Azure Data Explorer (Kusto)

CALGARY AZURE ANALYTICS USER GROUP
19-10-2020

VINCENT-PHILIPPE LAUZON
CLOUD SOLUTION ARCHITECT (CSA, DATA & AI)
MICROSOFT CANADA

What is
Azure Data Explorer (ADX)?

Factually…

New product (2018), used internally for 5 years

- Foundation of multiple Azure Services
  - Azure Monitor
  - Azure Security
  - Time Series Insights (TSI)
  - Xbox Playfab
  - Microsoft Connected Vehicle Platform (MCVP)

MS Proprietary Technology

Analytic Database (ingested data) with highly optimized ad hoc analytic queries capabilities

Big Data in Modern Business Environment

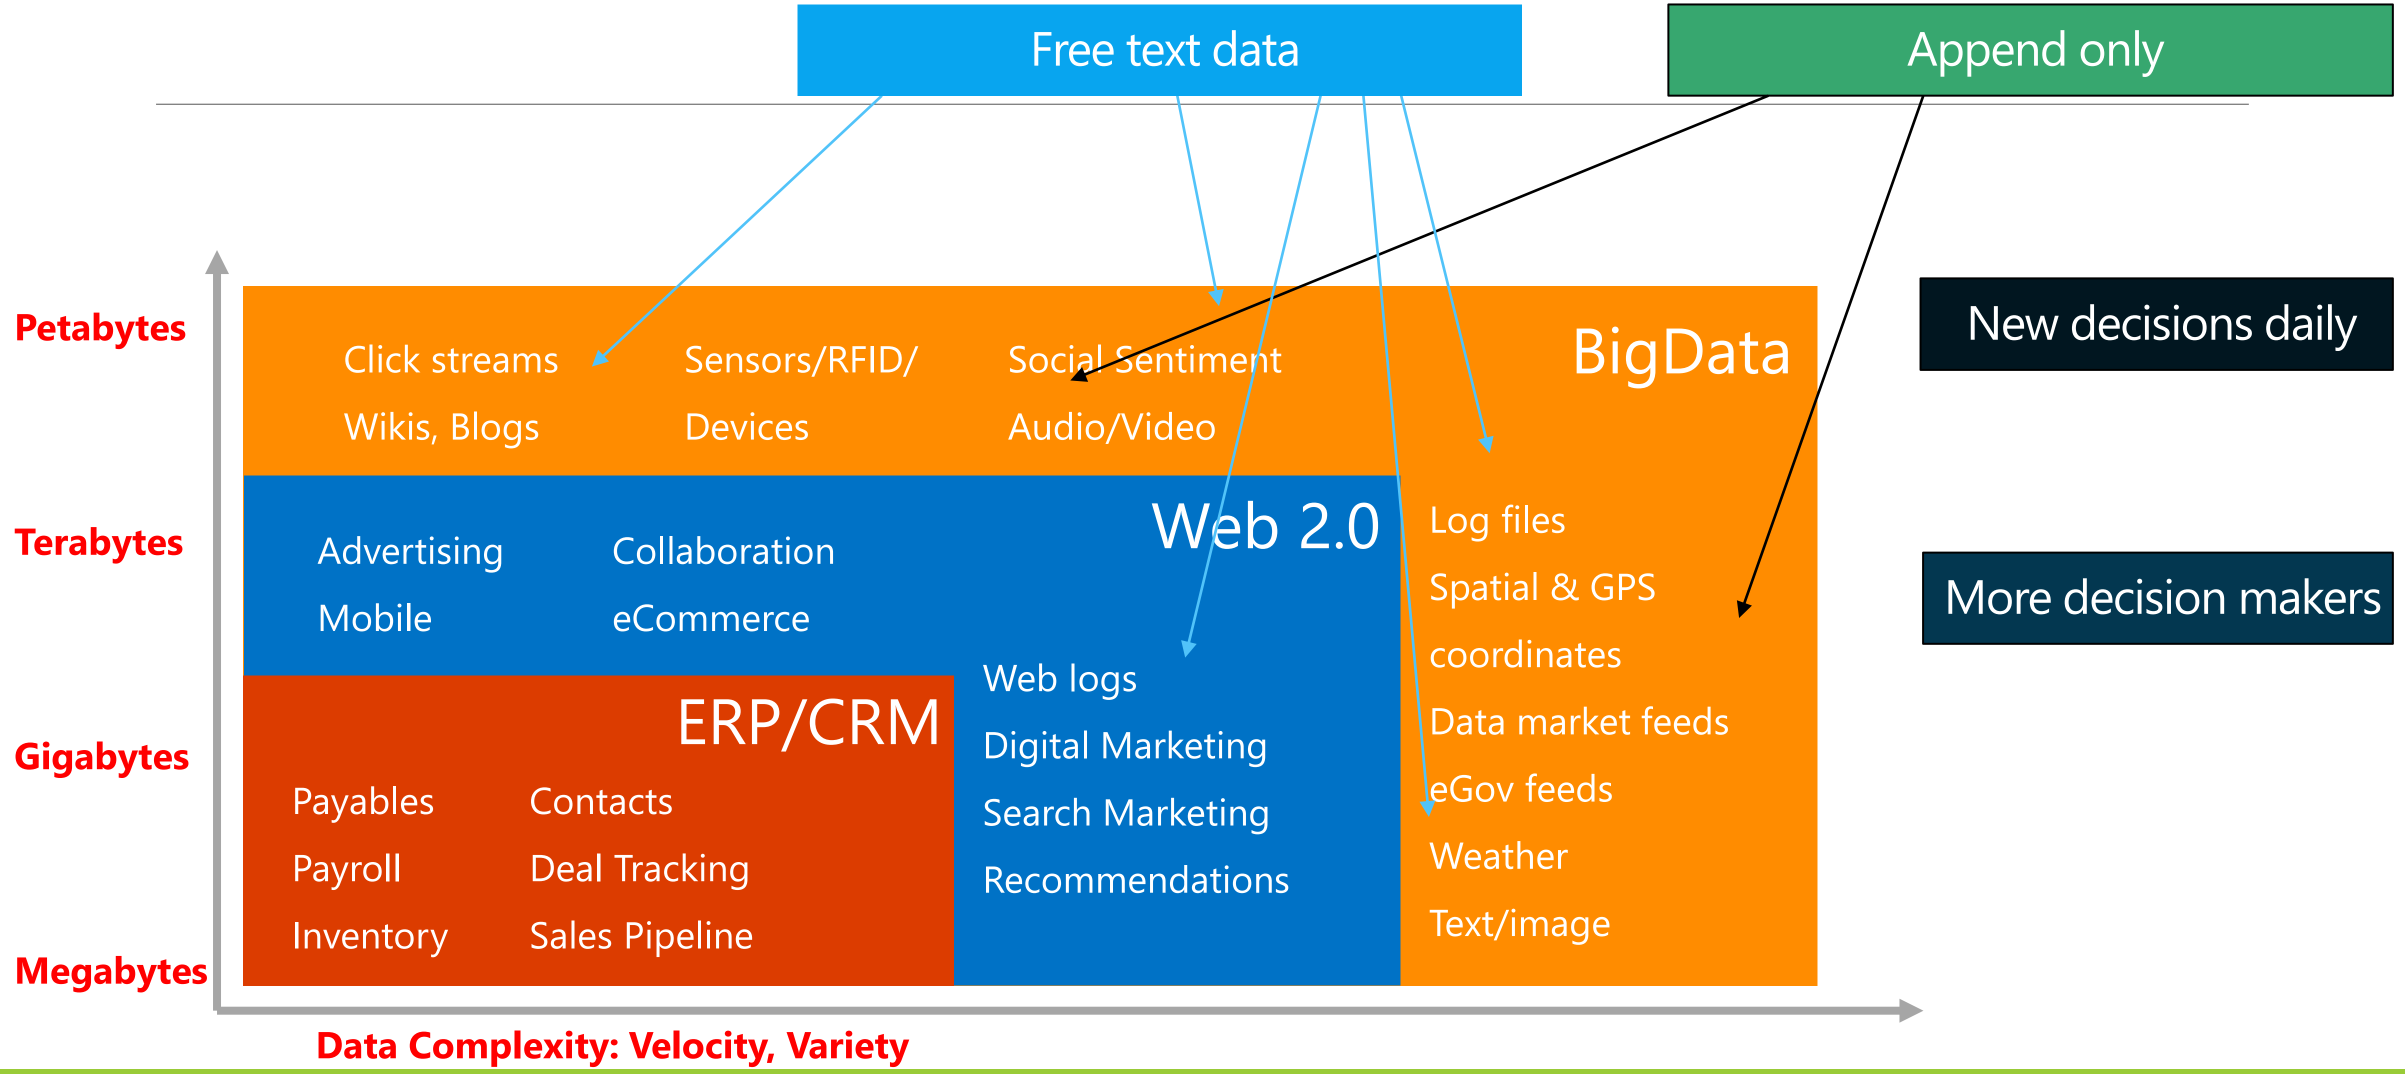

Demo

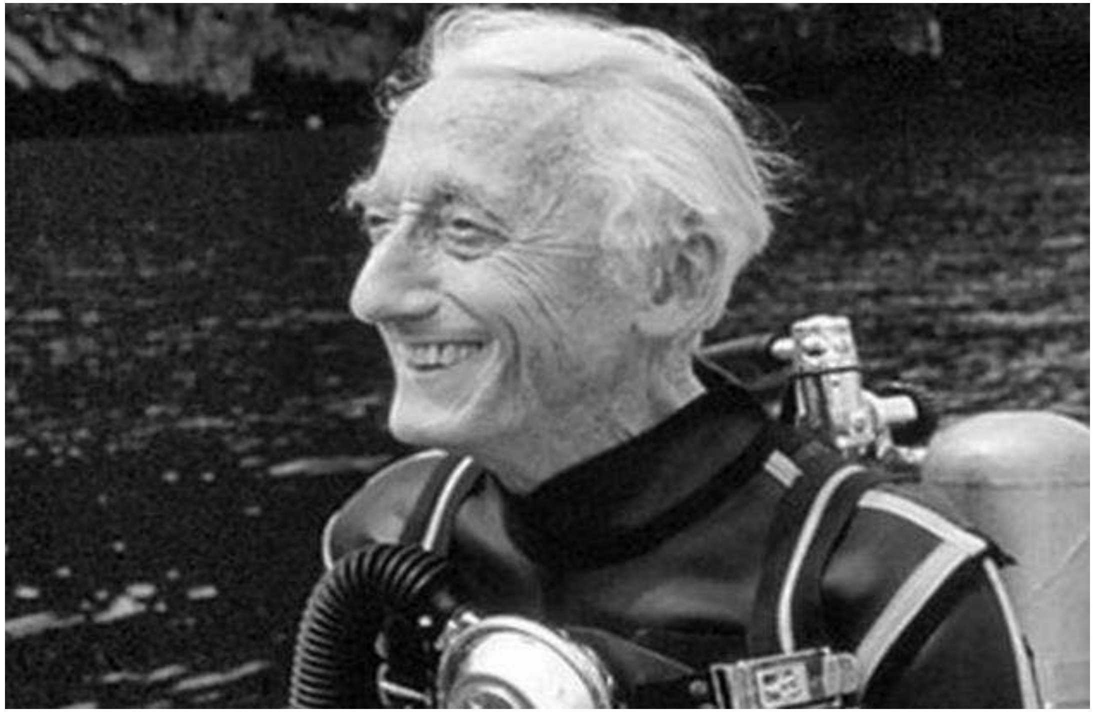

