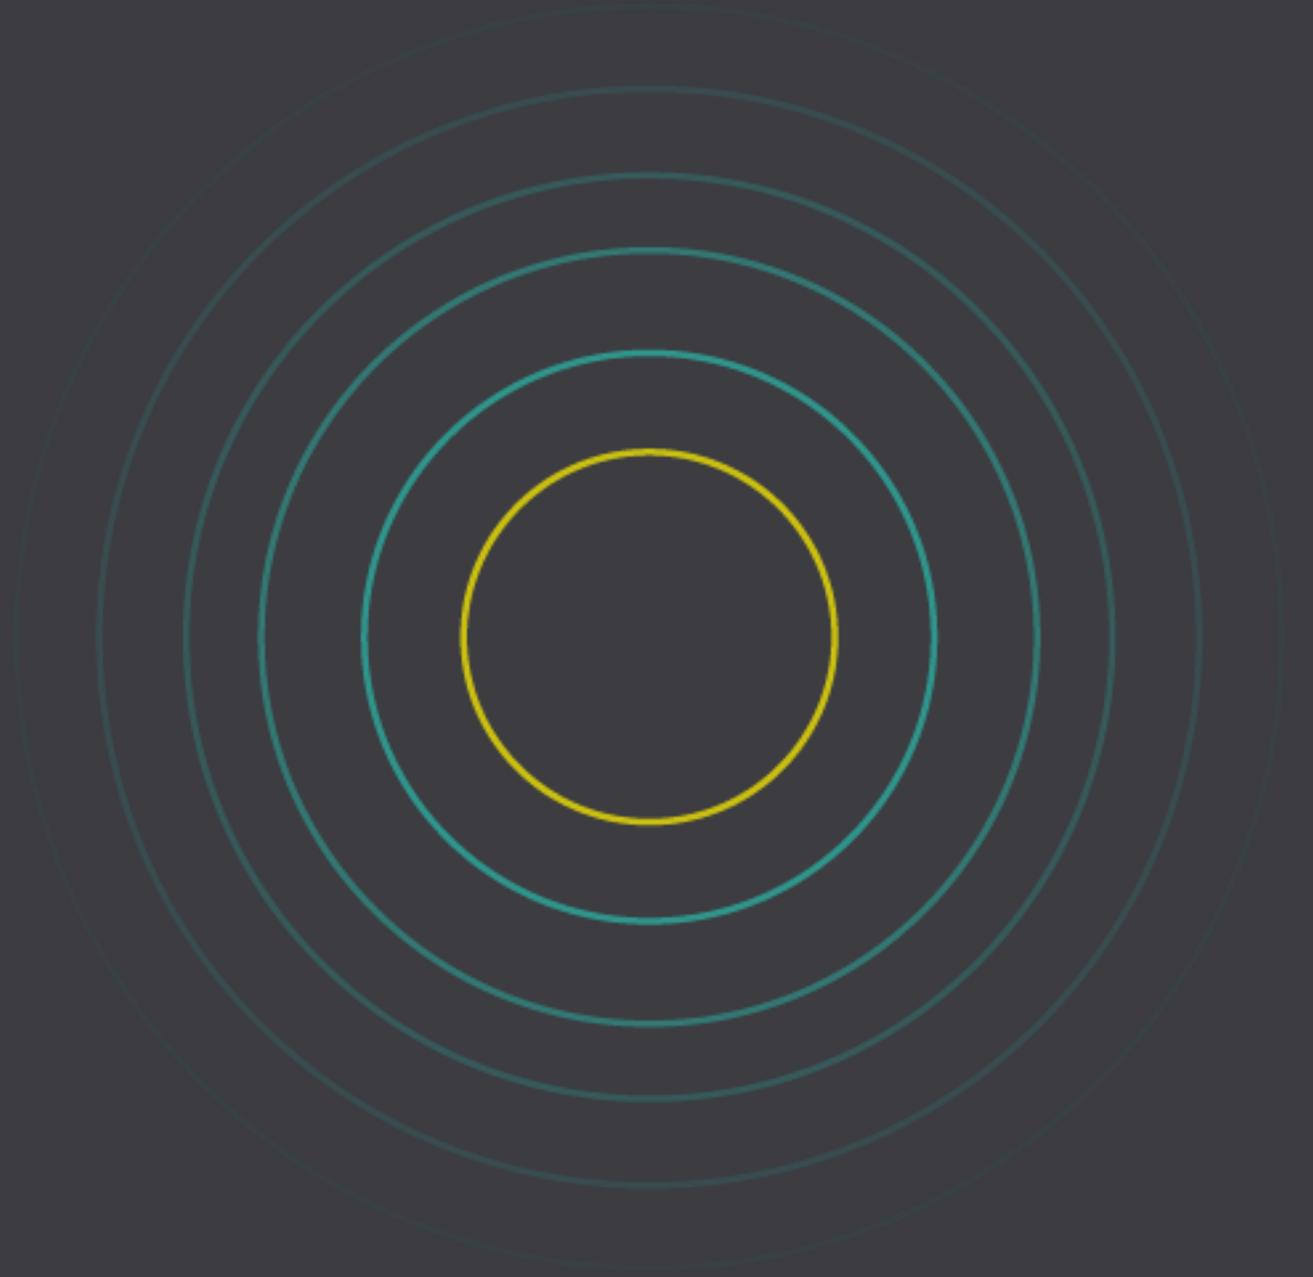
1. Azure Data Lake Storage Pricing

- Scalable, secure data lake for high-performance analytics.
- Pay-as-you-go pricing model.
- More information [here](#).

2. Azure Synapse Analytics Pricing

- Limitless analytics service with unmatched time to insight.
- Pay-as-you-go pricing model.
- More information [here](#).

Chapter 2: Planning

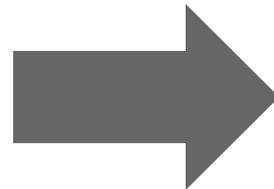


Assess Existing Data Export Service Scenarios

Considerations

Before transitioning to Azure Synapse Link for Dataverse, it is important to understand the current usage patterns of Data Export Service, so that you can effectively plan to move to Azure Synapse Link for Dataverse.

The following sample questions will help you identify key risk areas in the data consumption scenarios and align to the new transition points.



What are the key business requirements for Data Export Service?

Having a clear understanding of the data use cases will help you understand the key purposes of using DES and how to transition to Azure Synapse Link for Dataverse. For example, a scenario can be for analysis, feeding data to line of business applications, reporting, complex data transformations, or a combination of these.

Do you have a high-level architecture diagram of how Data Export Service data is consumed from Azure SQL?

Having this visualization will help you understand the different components used in DES Service and how effectively these components can be transitioned to Azure Synapse Link for Dataverse.

When using Azure SQL, are there any SQL Stored Procedures, functions, or triggers used to transform the exported data in Azure SQL?

This will help you understand if there are data transformations happening and how to approach that Azure Synapse Link for Dataverse. Review the scenarios section for more guidance.

Are there any custom views on the exported data that are being utilized in Power BI or other consumers?

Having the granular information around how the exported data is being consumed helps you identify all the points that need transition and reduce failures.



Please refer to the 'Anti-Patterns' section in Appendix for usage scenarios where neither DES nor Azure Synapse Link are recommended.



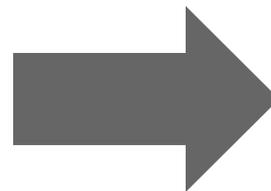
Identify all the data consumers and transformations happening on your exported Dataverse data to align the implementation to Azure Synapse Link for Dataverse.



Considerations

At this stage, you have already identified the key scenarios and use cases for Data Export Service. For any transition project to be successful, there are a few steps that need to be followed, and transitioning to Azure Synapse link for Dataverse is no exception.

In this section, we will discuss some important considerations for a successful transition to Azure Synapse Link for Dataverse.



Step 1

Step 2

Step 3

Step 4

Step 5

Step 6

Step 7

Step 8

Step 1: Assess Business Impact

As you have already identified the key scenarios, it is important to understand the business impact of the transition. Typically, Data Export Service is used to feed Dataverse data to downstream analytical systems. Although this is the most common use case, there are scenarios where exported data is also used for line of business applications and as a mode of integration, so unavailability of exported data for even a brief period can lead to critical business impact across the organization. There are also customer scenarios where Dynamics 365 data needs to be available on-premise on a regular interval to comply with the company compliance policy.

To assess the right business impact, consider the following when planning for a transition:

- Disruption over sales
- Violations with compliance
- Impact on external vendors
- Impact on other line of business applications

Step 2: Define Scope

A well-defined project scope should consider all aspects of the project including, but not limited to, all activities and deliverables, timeline, budget, resources, and key stakeholders. When you are planning to transition from DES to Azure Synapse link for Dataverse, it is critical to ensure a well-defined scope is in place before starting the transition activities. We also recommend assessing the business requirements for using DES and architect the right pattern while transitioning to Azure Synapse Link for Dataverse.

A well-defined project statement should have the following parts:

- Breakdown of all work items
- Project constraints which can impact the overall success of the project
- Key milestones
- Assumptions
- Key stakeholders
- Scope elements



Step 1

Step 2

Step 3

Step 4

Step 5

Step 6

Step 7

Step 8

Step 3: Define Dependencies

A well-defined project plan should clearly identify and list the dependencies which might impact the overall success of the project. For transitioning to Azure Synapse Link for Dataverse, there can be internal dependencies like task completion order, resource availability, or external dependencies like availability of a specific team (i.e., reporting team for transitioning the reports). You can identify the dependencies with such as:

- Compliance requirements
- Data governance approvals
- Architecture patterns for your organization
- Work requirements for internal teams and vendors
- Conducting testing
- Approvals

Step 4: Define Ownership

A critical criteria for a project to be successful is defining clear ownership of the various project elements. For example, the transition from DES to Azure Synapse Link for Dataverse will be comprised of various project activities:

- Analysing usage
- Impacting assessment
- Defining scope
- Allocating resources and budget
- Implementation (internally or with vendors)
- Testing
- End user communication
- Approval decision followed by go-live

The project plan should include a clear ownership for each of all the key activities to ensure adequate steps are taken to mitigate any impact towards successful transition.

Step 1

Step 2

Step 3

Step 4

Step 5

Step 6

Step 7

Step 8

Step 5: Testing Strategy

Testing the transition to Azure Synapse Link for Dataverse from DES is essential to the success of its implementation. It is recommended to create a test strategy in the planning phase. Ideally, the test strategy will provide the blueprint for the approach to testing and includes the high-level overview of the plan that will be followed to support the validation of the transition.

The transition test strategy should include:

- High-level overview of the plan
- Different test types required for the transition (i.e., end-to-end, SIT, UAT, etc.)
- Approach for non-functional testing (i.e., performance, availability, and security)

Step 6: User Communication

Since you should have already concluded the dependencies, analysed the impact, and determined ownership, it is time to start drafting the user communication plan to include what, when, and how information will be shared across different teams and stakeholders.

The user communication information may include:

- Regular updates on the progress
- Major milestones
- Proposed vs. actual timeline
- Planned downtime



Step 1

Step 2

Step 3

Step 4

Step 5

Step 6

Step 7

Step 8

Step 7: Go-Live Strategy

A detailed plan for the go-live should be developed during the planning phase for a successful transition. Even though it seems to be the last activity during implementation, many of the tasks performed during go-live takes time to prepare and should be planned during the early implementation stage.

A successful go-live plan should include:

- Go-live schedule
- Detailed cut-over activity
- User communication plan
- Go/No-Go strategy
- Number of dress rehearsals
- Contingency plan

Step 8: Pilot the Transition

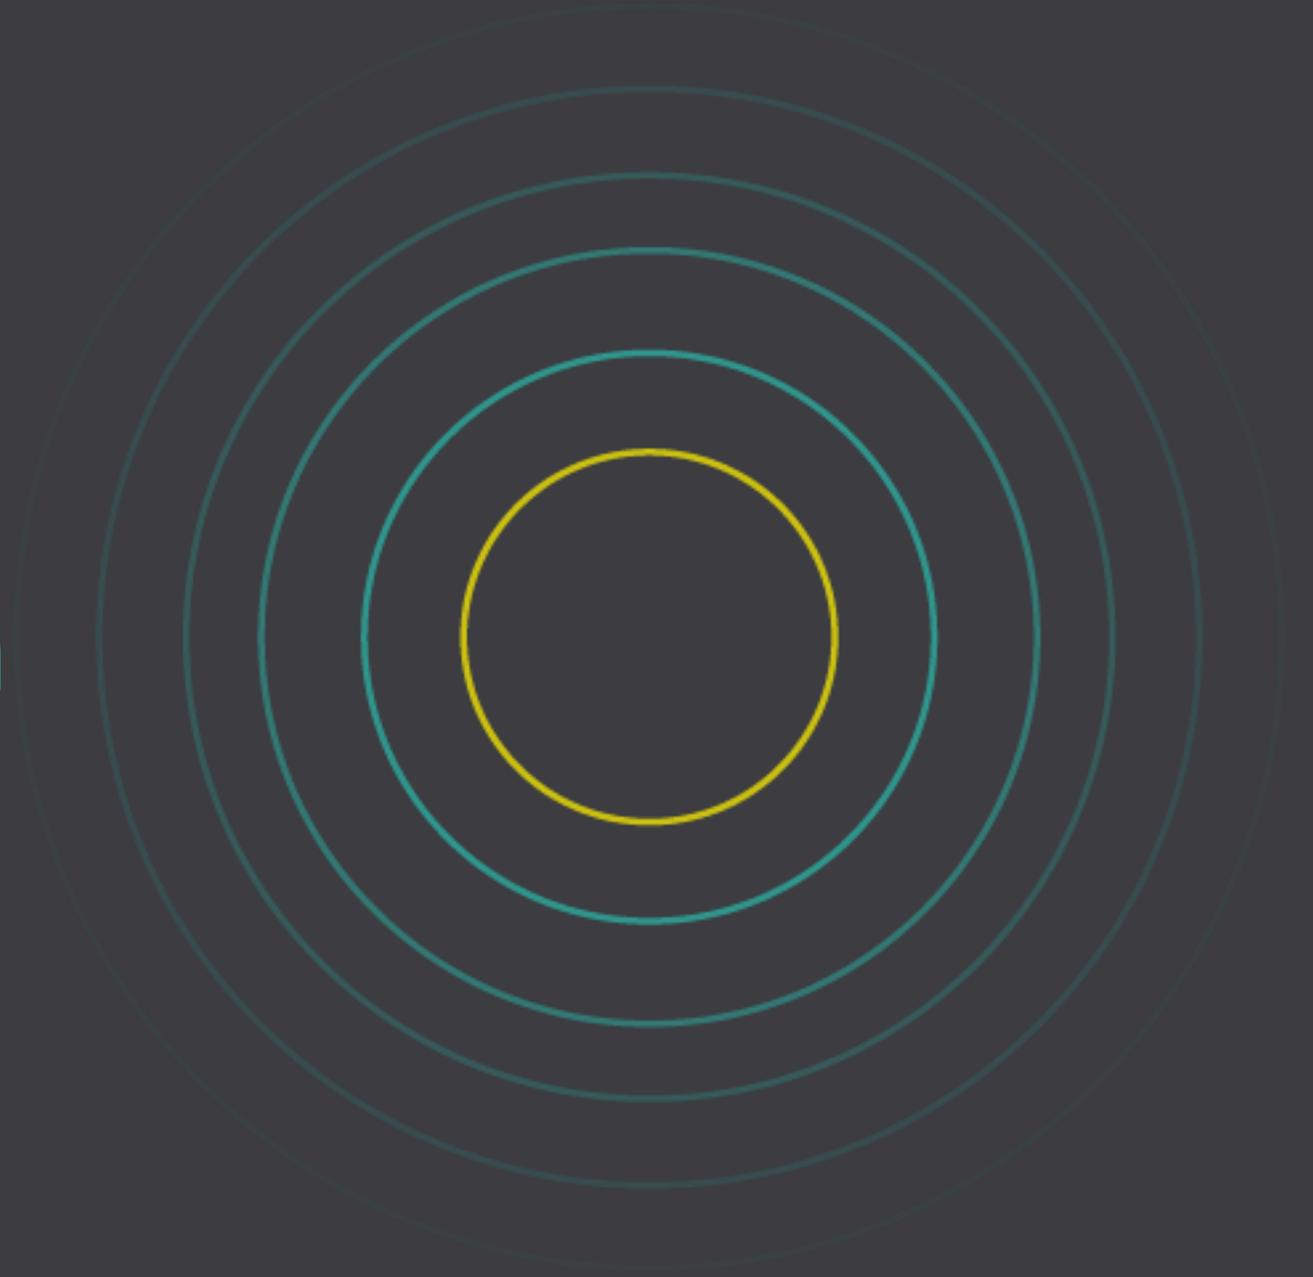
An important step in any transition project is performing a pilot. Once you have completed all the prerequisite requirements, defined the scope, obtained business approval, identified the transition scenarios, it is highly recommended to perform a pilot transition.

