



Modern Analytics Academy “Vignettes”

Follow up to Modern Analytics Academy: <https://aka.ms/maa>

Topical sessions that addresses:

- New features and how to use them
- Frequent questions we receive
- Updates to a previous session



Pardeep Singla

Cloud Solution Architect



Modern Analytics Academy “Vignettes”

Azure Synapse Link for Cosmos DB

<https://aka.ms/maa>



Pardeep Singla

Cloud Solution Architect

Academy Sessions



Acquisition & Storage

- Azure Data Lake
- Azure Cosmos DB
- Azure Event Hubs
- Synapse Link



Modeling

- Data Lake Structure
- Synapse Spark Pools
- Synapse SQL Pools
- Synapse Serverless SQL



Pipelines

- Azure Data Factory
- Synapse Pipelines
- Power BI Dataflows
- Azure Stream Analytics



Security & Governance

- Auditing
- Security
- Azure Purview



Visualization

- Power BI
- Paginated Reports
- Power BI Embedded

Agenda

Azure Synapse Link for Cosmos
DB

- **Overview & Use Cases**
- Demo

Overview

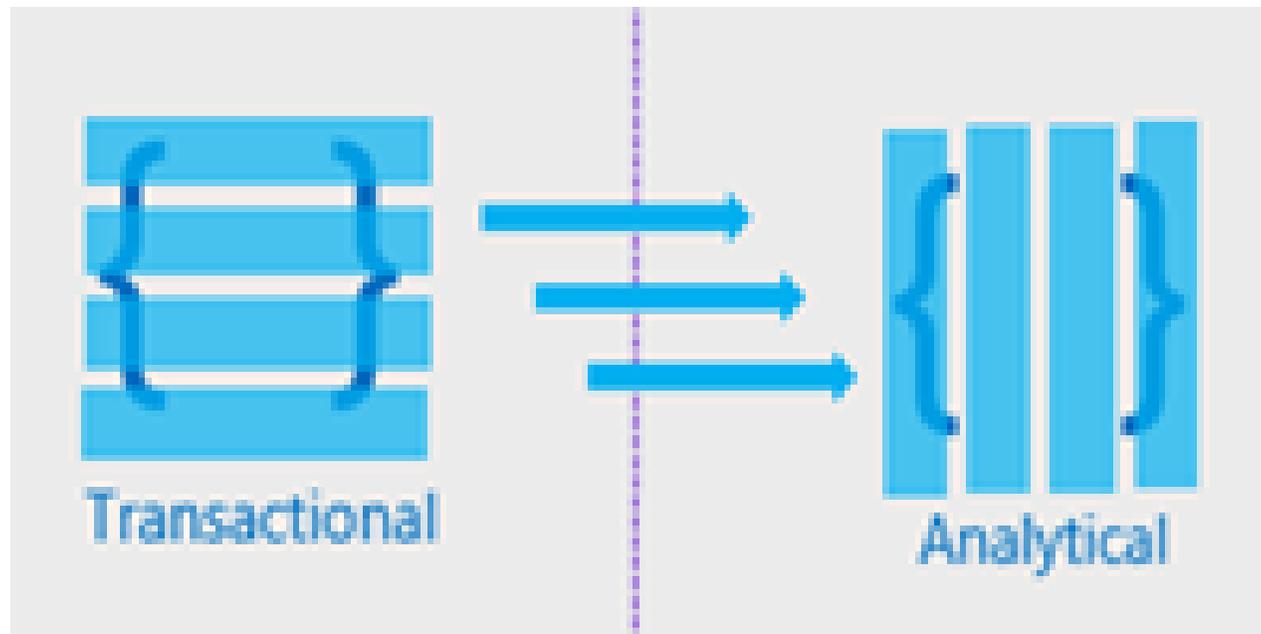
- Azure Synapse Link Availability
- Understand HTAP Architecture
- Synapse link for Cosmos DB Architecture

