## BENNETT ADELSON SELF SERVICE ANALYTICS

Leveraging the Power of the Azure Data Platform and Microsoft's Power Platform

October 28, 2021



## Participants

- Kevin Karlik, Solutions Sales Executive
- Narayan Sujay Somasekhar, Data Architect



#### BENNETT ADELSON<sup>®®®</sup>

#### BENNETT ADELSON<sup>®®</sup>

## Agenda

- Discuss the benefits of combining the use of Azure Data Services and the Microsoft Power Platform
  - How these architectures helps ingest, store, process, enrich, and deliver data and insights to users across your organization
  - The Combination of Azure Data Platform and Power Platform and how it allows engineers to build fully customizable end-to-end, self-serviced solutions with the end User
  - Relevant Use Cases where large amounts of data needed to be processed in an ad-hoc fashion without significant IT investment

### Overview

- In this presentation, we will go over an enterprise ready solution and various techniques that combines a range of Azure data services and Microsoft Power Platform services that will ingest, store, process, enrich, and serve data and insights from different sources to business users and power users.
- Azure data platform enables data engineers to pre-build enterprise ready data platforms working with business users to understand requirements that can be leveraged within Power Platform technologies like Power apps and Power automate to deliver a fully customized end to end self service solution.
- Relevant use cases of this would be in businesses that deal with large volumes of data that need to be processed in an ad hoc fashion without significant investments in IT resources and infrastructure.
- Few examples include Accounting, Tax , Finance and Audit service lines.



## **Design Principles we consider**

We focus on five pillars during our solution and architecture design:

- □ Reliability
  - $\checkmark$  The ability of a system to recover from failures and continue to function.
- □ Security
  - ✓ Protecting applications and data from threats.
- Cost Optimization
  - ✓ Managing costs to maximize the value delivered.
- Operational Excellence
  - ✓ Operations processes that keep a system running in production.
- Performance Efficiency
  - $\checkmark$  The ability of a system to adapt to changes in load.



#### Microsoft Azure data platform reference architecture





# Modern data warehouse for small and medium business

- Choices for Small and medium businesses (SMBs) when modernizing their on-premises data warehouses for the cloud are
  - Adopt big data tools for future extensibility
  - Keep traditional, SQL-based solutions for cost efficiency, ease of maintenance and smooth transition.
  - Hybrid approach combines easy migration of the existing data estate with the opportunity to add big data tools and processes for some use cases. SQL-based data sources can keep running in the cloud and continue to modernize as appropriate.
- This example workload shows several ways that SMBs can modernize legacy data stores and explore big data tools and capabilities, without overextending current budgets and skillsets. These end-to-end Azure data warehousing solutions integrate easily with Azure and Microsoft services and tools like Azure Machine Learning, Microsoft Power Platform, and Microsoft Dynamics.

#### BENNETT ADELSON<sup>sta</sup>

## **Potential use cases**

- Several scenarios can benefit from this workload:
- Migrating a traditional, on-premises relational data warehouse that's smaller than 1 TB and extensively uses SQL Server Integration Services (SSIS) packages to orchestrate stored procedures.
- Meshing existing Dynamics or Power Platform Dataverse data with batched and realtime Azure Data Lake sources.
- Using innovative techniques to interact with centralized Data Lake Storage data. Techniques include serverless analysis, knowledge mining, data fusion between domains, and end-user data exploration.

#### BENNETT ADELSON<sup>sm</sup>



## **SMB** Architecture

![](_page_8_Picture_1.jpeg)

![](_page_8_Figure_2.jpeg)

## **Cloud Data Lakehouse Architecture**

![](_page_9_Figure_1.jpeg)

10

#### **Cloud Lakehouse Architecture(Snowflake as DW)**

![](_page_10_Figure_1.jpeg)

11

## **Raw data complexity**

- Legacy SMB data warehouses might contain several types of data:
  - Unstructured data, like documents and graphics
  - Semi-structured data, such as logs, CSVs, JSON, and XML files
  - Structured relational data, including databases that use stored procedures for extract-transform-load/extract-loadtransform (ETL/ELT) activities

![](_page_11_Picture_6.jpeg)

## Ingestion of your choice based on data complexity

- Azure Synapse Analytics pipelines or ADF v2
  - The pipelines orchestrate the flow of migrated or partially refactored legacy databases and SSIS packages into Azure SQL Database. This lift-and-shift approach is fastest to implement and offers a smooth transition from an on-premises SQL solution to an eventual Azure platform-as-a-service (PaaS). You can modernize databases incrementally after the lift and shift.
  - The pipelines can also pass unstructured, semi-structured, and structured data into Azure Data Lake Storage for centralized storage and analysis with other sources. Use this approach when fusing data provides more business benefit than simply replatforming the data.
- Apache Spark Platform like Databricks for code-based pipelines for reading data from storage, processing with DBFS and pushing processed data to relational DW layer
- Using In database ELT approaches like below
  - Snowpipe if using Snowflake as the DW
  - Synapse Polybase if using Azure Synpase as the DW
  - Redshift spectrum if using AWS Redshift as the DW

![](_page_12_Figure_9.jpeg)

#### BENNETT ADELSON

13

## Ingestion of your choice based on data complexity - Alternatives

- Azure IoT Hub could replace or complement Event Hubs. The solution you choose depends on the source of your streaming data, and whether you need cloning and bidirectional communication with the reporting devices.
- You can use Azure Data Factory for data integration instead of Azure Synapse pipelines. The choice depends on several factors:
  - Azure Synapse pipelines keep the solution design simpler and allow collaboration inside a single Azure Synapse workspace.
  - Azure Synapse pipelines don't support SSIS packages rehosting, which is available in Azure Data Factory.
  - Synapse Monitor Hub monitors Azure Synapse pipelines, while Azure Monitor can monitor Data Factory.
- More information and a feature comparison between Azure Synapse pipelines and Data Factory is available Microsoft website -Data integration in Azure Synapse Analytics versus Azure Data Factory.
- Synapse Analytics dedicated SQL pools used for storing enterprise data, instead of using SQL Database.