A pilot transition will help you validate:

- High level plan
- Transition approach
- Transition volumes and timeline
- Transition architecture

A pilot transition will also help you reduce problems during the actual transition and refine your overall execution plan.

Chapter 3: Adoption



Scenario Support Overview

		Tool						
		Serverless SQL	Apache Spark for Synapse	Power BI	Azure Data Factory	Synapse Pipelines	Dedicated SQL	Azure SQL
Dataverse Data Action	View	✓ aka.ms/aslserverless	✓ aka.ms/aslspark	✓ aka.ms/aslpbi			✓ aka.ms/asldedicated	✓ aka.ms/aslsql
	Report	✓ aka.ms/aslserverless	✓ aka.ms/aslspark	✓ aka.ms/aslpbi				
	Transform	✓ aka.ms/aslserverless	✓ aka.ms/aslspark	✓ aka.ms/aslpbi	✓ aka.ms/asladftransform	✓ aka.ms/aslpipline	✓ aka.ms/asldedicated	✓ aka.ms/aslsql
	Execute	✓ aka.ms/aslserverless			✓ aka.ms/asladfexecute			
	Analyze	✓ aka.ms/aslserverless	✓ aka.ms/aslspark	✓ aka.ms/aslpbi				



Explore Azure Synapse Link in a Sandbox Environment.

Let's begin!

Implementation

- Set up your storage account.
 - More information on how to create a Storage Account can be found [here](#).
 - More Information on Azure Data lake Storage Gen2 can be found [here](#).
- Set up Azure Synapse Analytics Workspace.
 - More Information on how to setup the Azure Synapse Analytics Workspace can be found [here](#).
- Configure your Azure Synapse Link for Dataverse from the Power Platform Maker Portal.
 - Consider which Dataverse Tables are required.
 - Consider your business requirements and create multiple profiles if necessary.
 - More information on how to Configure Azure Synapse Link with a Synapse Workspace can be found [here](#).
 - More information on how to Configure Azure Synapse Link with a Data Lake can be found [here](#).

Considerations

Timing

- Plan your transition to Azure Synapse Link to have a smooth transition and avoid any last minute's surprises.
- Transition DES to Synapse Link by November 2022.

Quick Start

- Check the [prerequisites](#).
- The Synapse workspace and your Azure Data Lake Storage Gen2 account must be in the same region.
- Refer to the Advanced configuration settings [here](#).
- Consider multiple profiles with Synapse Link for the below scenarios:
 - Multiple line of business
 - Sync to different ADL Gen2
 - Isolation for tables to comply with security

Best Practice

- Check this [playbook](#) for sample transition scenarios.
- Map your end-to-end scenarios with Synapse Link.
- Understand and adopt strategy to meet your organizational need like Monthly vs Yearly partition, Append Only mode, etc.
- Check the [FAQ](#) Azure Synapse Link for Dataverse for frequently asked questions.



Transitioning from DES? New to Azure Synapse Link for Dataverse?

Planning

Plan

- Identify the number of Dataverse profiles that need to be transitioned to Synapse Link.
- Understand the nature and volume of data in the profile and identify the Dataverse tables needs to be exported.
- Ensure all the compliance requirements are met for using Azure Synapse Analytics workspace.
- Ensure all the prerequisites for setting up the Azure Synapse Link are met.
- Consider whether you need multiple export profiles with Azure Synapse Link to comply with security, compliance, and/or large data volumes.
- Determine a strategy on whether you need [monthly or yearly partition](#), and if your organization require [Append Only](#) feature for auditing.
- Start with a small volume of data to Pilot the transition and ensure all major scenarios working.

Implementation

Implementation

- Start transition with your lower (sandbox) environments before moving to production.
- Setup your Azure Synapse Analytics workspace and Azure Data Lake Storage Gen2.
- Configure the Azure Synapse Link in the Power Platform Maker Portal and add the required Dataverse tables for synchronization.
- Ensure all the tables have completed syncing and there no failures from Power Apps Maker portal.
- Ensure end-to-end scenarios are tested post transition to Azure Synapse Link.
- Consider having DES and Synapse Link to co-exists post successful transition for a brief "Stabilization" period.

Operation

Operation

- Monitor Azure Synapse Link from Power App Maker portal to ensure all data is transitioning to ADL as expected.
- Transfer any existing external monitoring on DES to Azure Synapse Link for Dataverse.
- Ensure all downstream business processes are working as expected.
- Plan for a regular communication to key users and stakeholders post transition to update key challenges / issues / achievements.
- Decide a stabilization date post transition to remove DES profiles from your Dataverse organizations and delete the SQL server instances.



Incremental vs. Full Transition from DES to Azure Synapse Link

- Pilot and address feedback
 - Pilot the Azure Synapse Link for Dataverse with a limited set of Dataverse tables and integrate the data to any downstream data warehousing or reporting tools.
 - Work on your transitional scenarios. Refer to Scenarios in the playbook for additional considerations.
 - Ensure end-to-end business scenarios are working as expected.
- Deployments of different sizes and complexity may require different approaches for transitioning data. A large organization with very high volumes of table sizes and requires complex transformations may choose an incremental roll out. Smaller organizations with smaller amounts of data may choose to sync all the data at once.

Incremental Transition for Large Databases Full Transition for Small Databases

Benefits

- Minimal impact on the data integration systems.
- Customers can begin implementing the end-to-end scenarios with low impact tables followed by high impact tables.

Challenges

- It will take longer time to replicate Dataverse data to ADL.
- Customers must maintain both DES and Azure Synapse Link for Dataverse during the transition.



Benefits

- One time data transition which requires creating the profile/s and configuring all the tables that needs to be replicated from Dataverse.
- All the integrations can be tested post data replication finishes.

Challenges

- “Big bang” approach, all issues, risks, gaps needs to be considered in advance.
- Large migrations can take multiple days to complete.



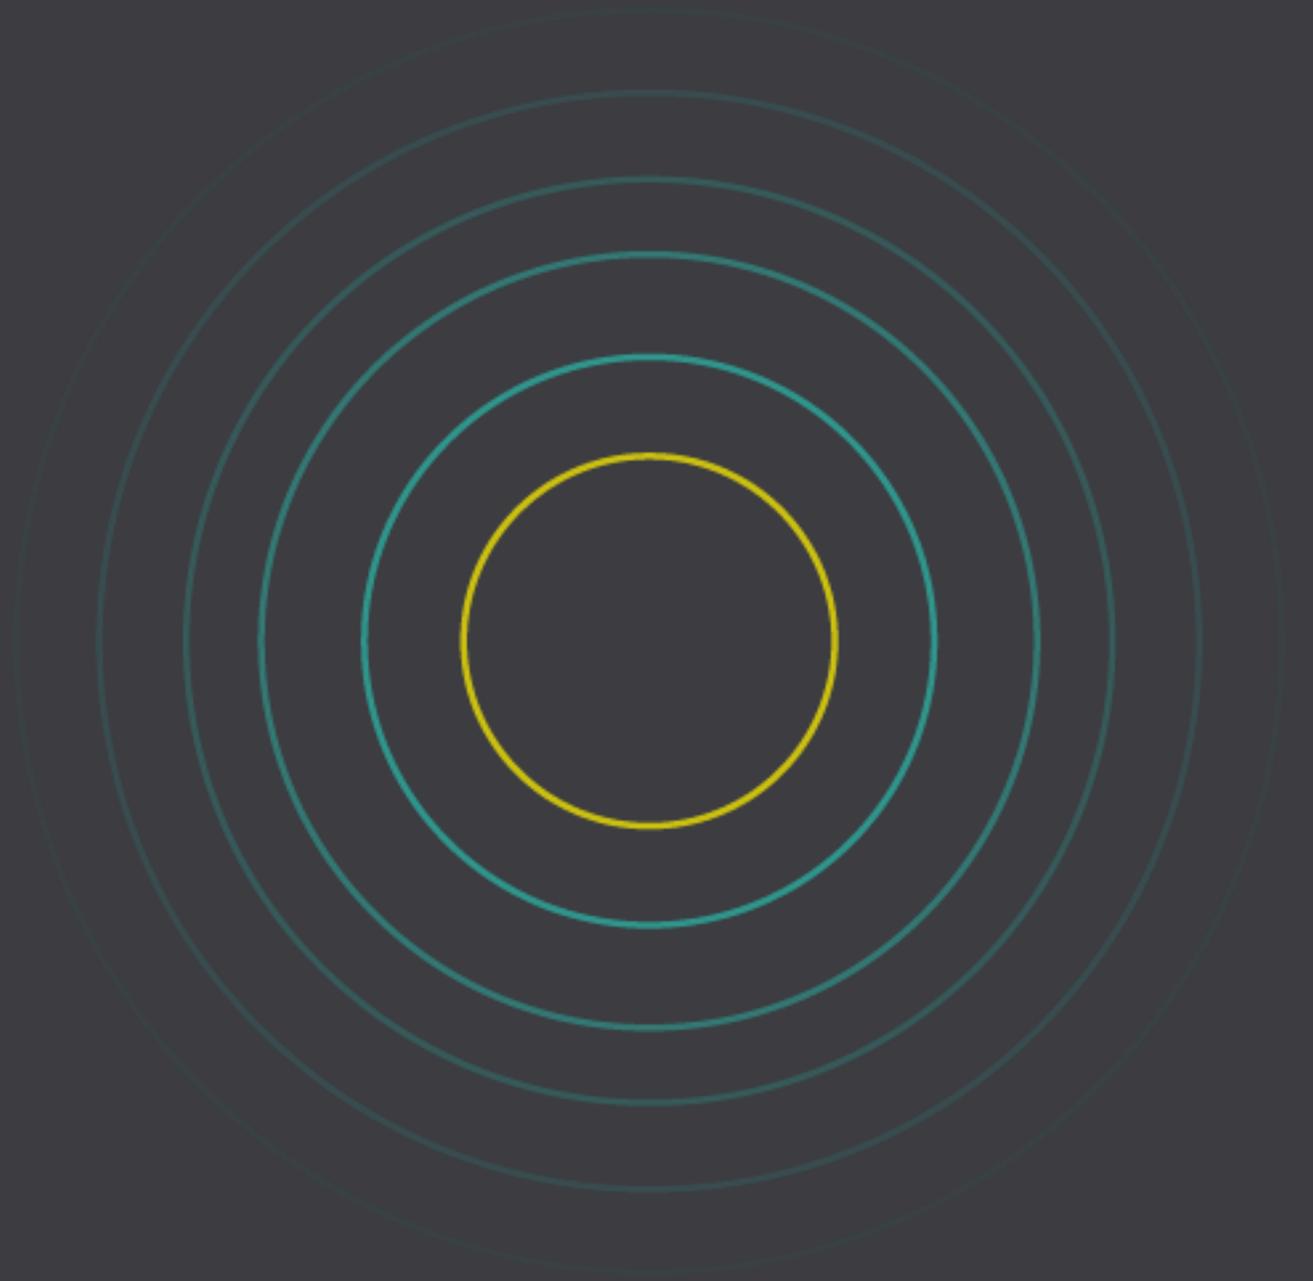
Raising and Tracking Issues in Production

When it's time to find answers to your Microsoft Dynamics 365 questions, there are a variety of self-support and assisted-support options to provide the help you need.