Synapse Link Availability

- Synapse link for Cosmos DB
- Synapse link for SQL
- Synapse link for Dataverse

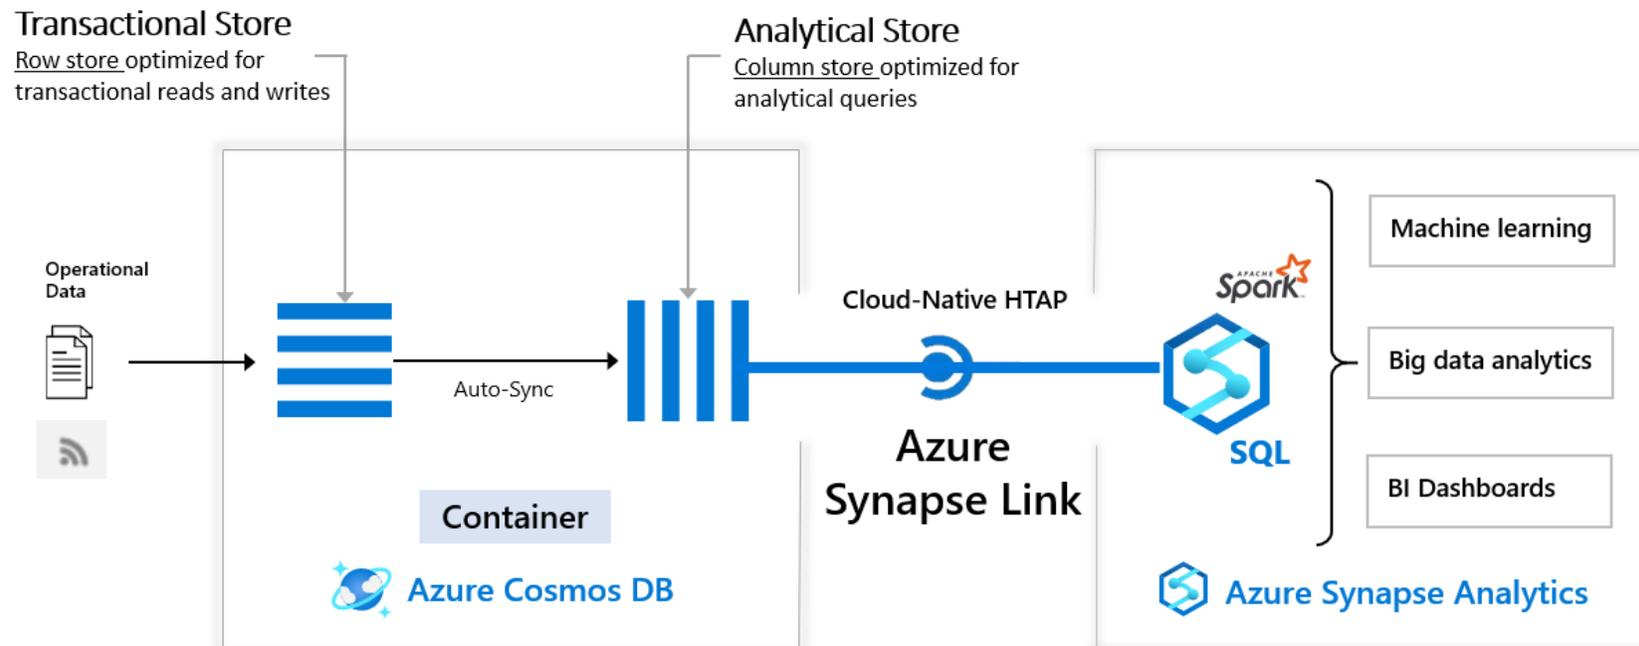
Understand hybrid transactional and analytical processing patterns

Many business application architectures separate transactional and analytical processing into separate systems with data stored and processed on separate infrastructures.



Azure Synapse Link for Azure Cosmos DB

Azure Synapse Link for Azure Cosmos DB is a cloud-native HTAP capability that enables you to run near-real-time analytics over operational data stored in Azure Cosmos DB.