#### BENNETT ADELSON"

![](_page_13_Picture_9.jpeg)

## Storage

- Primary Storage layer we typically adopt is Azure Storage
- Based on scenario of SMB architecture and requirements we use the following
  - ADLS Gen2
  - Blob Storage
  - Table Storage
- In a multi cloud architecture
  - AWS S3
  - GCP Cloud Storage

![](_page_14_Picture_9.jpeg)

![](_page_14_Picture_10.jpeg)

## Storage Cost Considerations

- Based on application Requirements we categorize data into 3 main categories
  - Hot
  - Cold
  - Archive
- Architecture design focus on
  - limited data movement to reduce ingress and egress cost
  - Redundancy (LRS, ZRS, GRS, RA-GRS, GZRS, RA-GZRS)

Not available	Premium	Hot	Cool	Archive
First 50 terabyte (TB) / month	\$0.15 per GB	\$0.018 per GB	\$0.01 per GB	\$0.00099 per GB
Next 450 TB/month	\$0.15 per GB	\$0.0173 per GB	\$0.01 per GB	\$0.00099 per GB
Over 500 TB/month	\$0.15 per GB	\$0.0166 per GB	\$0.01 per GB	\$0.00099 per GB

![](_page_15_Picture_9.jpeg)

![](_page_15_Picture_10.jpeg)

## Processing

#### BENNETT ADELSON<sup>sta</sup>

- Architecture Design pattern primarily based on batch processing
- Technology choices for batch processing
  - Azure Databricks
  - Azure Synapse
  - Snowflake
  - AWS Redshift
  - Azure SQL DB
- Key selection criteria we consider
  - To narrow the choices, start by answering these questions:
  - Do you want a managed service rather than managing your own servers?
  - Do you want to author batch processing logic declaratively or imperatively?
  - Will you perform batch processing in bursts? If yes, consider options that let you auto-terminate the cluster or whose pricing model is per batch job.
  - Do you need to query relational data stores along with your batch processing, for example to look up reference data?

## Data Modeling Layer

- Power BI XMLA endpoint
- Azure Analysis Services
- Direct connections
- Import

![](_page_17_Picture_5.jpeg)

#### BENNETT ADELSON<sup>®®</sup>

![](_page_17_Picture_7.jpeg)

## Data Modeling Layer

![](_page_18_Figure_1.jpeg)

![](_page_18_Picture_2.jpeg)

- Relevant use cases
- This approach can also be used to:
- Establish a data warehouse to be a single source of truth for your data.
- Integrate relational data sources with other unstructured datasets.
- Use semantic modeling and powerful visualization tools for simpler data analysis.

## Presentation & Visualization with Microsoft Power Platform

- PowerApps
  - More control for end users
  - Trigger data pipelines
- Power Automate
  - Build complex flows that go along with the batch processing of data
  - Integration with document management platforms like Teams / SharePoint
- Power Bl
  - Advanced dashboarding and visualizations

![](_page_19_Figure_9.jpeg)

#### BENNETT ADELSON"

### Differences in our development methodology

• With traditional app development such as the "waterfall" model, there's a long lead time before a user actually sees the working app. As a result, there's an increased risk that there will be gaps between what the user initially requested as a requirement and what the app developer created.

![](_page_20_Figure_2.jpeg)

• Even with more modern development approaches such as agile development, a significant amount of time can pass before the first minimum viable product (MVP) is delivered to

![](_page_20_Figure_4.jpeg)

With Power Platform, you can quickly create a usable version of your app, because Power Apps provides a WYSIWYG (what you see is what you get) development experience. Users experience the actual working app very early in the development process, and if new requirements arise, new features can be added to the next version.

![](_page_20_Figure_6.jpeg)

#### BENNETT ADELSON"

![](_page_20_Picture_8.jpeg)

#### Demo

• Trigger ADF Pipelines and Data workflows via Power Apps

![](_page_21_Figure_2.jpeg)

![](_page_21_Picture_3.jpeg)

#### Demo

• Trigger Snowpipe and Data workflows via Power Apps & Azure Storage Even Subscriptions

![](_page_22_Figure_2.jpeg)

![](_page_22_Picture_3.jpeg)

#### Demo

• Image Processing with PowerApps and Azure Cognitive Services

![](_page_23_Figure_2.jpeg)

![](_page_23_Picture_3.jpeg)

## BENNETT ADELSON CAN HELP

## Thank You

Brian Connelly Senior Solution Sales Executive bconnelly@bennettadelson.com 216-338-1524

Grant Margrett Senior Solution Sales Executive <u>gmargrett@bennettadelson.com</u> 440-725-5215

Kevin Karlik Senior Solution Sales Executive <u>kkarlik@bennettadelson.com</u> 216-256-6796

> BENNETT ADELSON