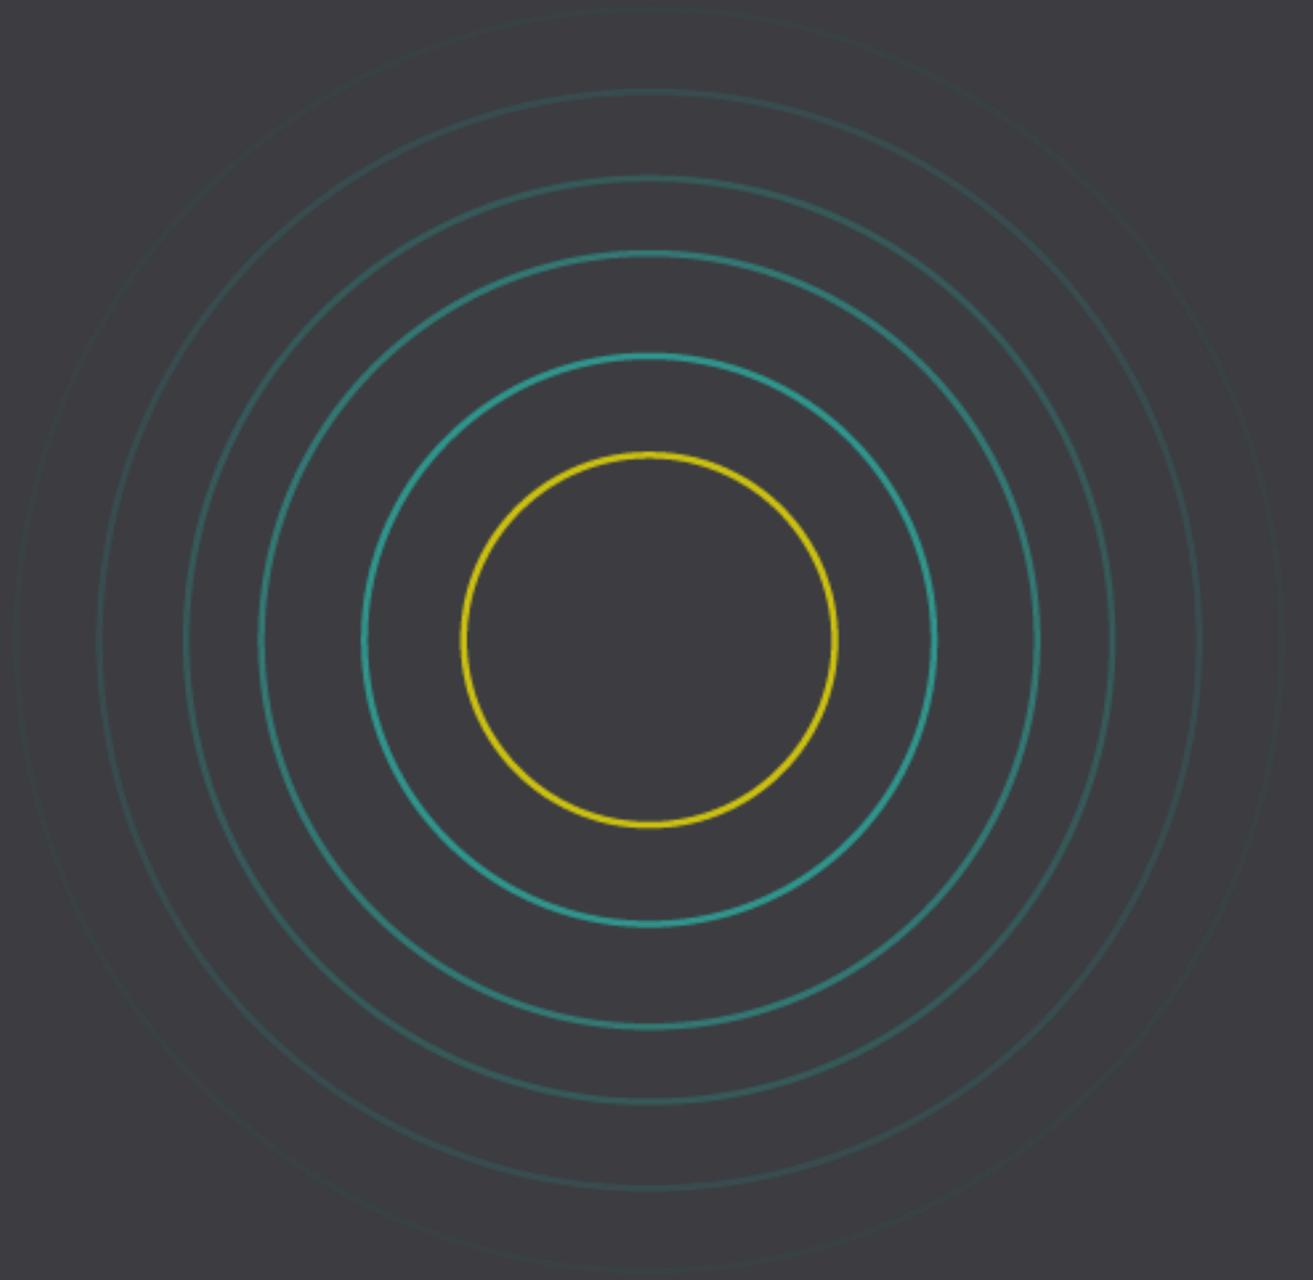
- **Can't find a solution yourself?** Send a question to the [Community](#) and receive answers from other customers, partners, MVPs, and Microsoft employees at <https://community.dynamics.com/>.
- **Raise support requests for expedited live site support.** Your support ticket will be tracked and routed to the appropriate team of experts that will help you resolve your problem(s). You can create the support ticket in [Power Platform Admin Center](#).



Appendix



Sample Scenarios



Sample Transition Scenarios - Considerations

- These sample DES scenarios are commonly used across our customers and partners. There can be additional scenarios and this section is not intended to cover all DES transition scenarios.
- The purpose of the sample scenarios is to provide guidance on how a DES scenario can be transitioned to Synapse Link.
- There may be many ways to implement the DES scenario with Synapse link.
- Every customer's scenario and use case will be different and the appropriate architecture that will best fit the scenario should be chosen accordingly.



Scenario 1 – Analyze and Report on Dataverse Data

DES Use Case

- Several organizational reports are developed using Dataverse data.
- Data Export Service is used to export Dataverse data to Azure SQL.
- Power BI reports are used to report on the data from Azure SQL.
- Power BI reports are shared across organization to different stakeholders.



Sample Transition

Approach

- Configure Azure Synapse Link with a Synapse Workspace. More information [here](#).
- Configure Power BI with Azure Synapse Analytics (DW) connector. More information [here](#).



Scenario 2 – Dataverse Data is Replicated

DES Use Case

- Data Export Service is used to export Dataverse data to Azure SQL where downstream systems read the data.



Sample Transition

Approach 1: Serverless SQL Pool

- Configure Azure Synapse Link with a Synapse Workspace. More information [here](#).
- Query your data with Serverless SQL. More information [here](#).

Approach 2: Azure SQL

- Configure Azure Synapse Link with Data Lake . More information [here](#).
- Use the official Pipeline Template to copy exported Data Lake data to Azure SQL. More information [here](#).
- Add a trigger for your pipeline.

Approach 3: Dedicated SQL Pool

- Configure Azure Synapse Link with a Synapse workspace. More information [here](#).
- Copy data to Dedicated SQL with Synapse Pipelines. More information [here](#).
- Add a trigger for your pipeline.



Scenario 3 – Transform Dataverse data

DES Use Case

- Data Export Service is used to export Dataverse data to Azure SQL.
- SQL Stored Procedures are configured to transform exported data. For example, a stored procedure can execute a complex calculation on the values of the exported Dataverse data.
- SQL views are configured to view specific parts of the exported data. For example, a view is created to combine data from multiple exported Dataverse tables.



Sample Transition

Approach 1: Serverless SQL

- Configure Azure Synapse Link with a Synapse Workspace. More information [here](#).
- Create a new Serverless SQL Database in the same workspace. More information [here](#).
- Run your create view, stored procedure, or write SQL query by referencing your Dataverse database using the three-part naming convention – *[database-name].[schema-name].[table-name]*.

Approach 2: Synapse Apache Spark

- Configure Azure Synapse Link with a Synapse Workspace. More information [here](#).
- Create a new Apache Spark notebook. More information [here](#).
- Run your create view, stored procedure, or write SQL query by referencing your Dataverse database using the three-part naming convention – *[database-name].[schema-name].[table-name]*.



Scenario 4 – View the Incremental Updates of Dataverse data

DES Use Case

- Data Export Service is used to export Dataverse data to Azure SQL.
- Analytics requires reading only the updated records or the incremental updates in a specified time interval.



Sample Transition

Approach

- Configure Azure Synapse Link with Synapse workspace. Ensure, Append Only option is selected from Advanced Configuration. More information [here](#).
- With the tool of your choice (like Serverless SQL), maintain a pointer to the *SinkModifiedOn* field and set the preferred time interval. More information [here](#).



Scenario 5 – View Dataverse data via External Table in Azure SQL DB

DES Use Case

- Data Export Service is used to export Dataverse data to Azure SQL.
- The data is transformed and/or merged with other data sources.
- Virtual Tables are used to surface the transformed data in Dataverse from Azure SQL.



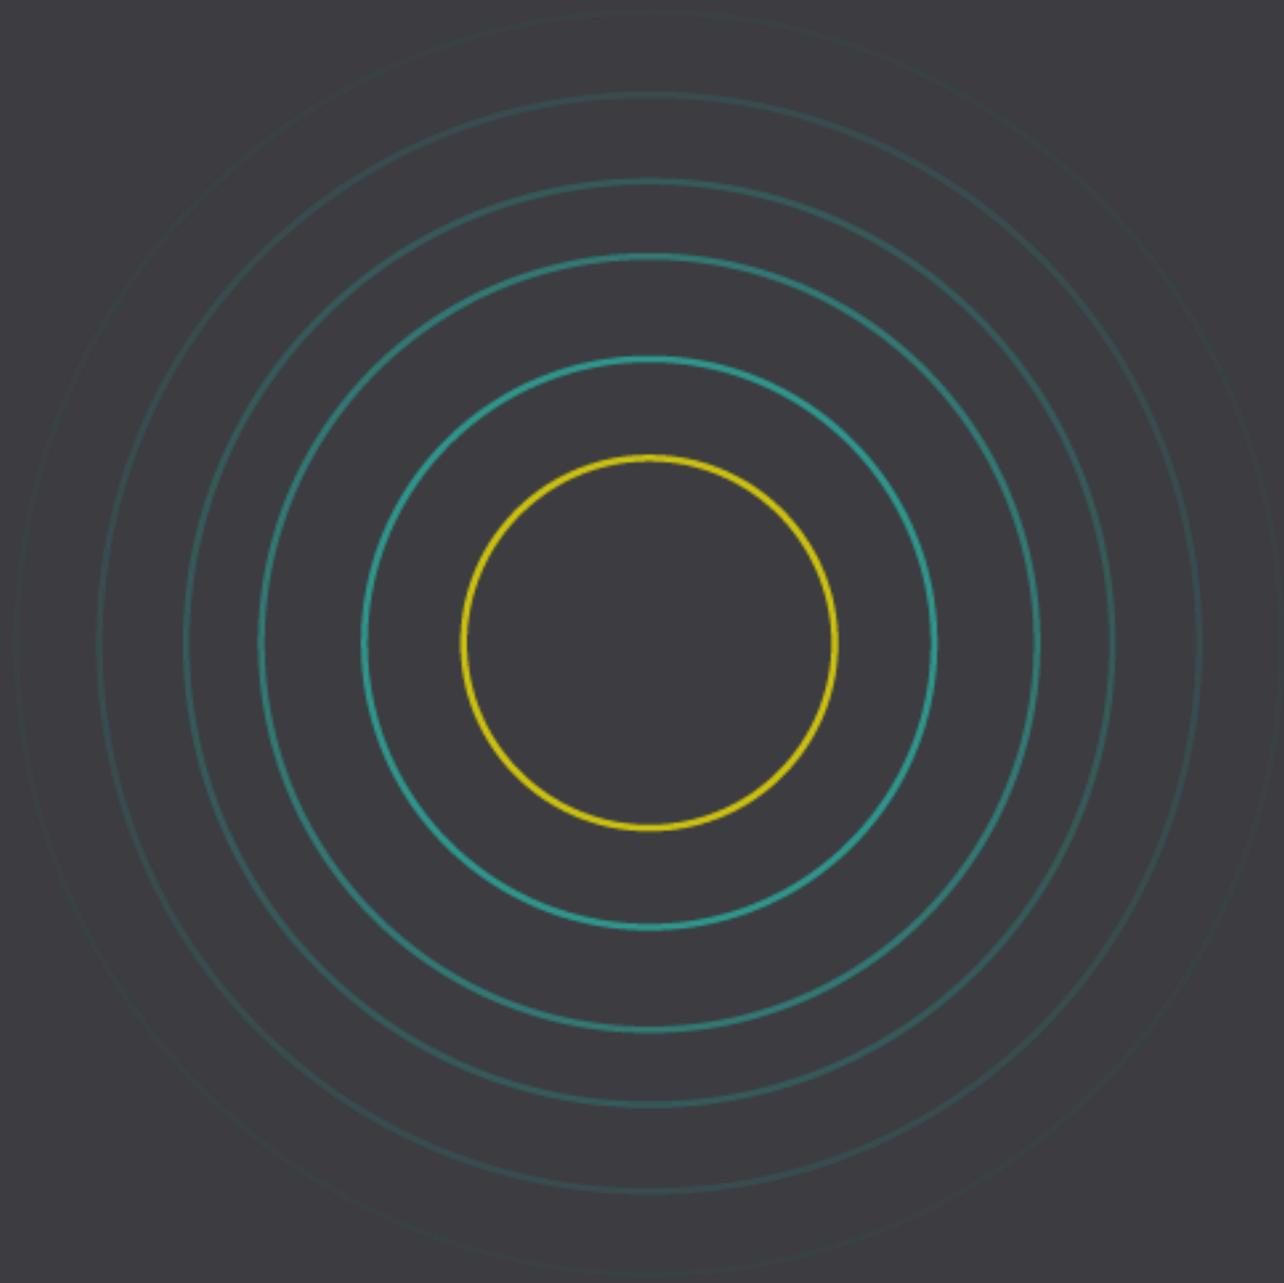
Sample Transition

Approach

- Configure Azure Synapse Link with a Synapse Workspace. More information [here](#).
- Add external tables pointing to Synapse SQL serverless endpoint.
- Read the Dataverse data via SQL query. More information refer to Appendix section.



Anti-patterns



Sample Antipatterns for DES & Azure Synapse Link for Dataverse

- The purpose of the sample antipatterns is to provide guidance on scenarios where Azure Synapse link for Dataverse may not be the right fit to transition from Data Export Service
- While DES was envisaged to support the Operational/ Enterprise Reporting workloads, there were scenarios where DES was used for what it is not meant for. Such scenarios are equally not applicable for Azure Synapse Link for Dataverse as well.



There can be additional antipatterns as well. This section is not intended to cover all the antipattern scenarios related to DES & Azure Synapse link for Dataverse.

Scenario 1 – Realtime or very near-time(~secs of latency) Integration

DES Use Case

- Data Export Service is used to replicate Dataverse data to Azure SQL
- Downstream systems are integrated with Azure SQL
- Customer business have a dependency **to real time data availability (within seconds)** on the integrating systems.
- Exporting data for near real time OLTP queries.