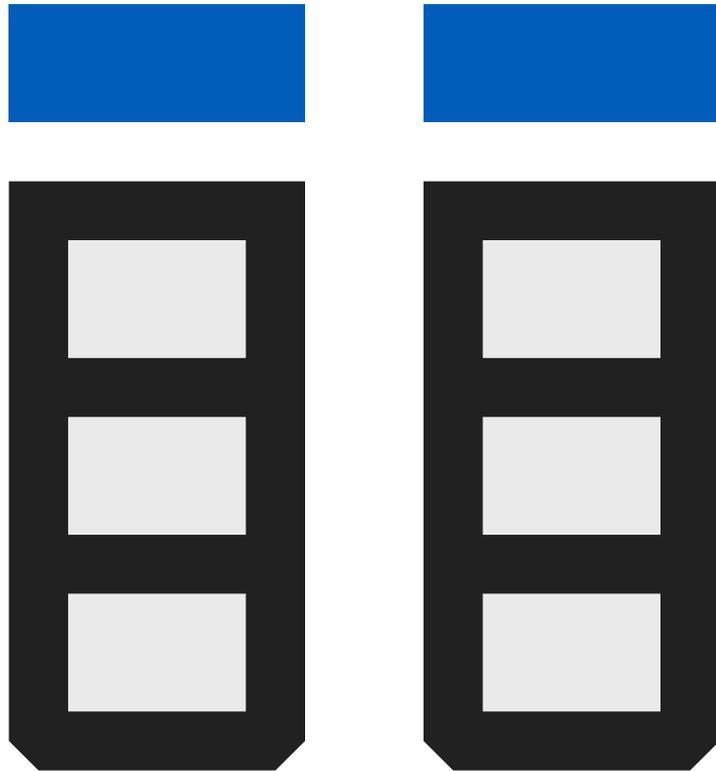


Row-oriented transactional store

Operational data in an Azure Cosmos DB container is internally stored in an indexed row-based "transactional store".