Azure Data Explorer
(aka Kusto)
2018-

Jacques-Yves Custeau
French Explorer
1910-1997

Azure Data Explorer – By the Numbers

51
Regions in Azure

1.9 EB
Total data size

16.3B
Total queries

1M
Cores

30 PB
Data ingested per day

35K
Explorer distinct users

3 Main scenarios

Pick the Best Tool for the Job

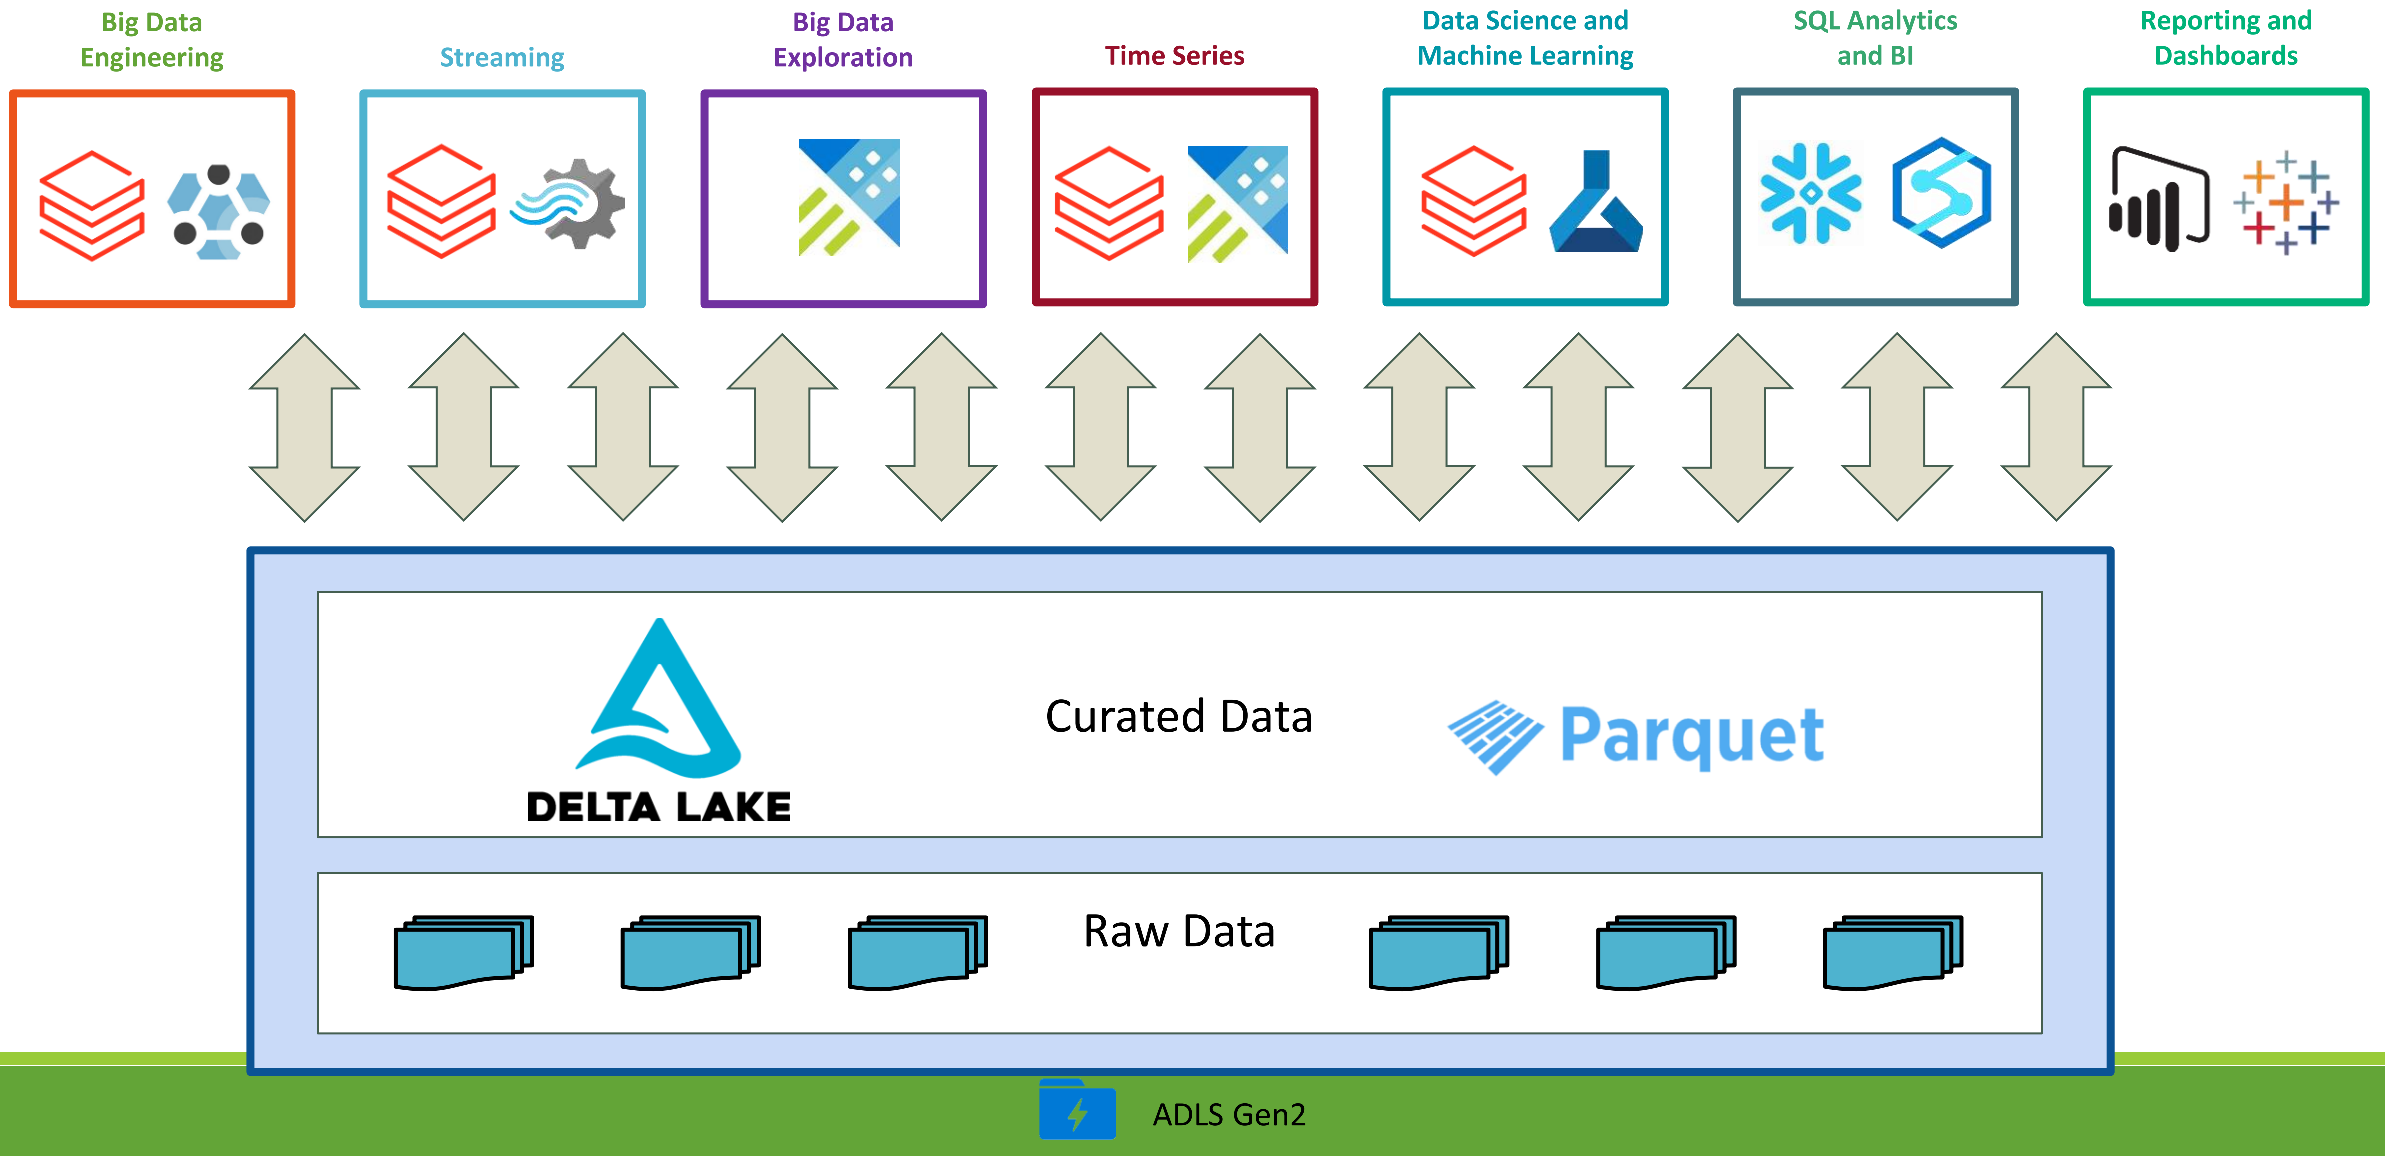

3 main scenarios

Data Exploration

Real Time Analytics

Fast moving data / time series

Why is it so fast?