Facts :

- Data Export Service & Azure synapse link for Dataverse are designed for near real-time data replication.
- Azure synapse link for Dataverse has soft SLA of 15 mins
- The delta sync latency can vary depending on volume of data being synched, number of profiles, tables, and #lookups associated to the tables etc.



Sample Transition

Recommended Approach

- Evaluate [TDS end point](#)
- Consider evaluating [publisher-subscriber](#) model to get the data out of Dataverse using any of the [azure integration](#) capabilities with Dataverse.
- You may also evaluate Azure platform capabilities like [Logic Apps](#) to get near-real-time update from Dataverse.
- Another alternate approach might be [Power Platform Business Events](#).

Considerations

- [Service protection API limit](#) is applicable when interacting with Dataverse API.
- If using Power Automate, check the requests limit [here](#).
- Similarly, any connector used with Power Automate / Logic App will be governed by the connector limit.

Scenario 2 – Integration requiring transactional Concurrency

DES Use Case

- Data Export Service is used to replicate Dataverse data to Azure SQL
- Downstream systems are integrated with Azure SQL
- Incremental data updated should be received in Azure SQL in same order as it was updated in Dataverse

Facts :

- Both Data Export Service and Azure Synapse Link services are built for eventual data consistencies. There is no guarantee that data edits are replicated in the order of user actions on Dataverse



Sample Transition

Recommended Approach

- Consider evaluating [publisher-subscriber](#) model to get the data out of Dataverse using any of the [azure integration](#) capabilities with Dataverse.
- You can also leverage [change tracking](#) to get the deltas from a Dataverse table.
- Dataverse recently introduced [Business Events](#) which can be considered for another additional option.

Scenario 3 – Dataverse data is required for real time reporting

DES Use Case

- Data Export Service is used to replicate Dataverse data to Azure SQL
- Reporting tool like Power BI is used to interface business critical near real time reporting from Azure SQL Data.

Facts :

- Both Data Export Service and Azure Synapse Link services are built for near real time data.

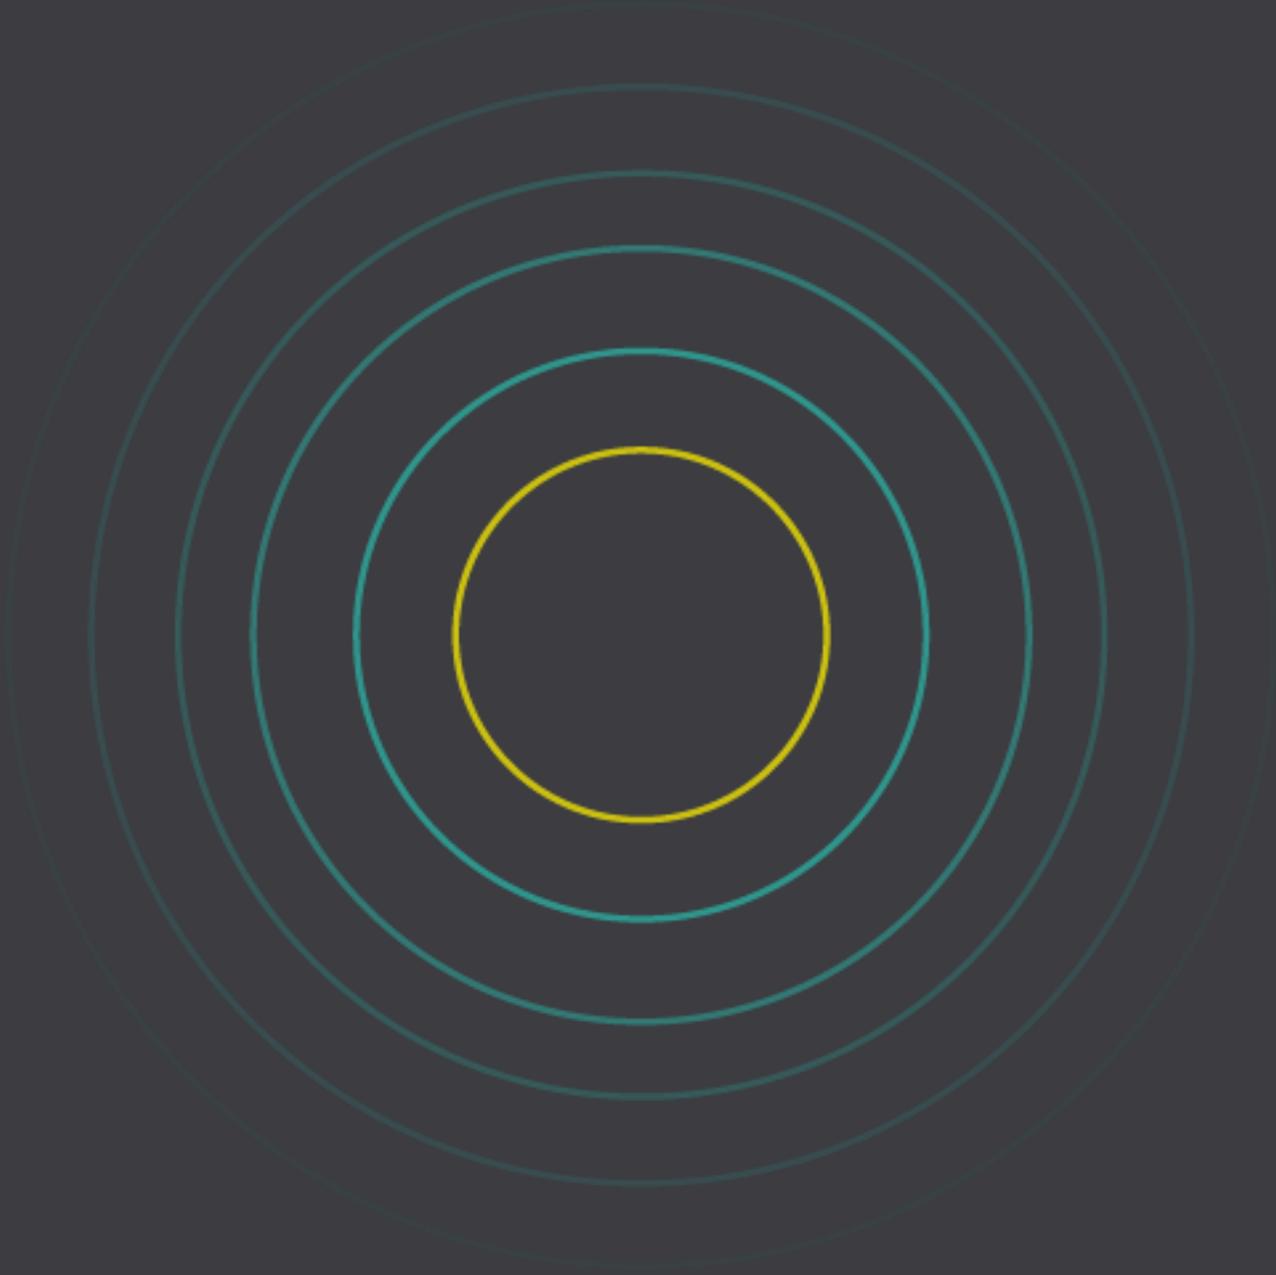


Sample Transition

Recommended Approach

- Consider [Power BI Direct Query](#) support for real time reports with Dataverse data.
- Please be aware of the known [limitations](#) while using this approach.

FAQ



💡 SLA for replicating Dataverse data to Azure Synapse

Near real time with soft SLA of 15 minutes. Following are the dependencies which might impact:

- Initial data sync and delta data sync volumes.
- Number of Dataverse tables added in profile
- Number of profiles
- Number of columns in a Dataverse tables
- Network latency

💡 Dataverse Table Schema Changed

New Column Added

- For New Rows, the columns will be added at the end of the latest partition
- For existing rows, no change till updated.

Column Deleted

- Column name will exist with a Null value in latest partition
- Existing data preserved;

Column data type changed

- Need relink of Azure Synapse Link for Dataverse
- Not recommended



💡 Row deleted in Dataverse

In place update

- Row will be deleted

Append Only

- A new row will be added to the end of the partition file.
- "isDeleted" column set to "True" for deleted row

💡 Incremental data updates from Dataverse

Configure write behavior as Append Only to get incremental updates from any table

Manage tables for ~~XXXXXXXXXX~~

Manage tables
2 of 364 selected

When you add new table(s), all changes for existing tables are temporarily paused. We will resume writing changes for existing table(s) after we complete exporting data for new table(s) to the Azure data lake.

Advanced

Show advanced configuration settings

Filter by keyword

<input type="radio"/>	Table 1	Name	Append only	Partition
<input checked="" type="radio"/>	Account	account	<input type="checkbox"/>	Month
<input checked="" type="radio"/>	Accountleads	accountleads	<input checked="" type="checkbox"/>	Year
<input type="radio"/>	Action Card	actioncard	<input type="checkbox"/>	Month

Recommended write behavior configuration for Dataverse table having very frequent data updates

- For In-place update, row is scanned through all destination files and then insert or update is performed.
- For Append Only, row is added at the end of latest partition for insert or update operation in Dataverse.
- **Append Only is recommended** to gain better performance.

Credentials used for ongoing data sync process

- Export to Data lake service account gets created during first Link configuration
- Export to Data lake service account is used for data Synchronization

Data type mismatch for Create On and GUID columns in CSV after data synchronized with Data Lake

- Format of GUID column will be Varchar(36)

Column Name	Format	Example
SinkCreatedOn and SinkModifiedOn	M/d/yyyy H:mm:ss tt	6/28/2021 4:34:35 PM
CreatedOn	yyyy-MM-dd'T'HH:mm:ss.ssssssXXX	2018-05-25T16:21:09.0000000+00:00
All Other Columns	yyyy-MM-dd'T'HH:mm:ss'Z'	2021-06-25T16:21:12Z



Permissions required to modify link Configuration

- The owner role access of Azure Data Lake Storage Gen2 account is **only** required when creating a Synapse Link for Dataverse.
- Owner permissions are not required to modify the link configuration (use **Manage tables**) after the initial sync.

403 error received while configuring new link



- Error might be because prerequisites are not configured correctly. Refer below URL for details

[Create an Azure Synapse Link for Dataverse with your Azure Synapse Workspace - Power Apps | Microsoft Docs](#)

- If issue is not solved then consider logging service request with Microsoft.

💡 Considerations to read data from Trickle feed source vs snapshot from Azure Data Lake

Trickle Feed:

- Near real time
- Read/write issue might be observed at scale

Snapshot:

- Snapshot is taken every 1 hour
- No Read/write issue.

💡 Reference Dataverse read-only table in Azure Synapse Workspace

- Can be referenced using three-part naming convention: [database-name].[schema-name].[table-name].

💡 Considerations to query data across multiple databases within Azure Synapse workspace

- Data can be queried across same database types and in same region within Azure Synapse workspace.
- Data can be queried across databases of types either dedicated SQL or Serverless SQL pool.

💡 Restrict Network for Customer owned lake for Azure Synapse Link for Dataverse

- This is planned and expected to be available later this year, in 2022.

💡 Read Choice labels directly

- Refer below link for details:

<https://docs.microsoft.com/en-us/power-apps/maker/data-platform/azure-synapse-link-choice-labels>

💡 Read data from Azure Synapse

- Query Azure Synapse Link for Dataverse data with serverless SQL

<https://docs.microsoft.com/en-us/powerapps/maker/data-platform/azure-synapse-link-serverless>

- Transform Azure Synapse Link for Dataverse data with Apache Spark

<https://docs.microsoft.com/en-us/powerapps/maker/data-platform/azure-synapse-link-spark>

- Visualize Azure Synapse Link for Dataverse data with Power BI:

<https://docs.microsoft.com/en-us/powerapps/maker/data-platform/azure-synapse-link-powerbi>

- Copy exported Dataverse data to dedicated SQL pool:

<https://docs.microsoft.com/en-us/powerapps/maker/data-platform/azure-synapse-link-dedicated>

- For ETL, Azure Synapse Pipelines:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/get-started-pipelines>



💡 License implication of using Azure Synapse Link for Dataverse

- No additional cost for using Azure Synapse Link for Dataverse.
- Dataverse API entitlement limits will be applicable

[Requests limits and allocations - Power Platform | Microsoft Docs](#)

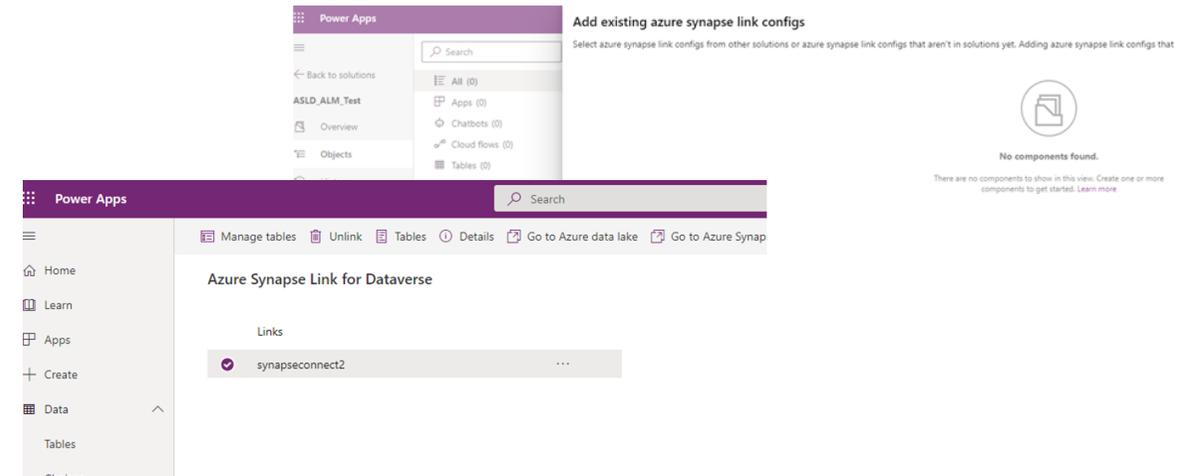
- Azure cost like ADLS storage and Azure Synapse pay per usage will be applicable.

[Pricing Overview—How Azure Pricing Works | Microsoft Azure](#)

💡 Dataverse APIs are consumed during data replication

- Dataverse APIs are called to pull latest changes from Dataverse
- Data manipulation events are aggregated & lesser API request are sent.

💡 Unable to find Azure Synapse link configuration while trying to add it in the solution



- Please install ALM solution in AppSource before creating the Synapse Link
- [Transport Azure Synapse Link for Dataverse configuration - Power Apps | Microsoft Docs](#)



💡 Alert mechanism to notify admins about failure

- Failures are shown in Azure Synapse link for Dataverse page.