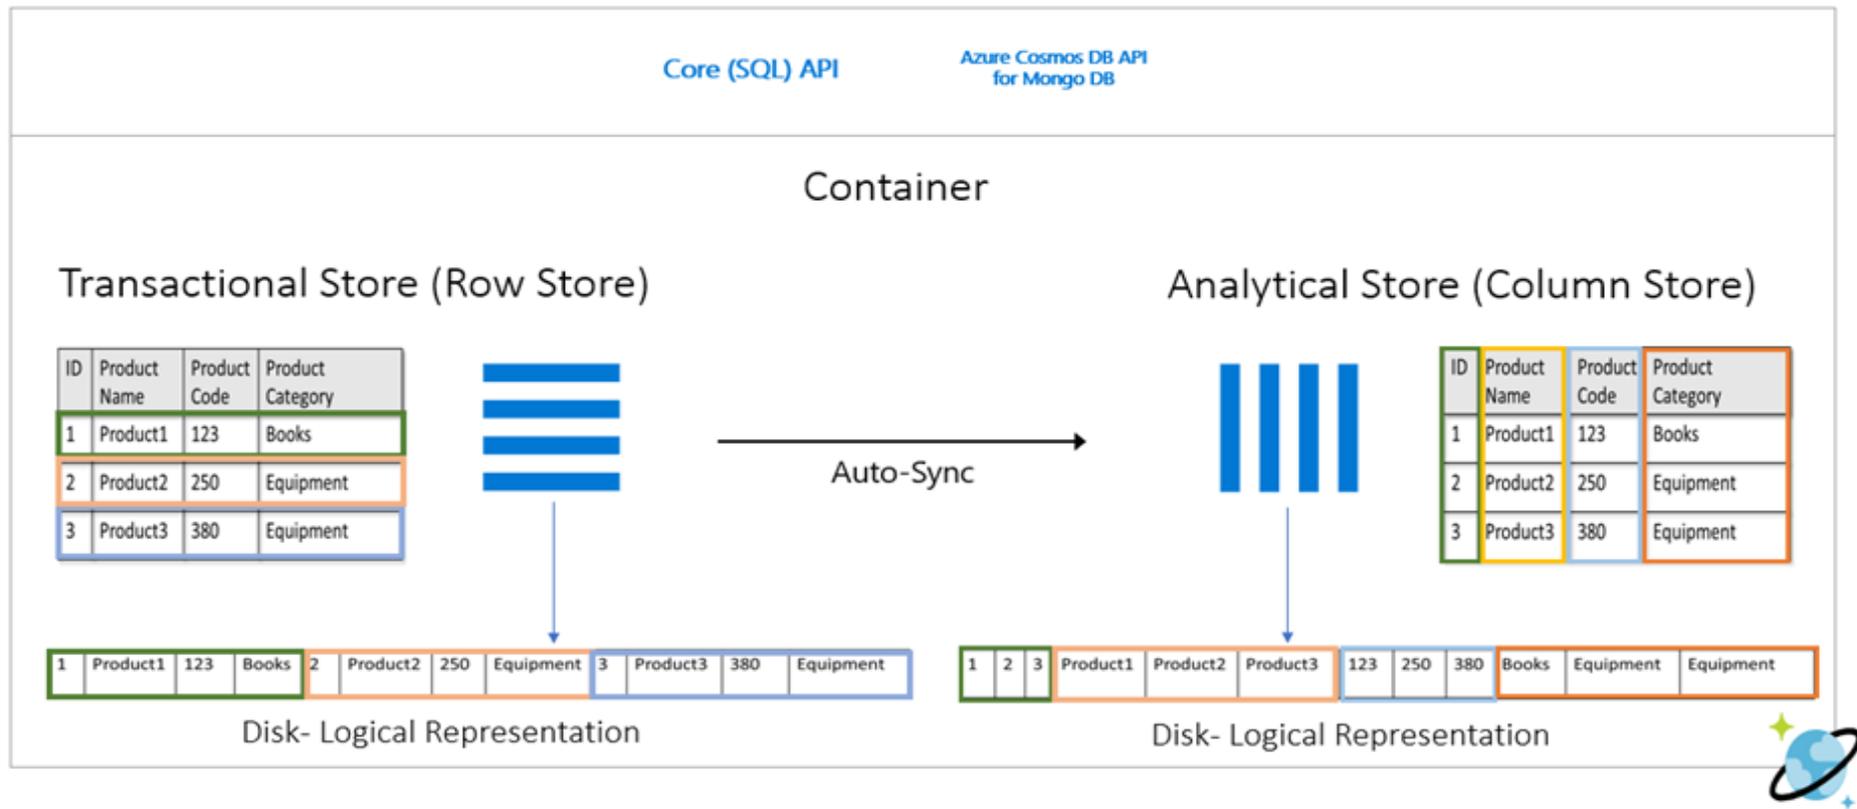
Column-oriented analytical store



Azure Cosmos DB analytical store is a fully isolated column store for enabling large-scale analytics against operational data in your Azure Cosmos DB, without any impact to your transactional workloads.

Providing the best of both worlds

Azure Cosmos DB analytical store addresses the complexity and latency challenges that occur with the traditional ETL pipelines.



Azure Synapse Link for Cosmos DB - Key Benefits

- Automatically sync transactional data to analytical store
 - No store or ETL management
 - No performance impact on transactional workloads
 - Near real-time: From 30 seconds to ~2 mins for most workloads
- Cost Reduction
 - No Cosmos DB OLTP RUs are used
 - Consumption-based pricing for analytical storage and queries
 - Move the expensive cross partition queries to Synapse Analytics
 - Data tiering
 - Hot data in transactional store (RUs / GB)
 - All data in analytical store (~90% compression and 10% of the cost)

Azure Synapse Link for Cosmos DB - Key Benefits

- 1-click integration with Synapse Analytics:
 - Native integration with Synapse Spark, unblocking Data Science on Cosmos DB data
 - SQL serverless T-SQL language: joins, stored procedures, aggregations, natural integration with BI Tools
- Data Modeling with focus on data ingestion since you can run your queries in Synapse Analytics, reading from analytical store instead of transactional store.