Azure Synapse Link for Dataverse > **Account**

Tables Details Discover hub

Table 1	Name	Sync status	Last synchronized on	Count	Append only	Partition
Account	account	Active	03/07/2022 8:46:12 AM	1	No	Month
Contact	contact	Active	03/07/2022 2:28:06 PM	3	No	Month
Lead	lead	Active	03/07/2022 2:27:06 PM	-	No	Month

💡 Considerations for querying data from Serverless SQL

- Max query duration 30min. (Timeout occurs)
- ~10TB max data that can be processed per query

Note: Use desktop tools (ADS, SSMS) instead of Synapse Studio if you are returning multi-GB of data per query. Web interface is not designed for huge data exports.

💡 How many Azure Synapse workspace can be configured with a Dataverse instance

- Each profile will need a new Azure Synapse Workspace so only one Azure Synapse Workspace can be configured with a Dataverse instance.

💡 Create statistics for query performance in Synapse Serverless SQL

- Statistics are not created automatically and need to create is manually.
- Refer below link for more details:

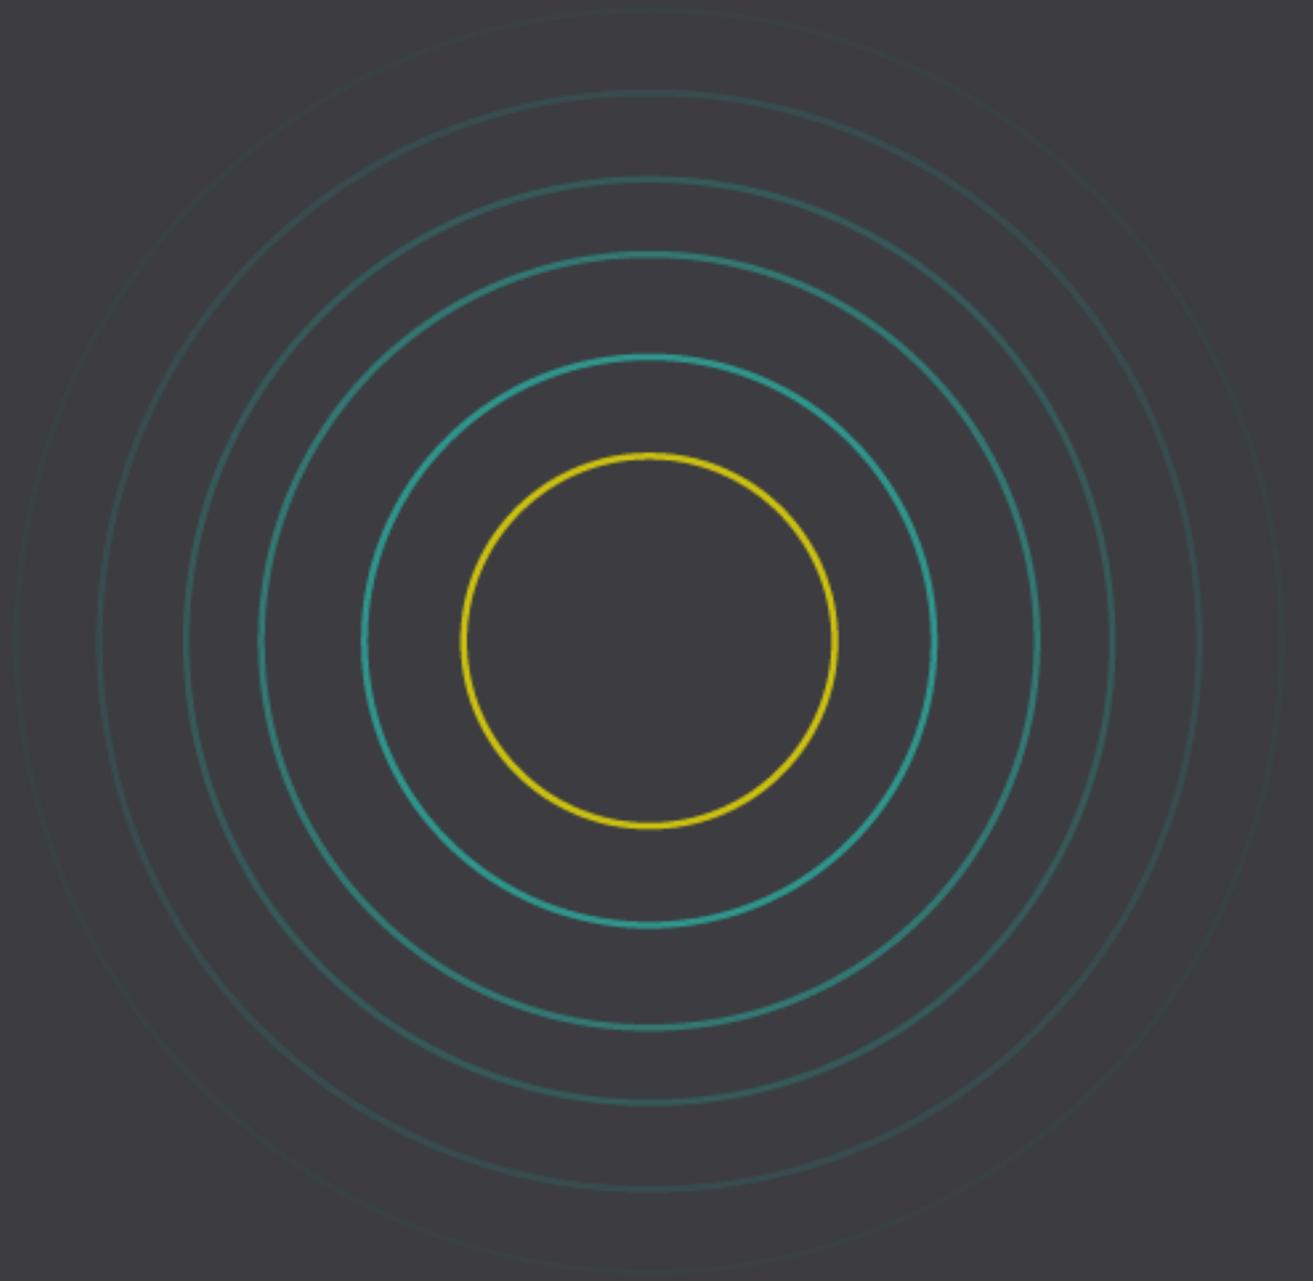
[Create and update statistics using Azure Synapse SQL resources - Azure Synapse Analytics | Microsoft Docs](#)

💡 I have implemented real-time using DES. How to replicate the same using Azure Synapse link for Dataverse?

- Evaluate [TDS end point](#)
- Consider evaluating [publisher-subscriber model](#) to get the data out of Dataverse using any of the [azure integration](#) capabilities with Dataverse.
- You can also leverage [change tracking](#) to get the deltas from a Dataverse table.
- Both DES & Azure synapse link for Dataverse are designed for near real time data export and it is not recommended for real-time integration.



Samples



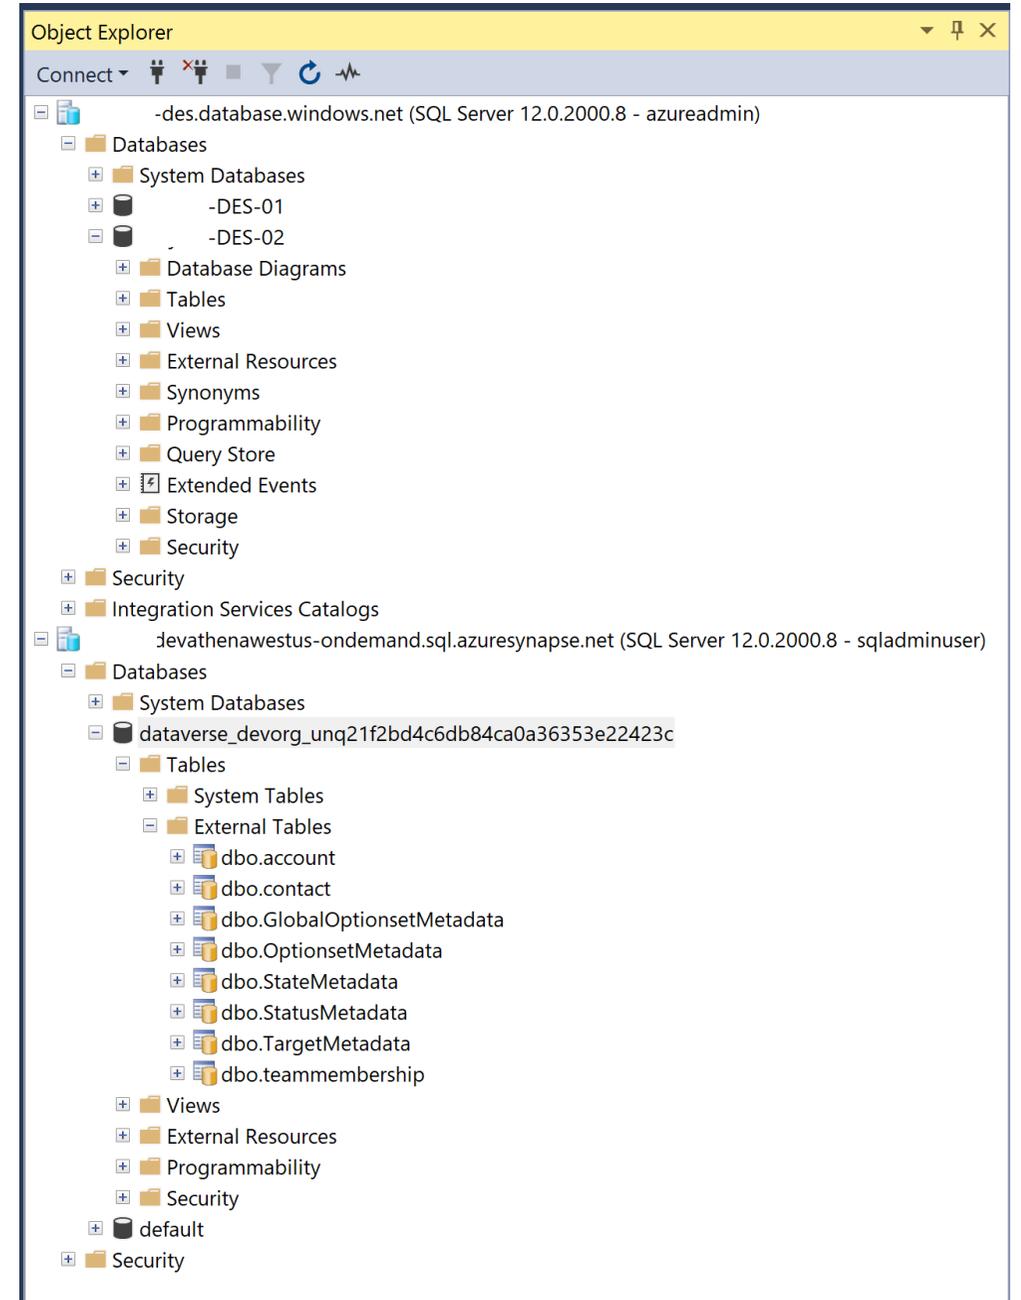


Pre-requisites:

- Azure SQL DB Administrator role
- Azure Synapse Administrator role
- Master key in the Azure SQL DB

Connect to Azure SQL database and Azure Synapse SQL on demand using SSMS. Please note that connecting via Azure portal query editor won't suffice. Existing Azure SQL DB cannot be used, if external table with same name as existing table needs to be created, due to name conflicts.

External table is created with static table schema. For any schema changes, external table can be safely deleted & re-created (ALTER statement is not allowed)





In Azure SQL database

Create a Database Scope Credential i.e., with the following SQL statement. The user details in this step should match the sqladmin credentials provided while deploying synapse workspace. If password is not known, please reset it from Azure portal.

```
CREATE DATABASE SCOPED CREDENTIAL SynapseSqlAdminUserCredential
    WITH IDENTITY = 'sqladminuser', SECRET = '***';
GO
```

Create an external data source as below. Location should point to the Serverless SQL end Point and Database name is the lake database name in Synapse

```
CREATE EXTERNAL DATA SOURCE SynapseSqlDataStore
WITH (
    TYPE = RDBMS,
    LOCATION='devathenawestus-ondemand.sql.azure.synapse.net',
    DATABASE_NAME='dataverse_devorg_unq21f2bd4c6db84ca0a36353e22423c',
    CREDENTIAL = SynapseSqlAdminUserCredential
);
GO
```



Please refer <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-external-table-transact-sql> for additional arguments.



In Synapse database

Generate create table script from the Synapse tables. Right click on pertaining table and select "Script Table as" -> Create To -> New Query Editor.

```

SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER OFF
GO

CREATE EXTERNAL TABLE [dbo].[account]
(
    [Id] [varchar](36),
    [SinkCreatedOn] [datetime2](7),
    [SinkModifiedOn] [datetime2](7),
    [statecode] [bigint],
    [statuscode] [bigint],
    [preferredcontactmethodcode] [bigint],
    [customersizecode] [bigint],
    [businessstypecode] [bigint],
    [industrycode] [bigint],
    [shippingmethodcode] [bigint],
    [preferredappointmenttimecode] [bigint],
    [address2_shippingmethodcode] [bigint],
    [ownershipcode] [bigint],
    [paymenttermscode] [bigint],
    [accountclassificationcode] [bigint],
    [territorycode] [bigint],
    [preferredappointmentdaycode] [bigint],
    [address1_shippingmethodcode] [bigint],
    [accountratingcode] [bigint],
    [address2_freighttermscode] [bigint],
    [address2_addresstypecode] [bigint],
    [address1_freighttermscode] [bigint],
    [customertypecode] [bigint],
    [accountcategorycode] [bigint],
    [address1_addresstypecode] [bigint],
    [merged] [varchar](6),
    [participatesinworkflow] [varchar](6),
    [donotpostalmail] [varchar](6),
    [donotbulkpostalmail] [varchar](6),
    [donotphone] [varchar](6),
    [donotemail] [varchar](6),
    [donotbulkemail] [varchar](6),
    [isprivate] [varchar](6),

```

```

[createdbyyominame] [varchar](640),
[address2_telephone1] [varchar](200),
[address2_telephone2] [varchar](200),
[address2_telephone3] [varchar](200),
[address1_line2] [varchar](1000),
[msdyn_segmentidname] [varchar](400),
[primarycontactidyominame] [varchar](640),
[createdbyname] [varchar](640),
[tickersymbol] [varchar](40),
[address2_primarycontactname] [varchar](600),
[importsequencenumber] [bigint],
[modifiedonbehalfbyyominame] [varchar](640),
[address1_fax] [varchar](200),
[address1_postalcode] [varchar](80),
[utconversiontimezonecode] [bigint],
[slainvokedidname] [varchar](640),
[preferredsystemuseridyominame] [varchar](640),
[telephone3] [varchar](200),
[address2_county] [varchar](200),
[IsDelete] [varchar](6)
)
WITH (DATA_SOURCE = [datasource_63ee4a45-1f02-416c-b7e5-3af9b5aba952], LOCATION = N'*.csv', FILE_FORMAT = [delimitedtext63ee4a45-1f02-416c-b7e5-3af9b5aba952])
GO

EXEC sys.sp_addextendedproperty @name=N'EDCTable', @value=N'True' , @level0type=N'SCHEMA',@level0name=N'dbo', @level1type=N'TABLE',@level1name=N'account'
GO

EXEC sys.sp_addextendedproperty @name=N'global_object_id', @value=N'63ee4a45-1f02-416c-b7e5-3af9b5aba952' , @level0type=N'SCHEMA',@level0name=N'dbo', @level1type=N'TABLE',@level1name=N'account'
GO

EXEC sys.sp_addextendedproperty @name=N'global_object_version', @value=N'1' , @level0type=N'SCHEMA',@level0name=N'dbo', @level1type=N'TABLE',@level1name=N'account'
GO

EXEC sys.sp_addextendedproperty @name=N'origin_type', @value=N'SPARK' , @level0type=N'SCHEMA',@level0name=N'dbo', @level1type=N'TABLE',@level1name=N'account'
GO

EXEC sys.sp_addextendedproperty @name=N'polaris_table_type', @value=N'csv' , @level0type=N'SCHEMA',@level0name=N'dbo', @level1type=N'TABLE',@level1name=N'account'
GO

```



In Azure SQL database

Viewing data from snapshots (preferred option)

Create a new query window connected to Azure SQL DB. Copy the create statement from generate create table step and update the table name to match the partition view (for example account_partitioned) and use WITH Data Source option (i.e., the external data source created in earlier step) at the end to create external table in Azure SQL DB.

```
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER OFF
GO

CREATE EXTERNAL TABLE [dbo].[account_partitioned]
(
    [Id] [varchar](36),
    [SinkCreatedOn] [datetime2](7),
    [SinkModifiedOn] [datetime2](7),
    [statecode] [bigint],
    [statuscode] [bigint],
    [preferredcontactmethodcode] [bigint],
    [address1_line1] [varchar](1000),
    [address1_postofficebox] [varchar](80),
    [lastusedincampaign] [datetime2](7),
    [createdbyyominame] [varchar](640),
    [address2_telephone1] [varchar](200),
    [address2_telephone2] [varchar](200),
    [address2_telephone3] [varchar](200),
    [importsequencenumber] [bigint],
    [modifiedonbehalfbyyominame] [varchar](640),
    [address1_fax] [varchar](200),
    [address1_postalcode] [varchar](80),
    [utcconversiontimezonecode] [bigint],
    [slainvokedidname] [varchar](640),
    [preferredsystemuseridyominame] [varchar](640),
    [telephone3] [varchar](200),
    [address2_county] [varchar](200),
    [IsDelete] [varchar](6)
)
WITH (DATA_SOURCE = SynapseSqlDataStore)
GO
```

Please note that columns have removed from above screenshot for brevity.



For snapshot feature, please refer <https://docs.microsoft.com/en-us/power-apps/maker/data-platform/azure-synapse-link-synapse#access-near-real-time-data-and-read-only-snapshot-data-preview>.



In Azure SQL database



Viewing data from live table

Create a new query window connected to Azure SQL DB. Copy the create statement from generate create table step and update it to use WITH Data Source option (i.e., the external data source created in earlier step) at the end to create external table in Azure SQL DB.

```
SET ANSI_NULLS ON
GO

SET QUOTED_IDENTIFIER OFF
GO

CREATE EXTERNAL TABLE [dbo].[account]
(
    [Id] [varchar](36),
    [SinkCreatedOn] [datetime2](7),
    [SinkModifiedOn] [datetime2](7),
    [address2_postofficebox] [varchar](80),
    [address1_utcoffset] [bigint],
    [address1_city] [varchar](320),
    [owneridyominame] [varchar](640),
    [modifiedbyexternalpartyname] [varchar](640),
    [onholdtime] [bigint],
    [modifiedonbehalfbyyominame] [varchar](640),
    [address1_fax] [varchar](200),
    [address1_postalcode] [varchar](80),
    [utconversiontimezonecode] [bigint],
    [slainvokedidname] [varchar](640),
    [preferreddsystemuseridyominame] [varchar](640),
    [telephone3] [varchar](200),
    [address2_county] [varchar](200),
    [IsDelete] [varchar](6)
)
WITH (DATA_SOURCE = SynapseSqlDataStore)
GO
```

Please note that columns have removed from above screenshot for brevity.



In Azure SQL database

Query Dataverse data from the Azure SQL database using SQL statements.

The screenshot displays the SQL Server Enterprise Manager interface. On the left, the Object Explorer shows the server structure, including the 'dbo.account_partitioned' external table. The central pane shows the following SQL query:

```
SELECT *
FROM [dbo].[account_partitioned]
GO
```

The bottom pane shows the results of the query, which is a table with 21 rows and 10 columns. The columns are: Id, SinkCreatedOn, SinkModifiedOn, statecode, statuscode, preferredcontactmethodcode, customersizecode, businesstypecode, and an unlabeled column with values 'N'. The data represents a list of account partitions with their respective sink creation and modification dates and codes.

Id	SinkCreatedOn	SinkModifiedOn	statecode	statuscode	preferredcontactmethodcode	customersizecode	businesstypecode	
1	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
2	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
3	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
4	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
5	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
6	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	6
7	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	1
8	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	4
9	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	5
10	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	6
11	2022-06-22 23:26:55.0000000	2022-06-22 23:26:55.0000000	0	1	1	1	1	4
12	2022-06-22 23:48:13.0000000	2022-06-22 23:48:13.0000000	0	1	1	1	1	N
13	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
14	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
15	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	8
16	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
17	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
18	2022-06-22 05:30:52.0000000	2022-06-22 05:30:52.0000000	0	1	1	1	1	N
19	2022-06-22 23:31:51.0000000	2022-06-22 23:31:51.0000000	0	1	1	1	1	8
20	2022-06-22 23:48:13.0000000	2022-06-22 23:48:13.0000000	0	1	1	1	1	N
21	2022-06-22 23:48:36.0000000	2022-06-22 23:48:36.0000000	0	1	1	1	1	N

The status bar at the bottom indicates: "Query executed successfully. sanjeek-des.database.window... azureadmin (74) sanjeek-DES-02 00:00:02 21 rows".



In Azure SQL database



Viewing latest data from snapshot



For Append Only

Create a view which uses the versionnumber to get the latest rows from the snapshot data. Below is a sample query. Below query include deletes. If deletes need to be removed, consider adding where clause (example WHERE [IsDelete] IS NULL) to remove deletes. Please note removing deleted won't remove deleted rows from the result set.

```
CREATE VIEW account_partitioned_latest
AS
SELECT *
FROM (
    SELECT
        *,
        rn = row_number() OVER (PARTITION BY [Id] ORDER BY [versionnumber] DESC)
    from [dbo].[account_partitioned]
) AS AP
WHERE rn = 1
```



In Azure SQL database

Viewing delete log from snapshot

For Append Only

Create a view which uses the uses IsDelete to get deleted rows. Below is a sample query to pull delete log from account & contact snapshots

```
CREATE VIEW [dbo].[DeleteLogSnapshot]
AS
SELECT CONVERT(nvarchar(64), 'account') AS [EntityName],
        CONVERT(nvarchar(64), [Id]) AS [RecordId],
        CONVERT(datetime, [SinkModifiedOn]) AS [SinkDeleteTime],
        [VersionNumber]
FROM [dbo].[account_partitioned]
WHERE IsDelete = 'True'
UNION
SELECT CONVERT(nvarchar(64), 'contact') AS [EntityName],
        CONVERT(nvarchar(64), [Id]) AS [RecordId],
        CONVERT(datetime, [SinkModifiedOn]) AS [SinkDeleteTime],
        [VersionNumber]
FROM [dbo].[contact_partitioned]
WHERE IsDelete = 'True'
```



In Azure SQL database

View data in different table name/schema than Synapse serverless

Below is a sample query to create table in different name/schema than Synapse serverless

```
CREATE EXTERNAL TABLE [dbo].[SynapseOptionsetMetadata]
(
  [EntityName] [varchar](128),
  [OptionSetName] [varchar](128),
  [Option] [bigint],
  [IsUserLocalizedLabel] [varchar](6),
  [LocalizedLabelLanguageCode] [bigint],
  [LocalizedLabel] [varchar](700)
)
WITH (
  DATA_SOURCE = SynapseSqlDataStore,
  SCHEMA_NAME = 'dbo',
  OBJECT_NAME = 'OptionsetMetadata',
)
GO
```



In Synapse SQL database

If there is permission issue with connecting with sqladminuser in SSMS with error message: Cannot find the CREDENTIAL 'filepath', because it does not exist, or you do not have permission, please create credentials using SAS token from container that is created by Azure Synapse Link for Dataverse

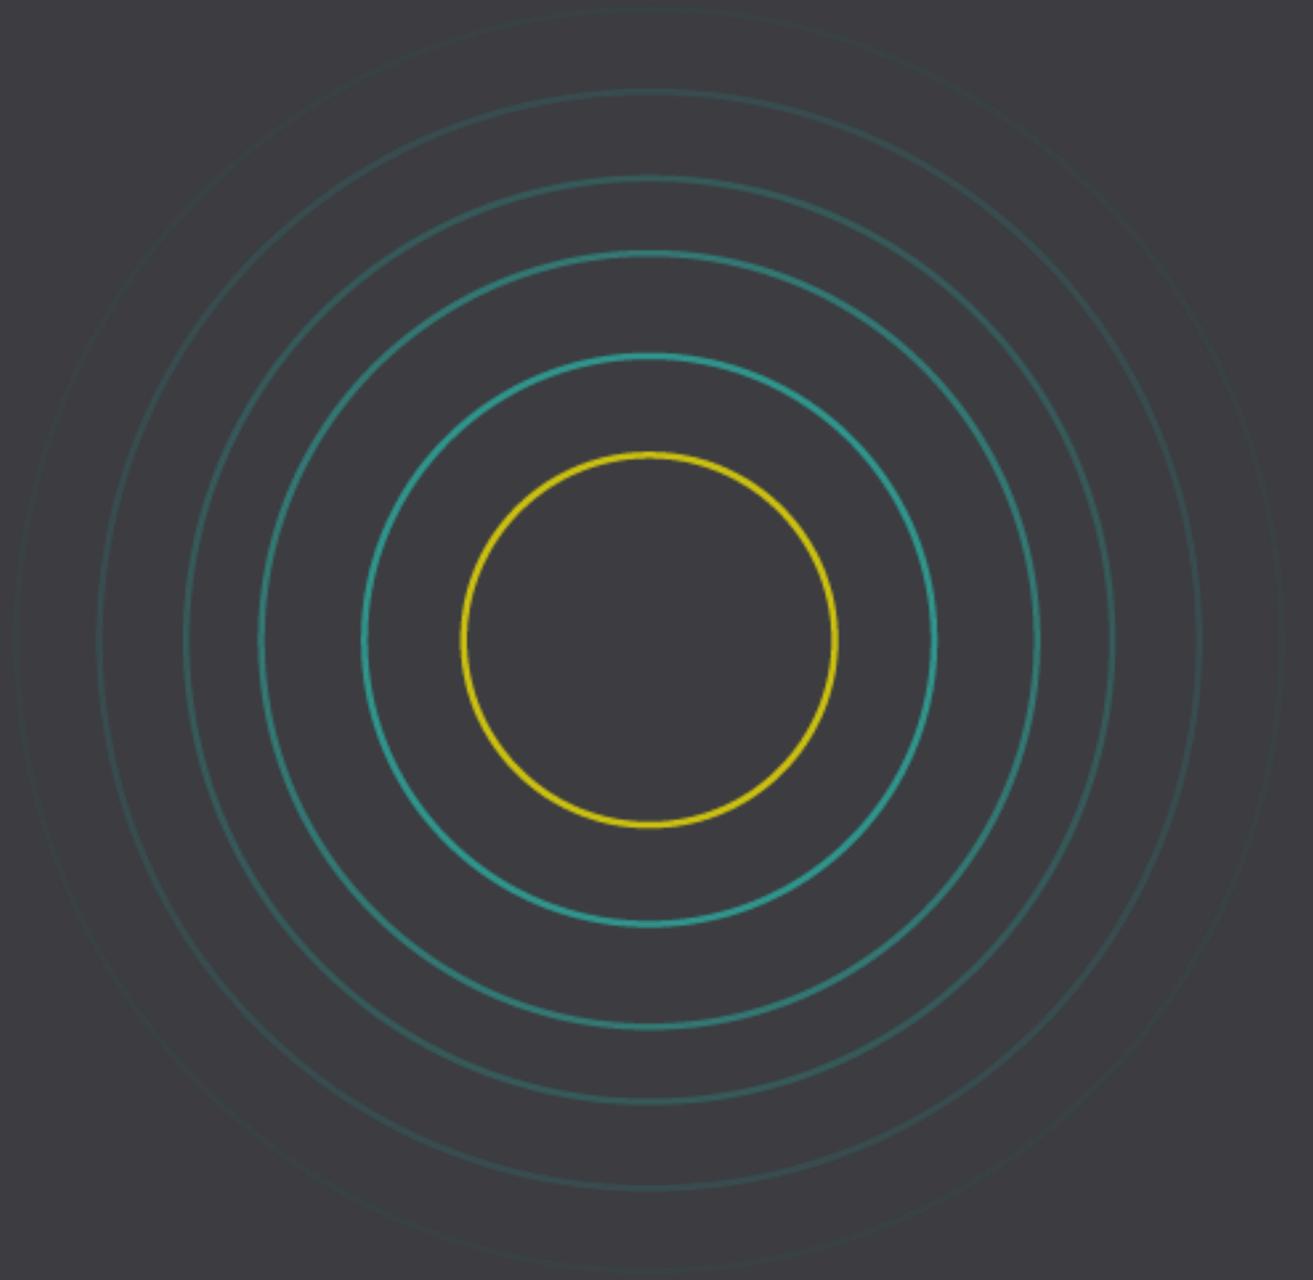
<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-scoped-credential-transact-sql?view=sql-server-ver16#b-creating-a-database-scoped-credential-for-a-shared-access-signature>

```
CREATE CREDENTIAL [https://devathenauswest.dfs.core.windows.net/dataverse-devorg-  
unq21f2bd4c6db84ca0a36353e22423c]  
WITH IDENTITY='SHARED ACCESS SIGNATURE' , SECRET = 'SAS Token FROM CONTAINER';  
GO
```



Please consider appropriate expiry of SAS token

Case Study



A real customer transitioning from DES to Synapse Link.