Features of the analytical store

Note: Azure Synapse Link is supported for the Azure Cosmos DB SQL (Core) API and for the Azure Cosmos DB API for MongoDB.

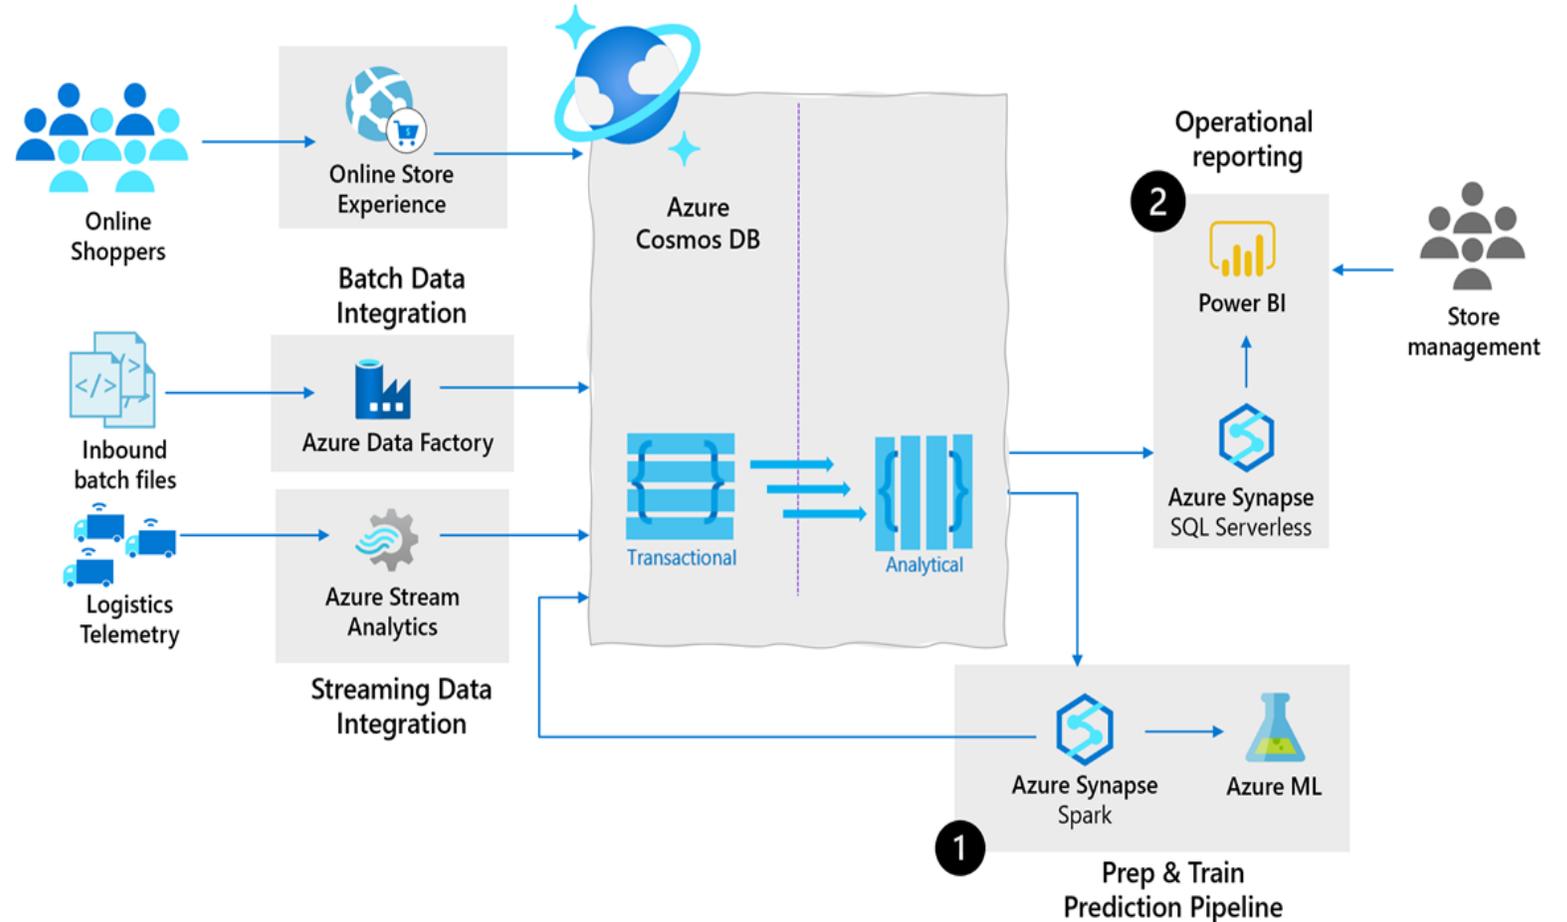
Note: Currently in MongoDB API, you can only enable analytical store at the time of creating a new container.

Use Cases

- Supply chain analytics, forecasting and reporting
- Retail real time personalization
- Predictive maintenance using anomaly detection with IOT

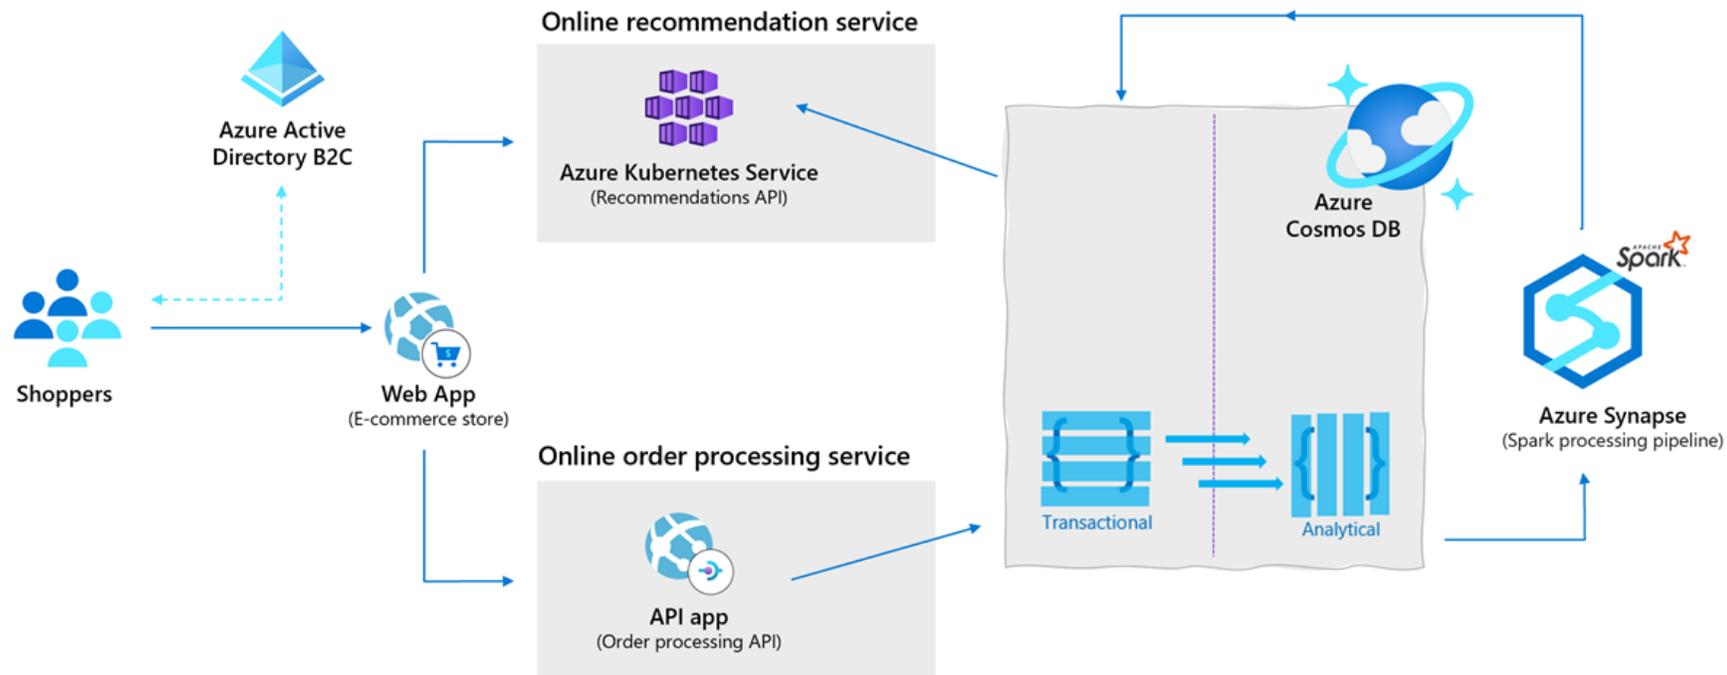
Supply chain analytics, forecasting and reporting.