Background

- A major customer has deployed Dynamics 365 (Customer Engagement and Finance Apps) to run their Dealer Operations across Sales and Service.
- The customer exported Dynamics 365 data to be used by their Downstream systems for Analytics and Upstream systems for data visibility.

Current Stage

- The customer had chosen Data Export Service (DES) for exporting Dynamics 365 CE data to Azure SQL DB.
- Exported data was then transformed and pushed to On-Premises Data Warehouse for reporting, analysis, and other systems consumption.
- The customer wished to push the data to Azure hyperscale for their upstream systems consumption.

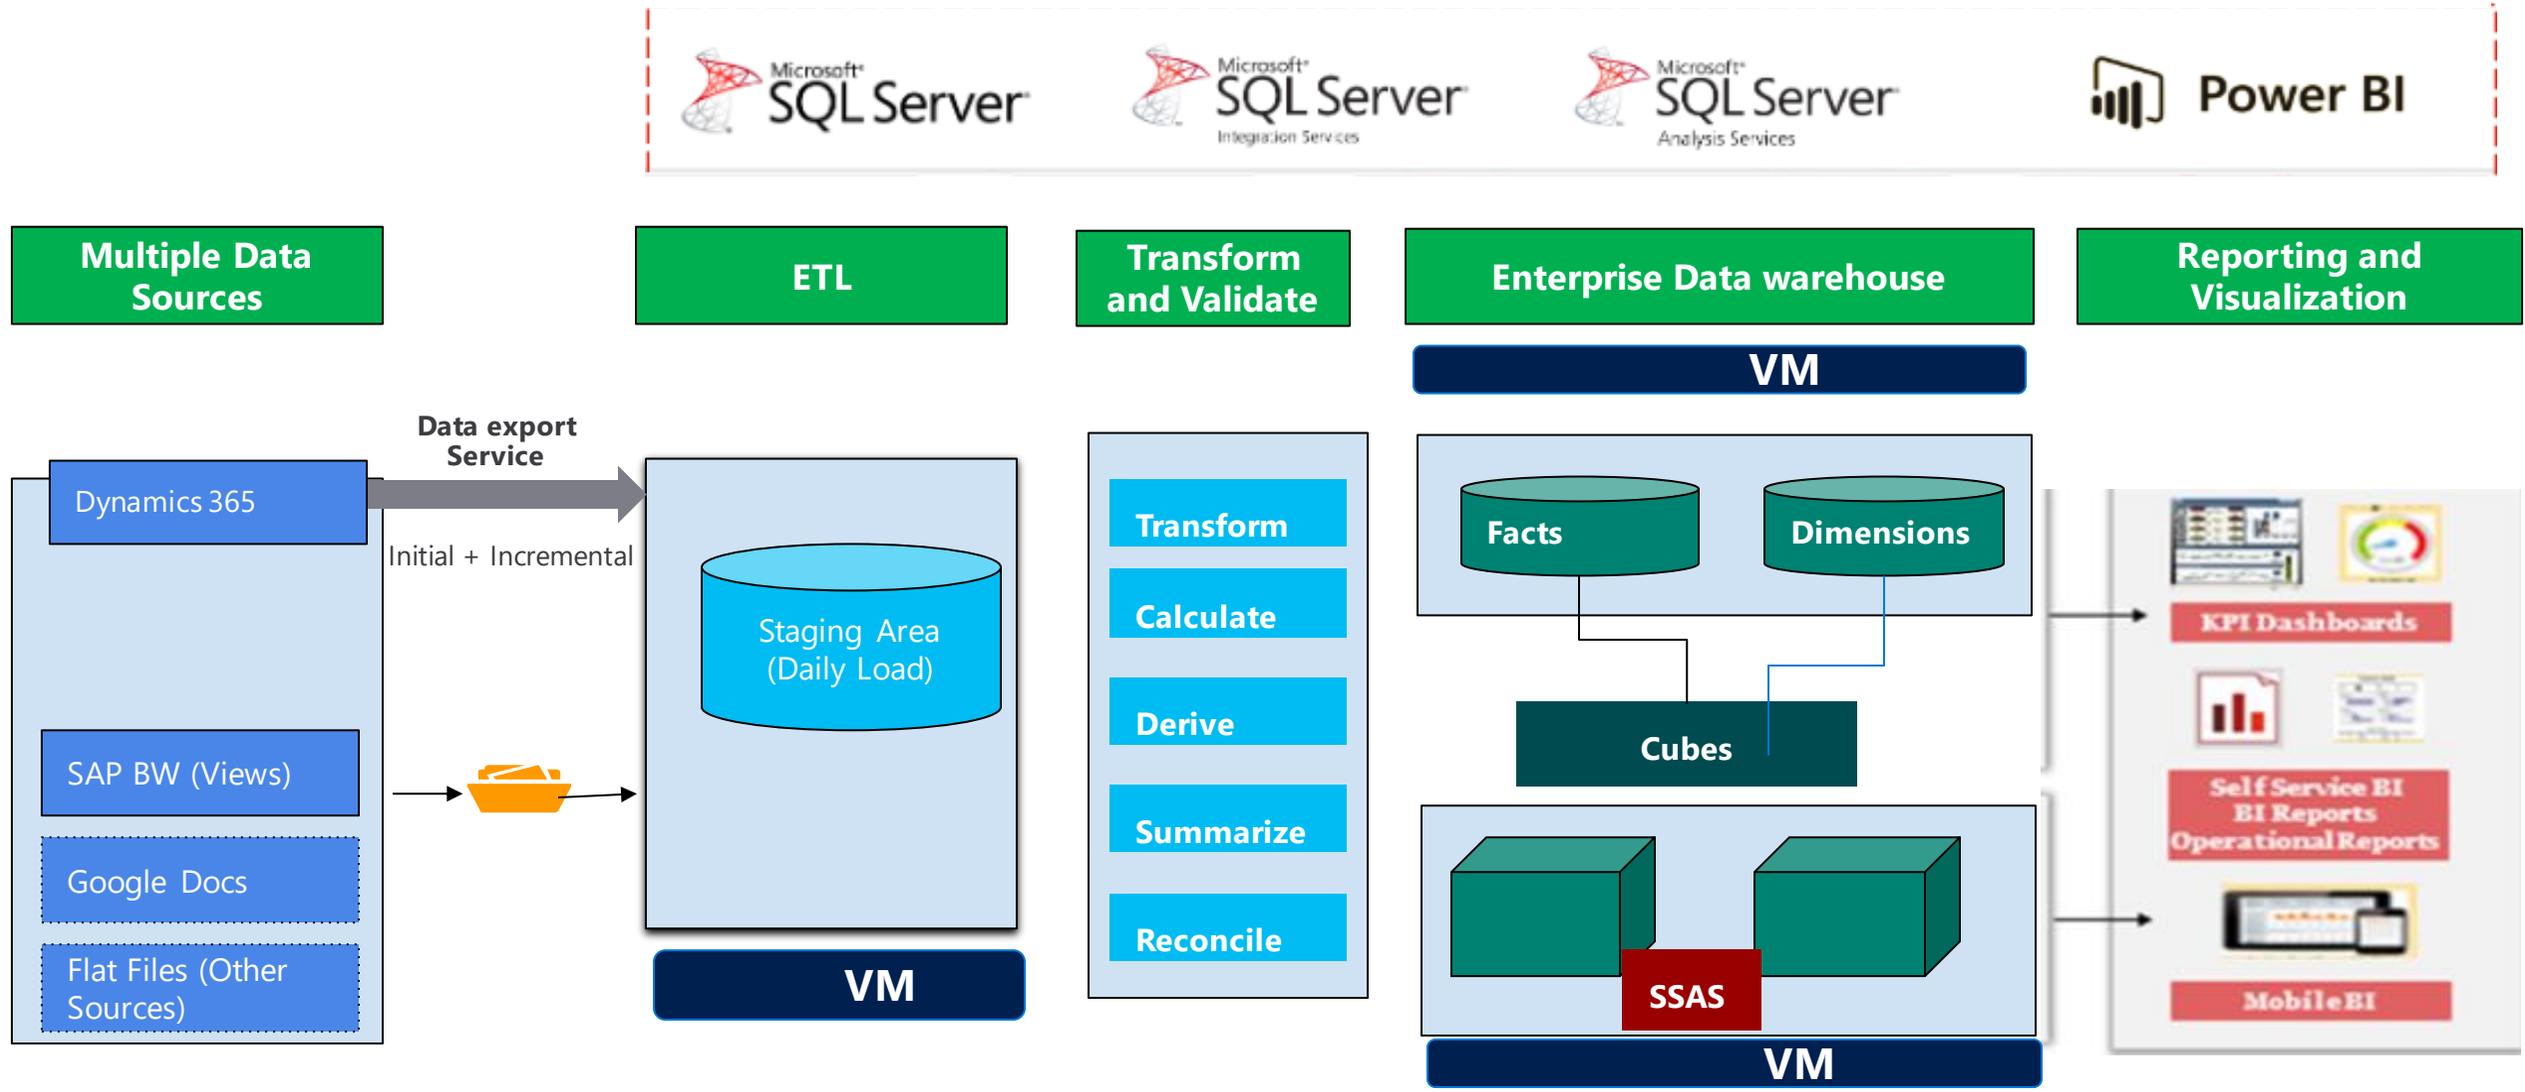
Challenges

- Dynamics data is more than 4TB which is a constraint for Azure SQL DB/VM.
- DES does not support pushing data to Azure Hyperscale.
- Azure SQL VM Replication options and storage costs are higher.
- Legacy Datawarehouse techniques give no agility to reporting or analytics.