With supply chains generating increasing volumes of operational data every minute for orders, shipments and sales transactions, manufactures and retailers need an operational database that can scale to handle the data volumes as well as an analytical platform to get to a level of real-time contextual intelligence to stay ahead of the curve.



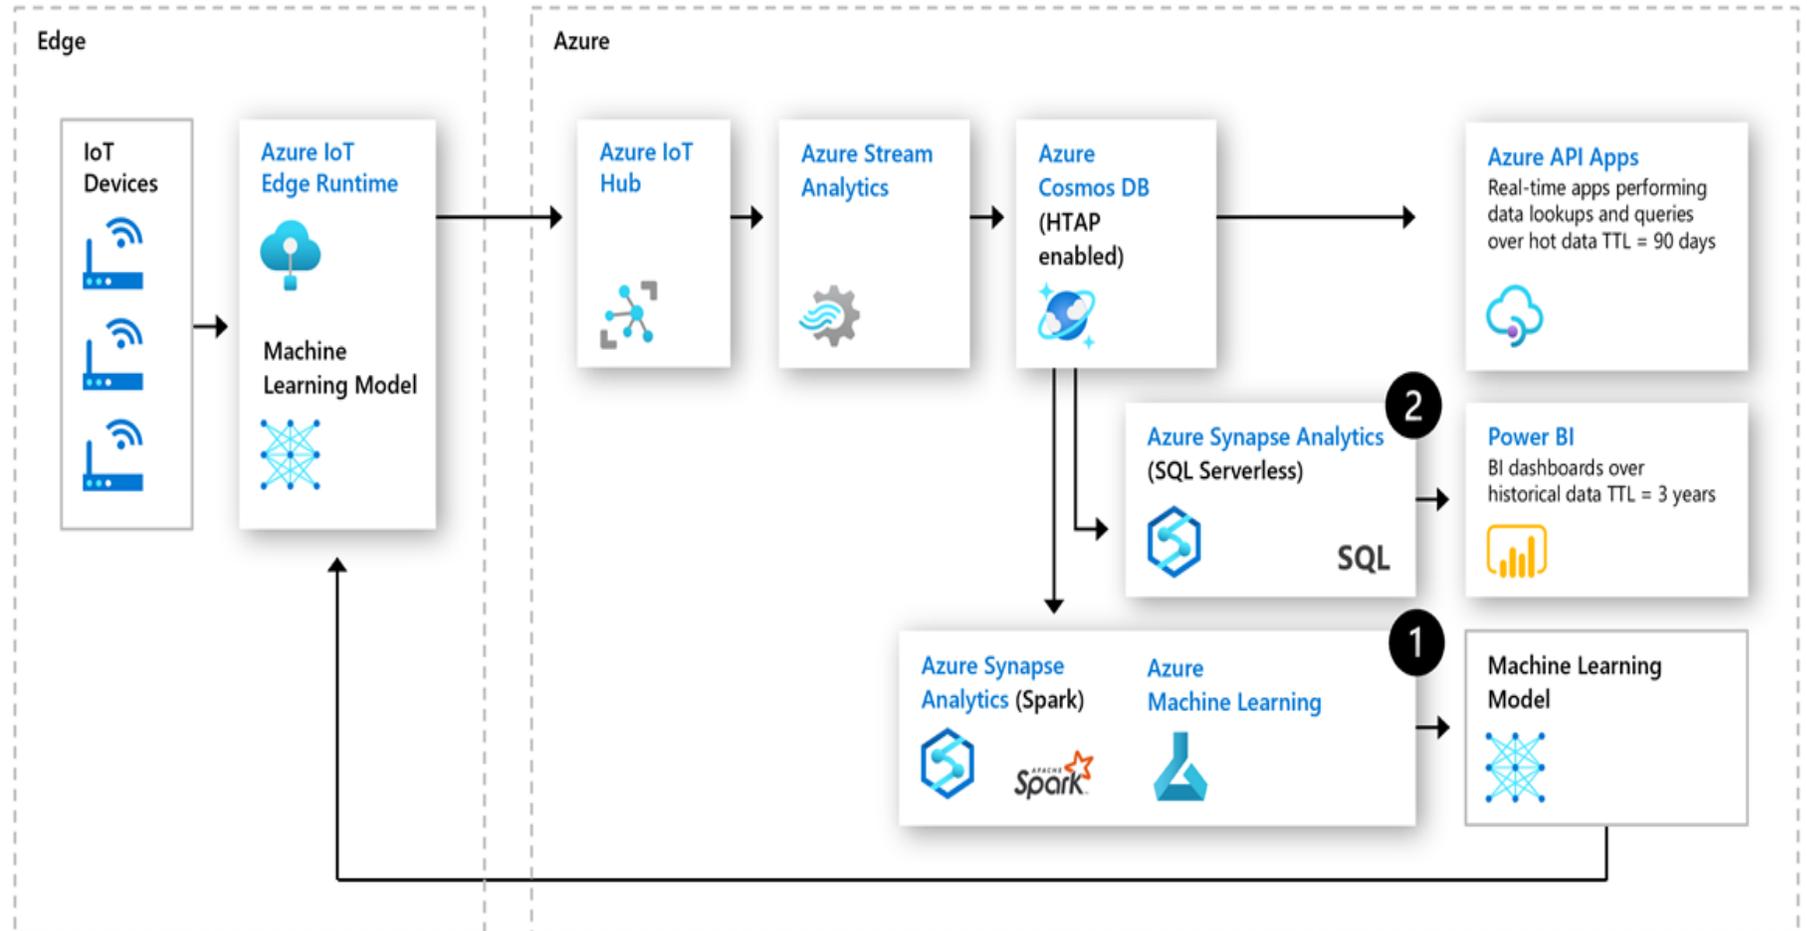
Retail real-time personalization.

In retail, many web-based retailers will perform real-time basket analysis to make product recommendations to customers who are about to purchase products.



Predictive maintenance using anomaly detection with IOT

Industrial IOT innovations have drastically reduced downtimes of machinery and increased overall efficiency across all fields of industry.



Agenda

Azure Synapse Link for Cosmos
DB

- Overview
- Use Cases
- **Demo**

Synapse Link for Cosmos DB– Demo

Pre-requisites

- Setup Azure Cosmos DB account
- Setup Azure Synapse Analytics

References

Cosmos DB main documentation

- <https://docs.microsoft.com/en-us/azure/cosmos-db/>
- All Cosmos DB documentation: APIs, concepts, migration, examples

Synapse Link for Cosmos DB documentation

- <https://docs.microsoft.com/en-us/azure/cosmos-db/synapse-link>
- Benefits, when to use, limitations, getting started



Modern Analytics Academy “Vignettes”

Azure Synapse Link for Cosmos DB

<https://aka.ms/maa>

Thank you!