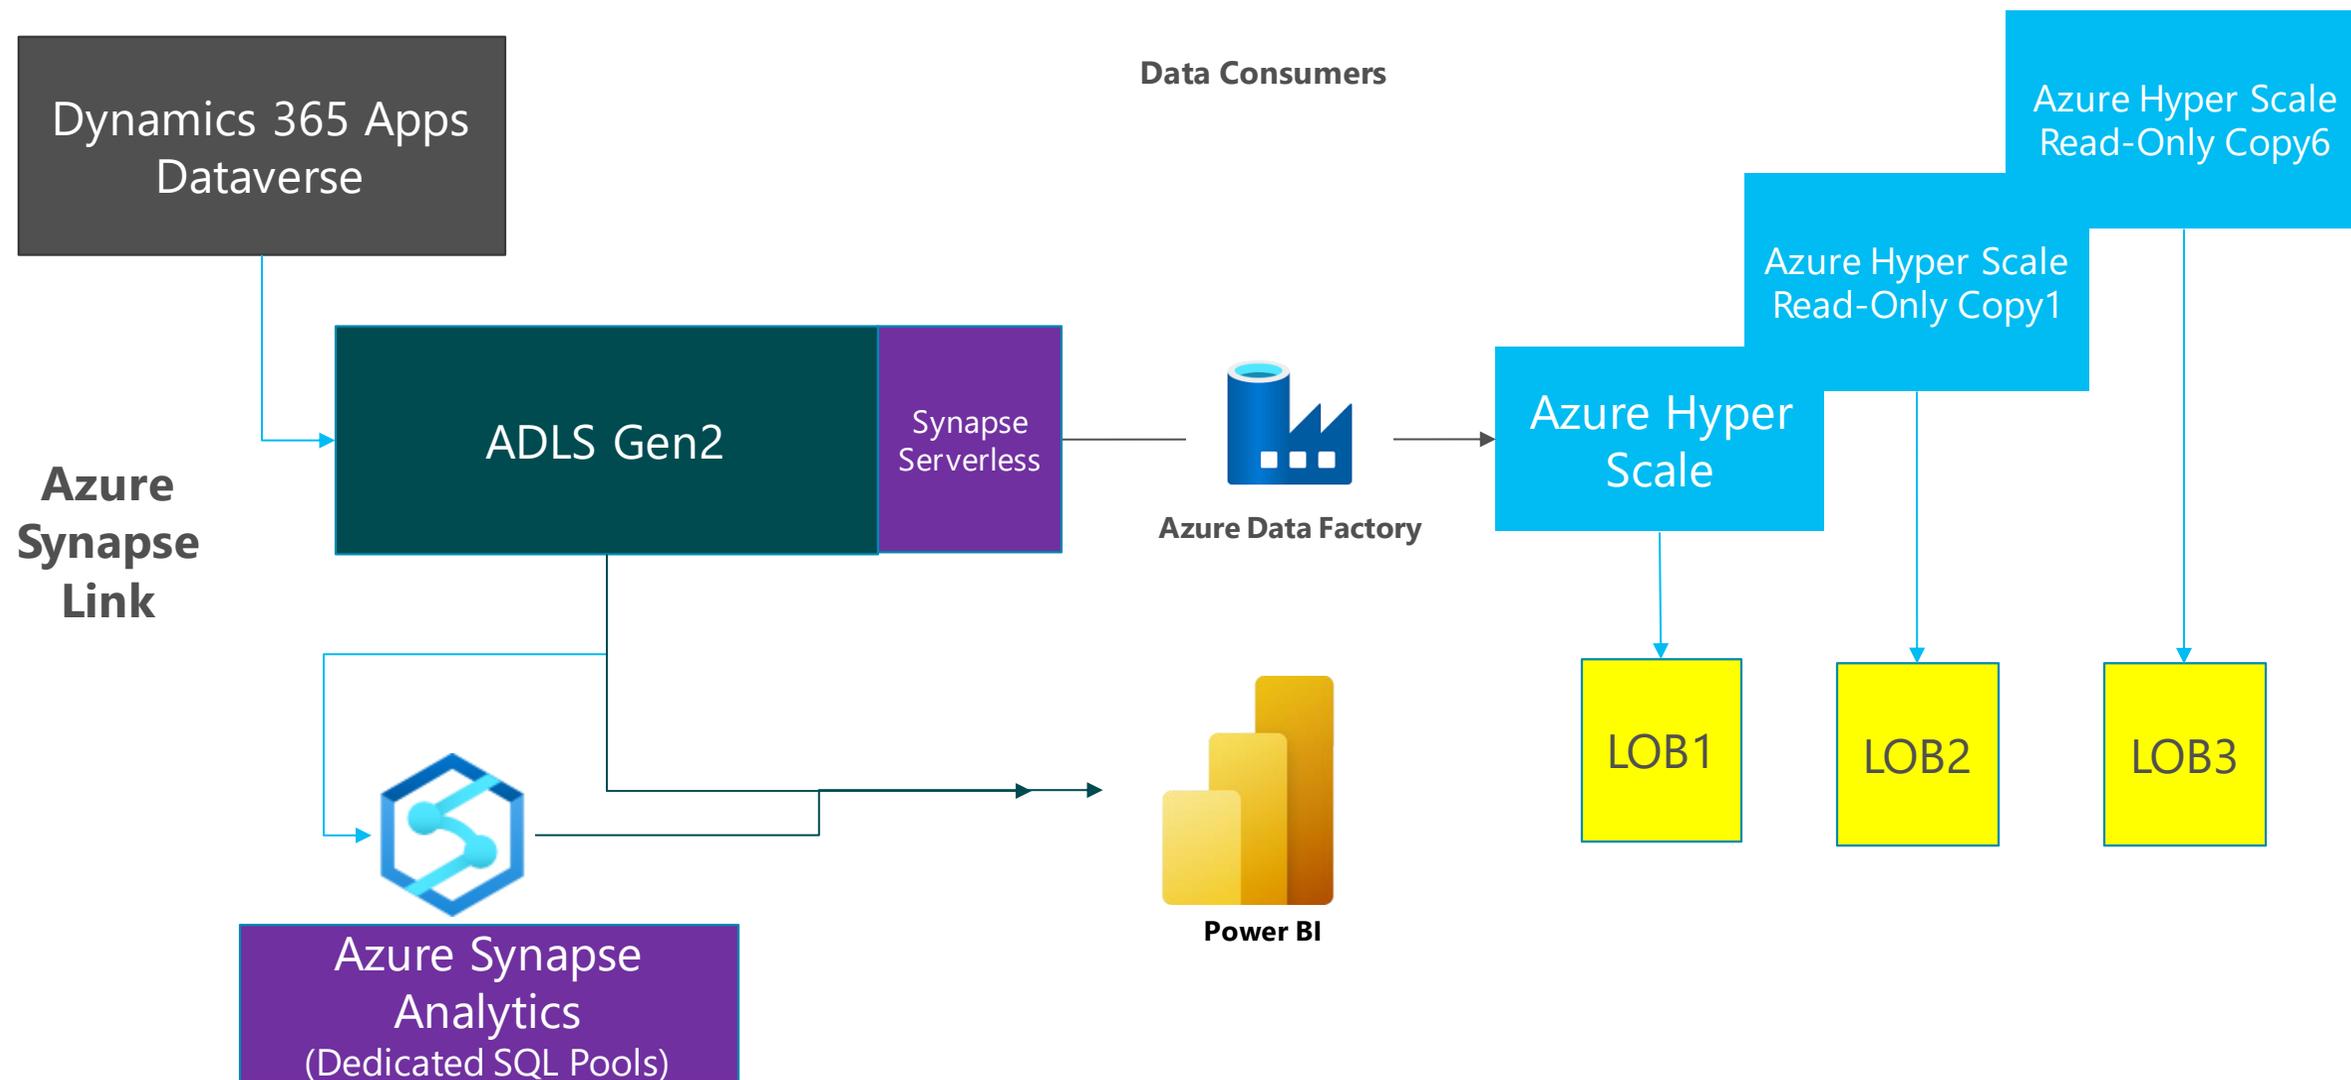
Future Stage

- The exported data from Dynamics 365 CE needs to be available in Azure Hyperscale for upstream systems consumption for frequent usage.
- Azure Hyperscale is chosen because it is scalable from both data growth and multi replica sets and the total cost of ownership is less. Please refer to the architectural diagram on the next slide which represents the current stage architecture.
- Azure Synapse Analytics is to be used for Advanced Analytics and KPI computations; Power BI connects to Azure Synapse Analytics and ADLS Gen II for various organizational reporting and trends.

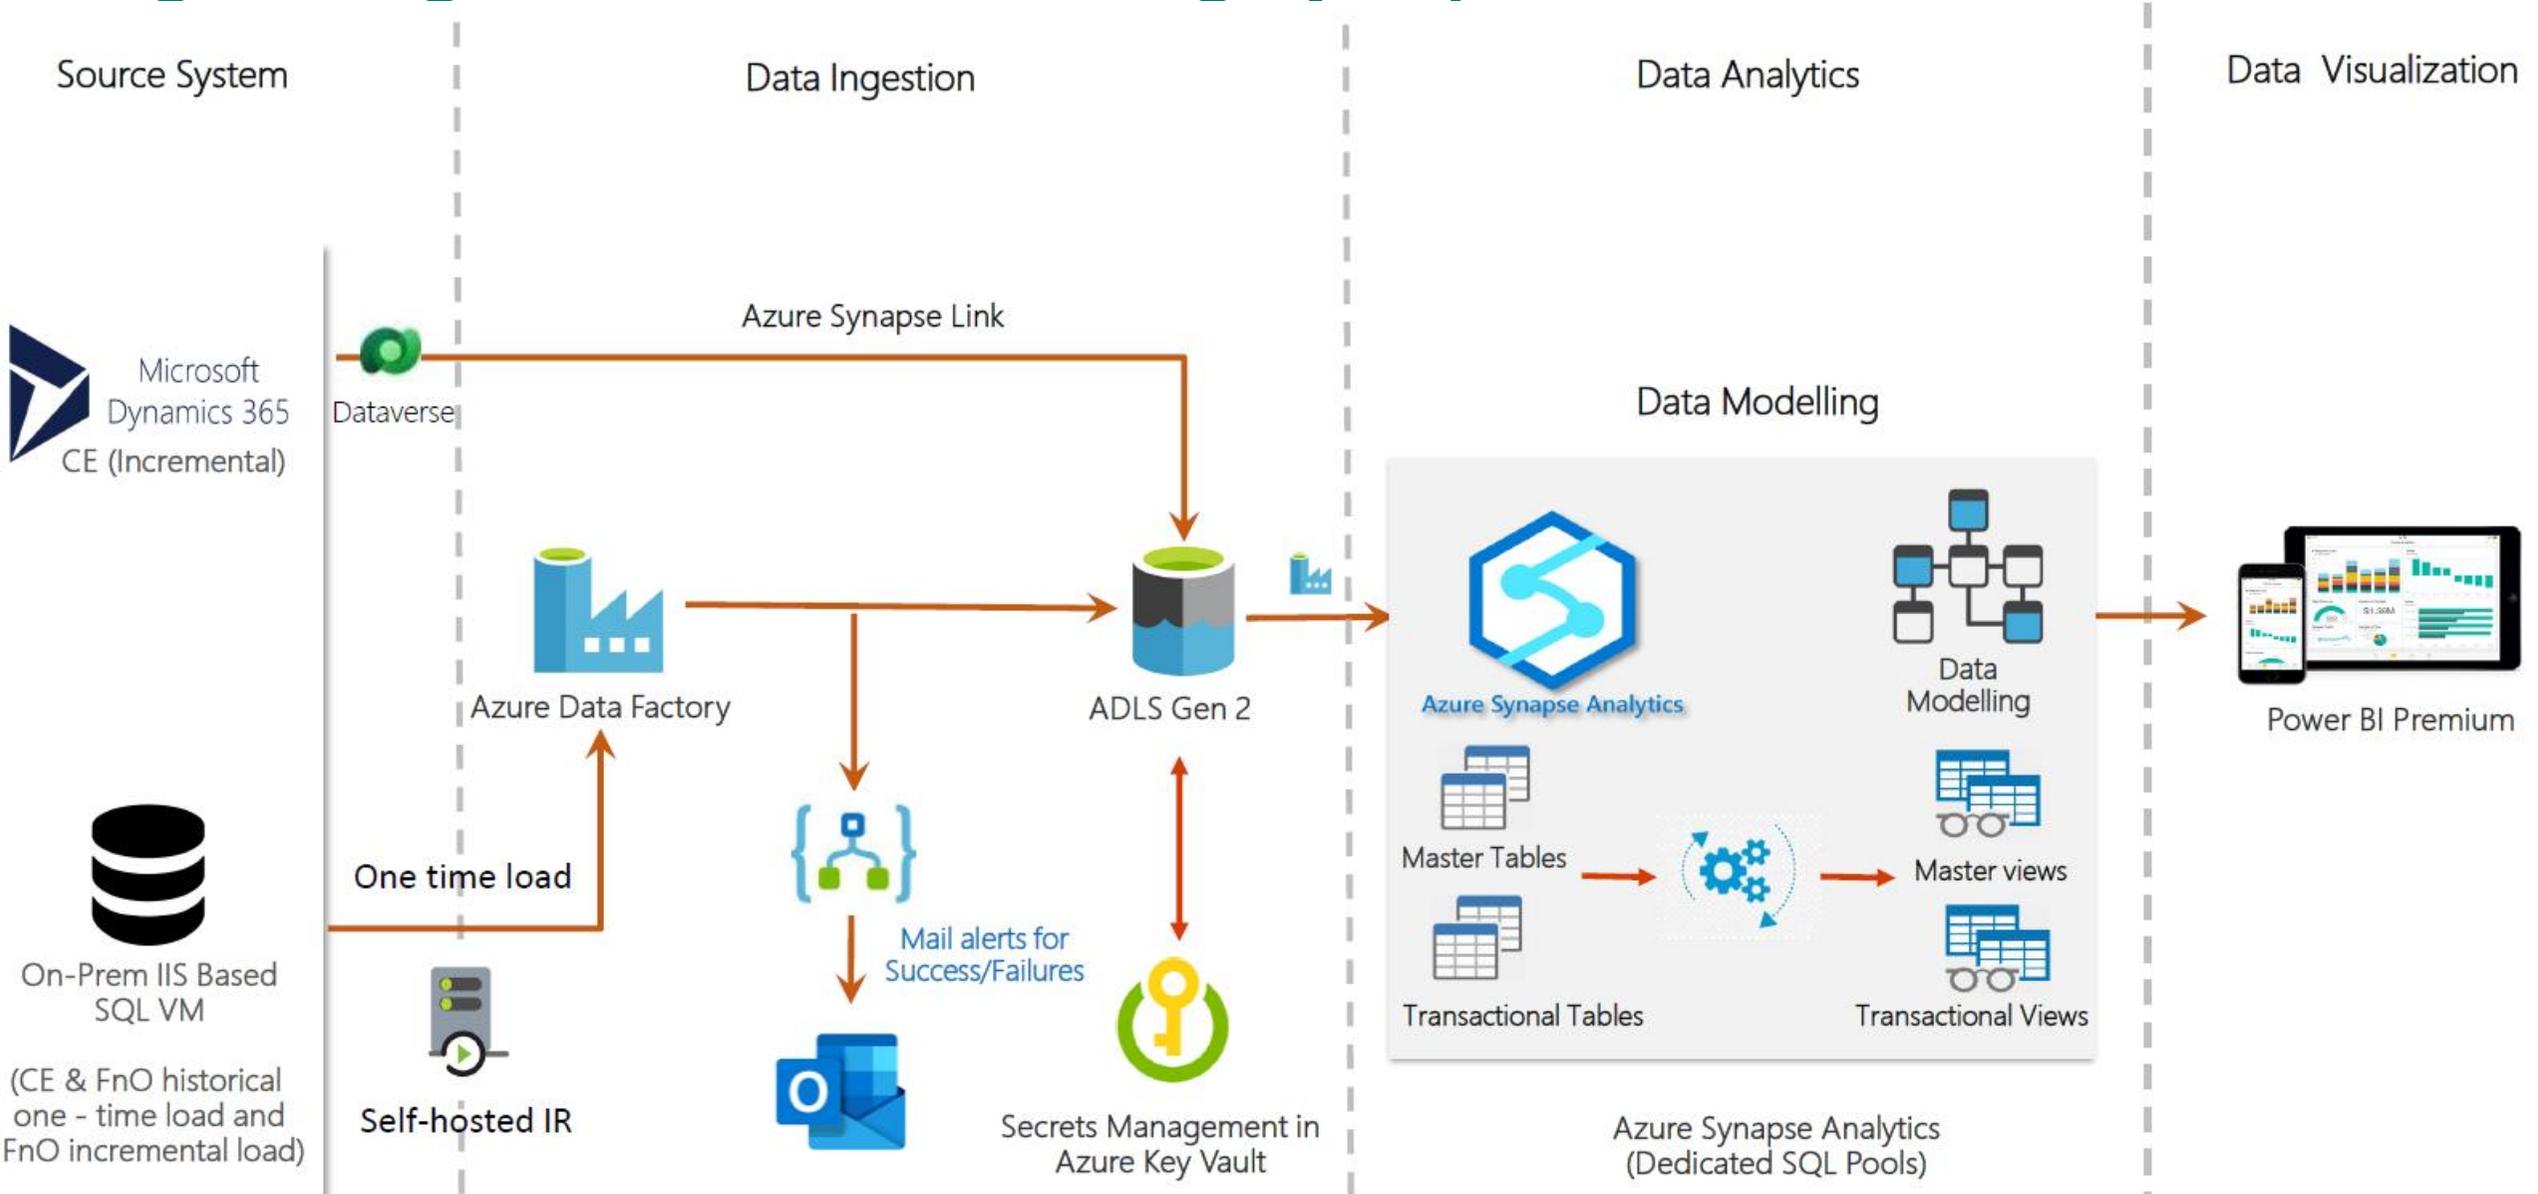
Previous Architecture Diagram



Target Stage Architecture using Synapse Link



Target Stage Architecture using Synapse Link



Key Sites and resources

- [Learn more about Azure Synapse Link for Dataverse](#)
- [Azure Synapse Link for Dataverse FAQ](#)
- [Learn more about Azure Synapse Analytics](#)
- [Configure Azure Synapse Link for Dataverse with your Azure Synapse Workspace](#)
- [Configure Azure Synapse Link for Dataverse with Azure Data Lake](#)
- [Azure Synapse Link for Dataverse Public Preview Announcement](#)
- [Azure Synapse Link for Dataverse General Availability Announcement](#)
- [Synapse Link for Dataverse – Power CAT Live](#)



This playbook is provided “as-is.” Information and views expressed in this playbook, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

Some examples are for illustration only and are fictitious. No real association is intended or inferred.

This playbook does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this playbook for your internal, reference purposes.

© 2021 Microsoft Corporation. All rights reserved.