ADX Deployment Model

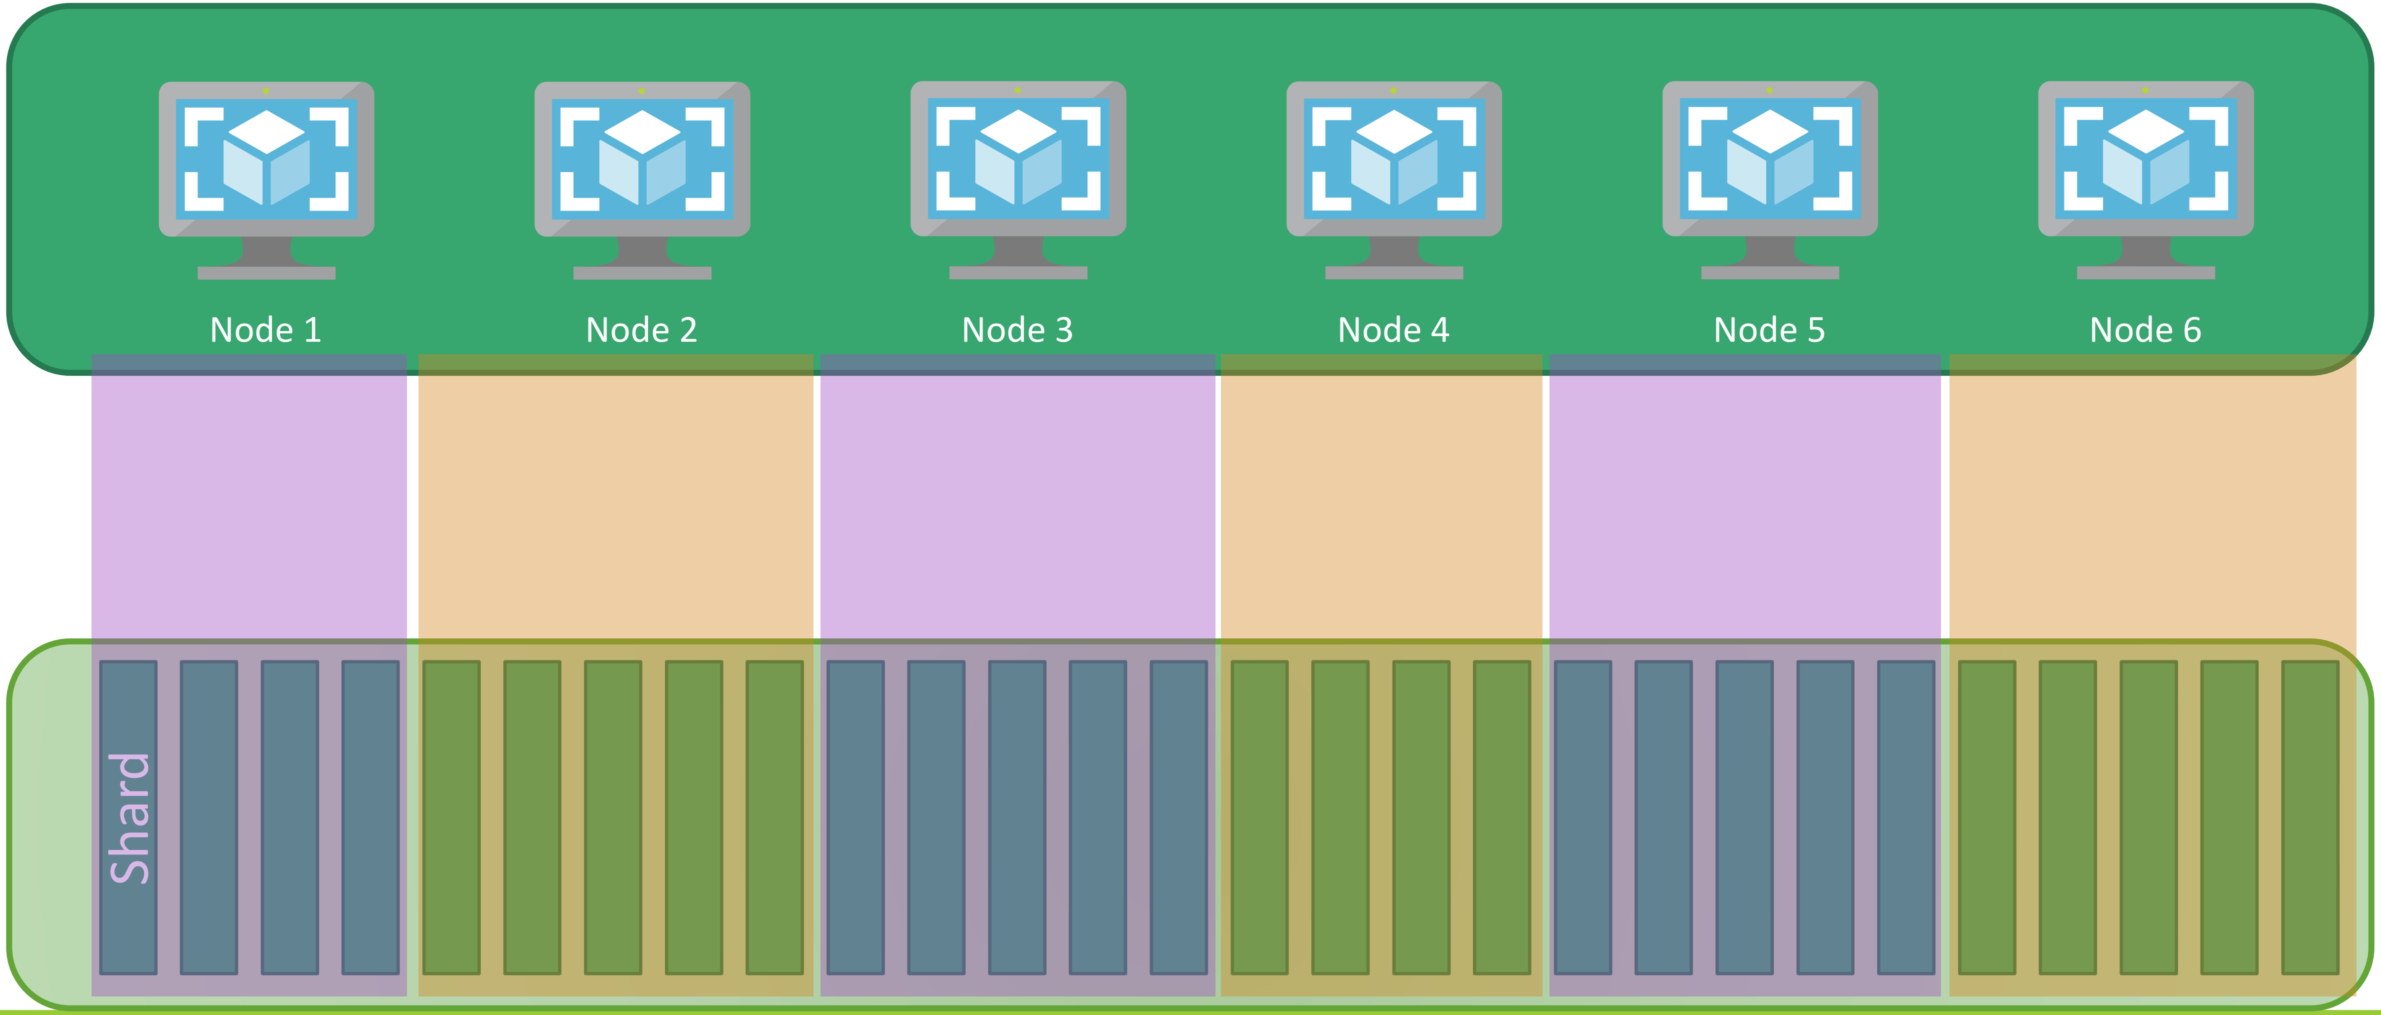

Storage Model

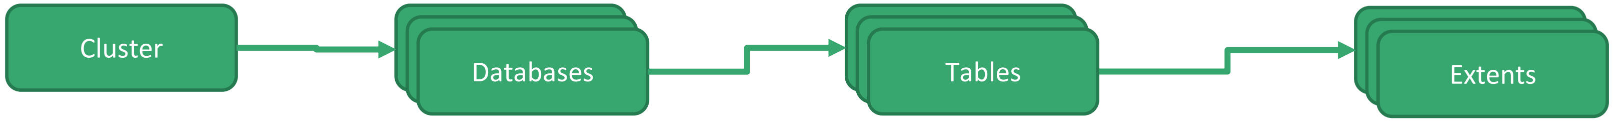

ADX cluster can have many databases

- Each database can have many tables

Table data is divided in **extents** (shards in Kusto Terminology)

An extent is

- Columnar & Compressed
- Fully Indexed
- Segmented
- Statistics (e.g. min, max, min / max size)
- Readonly / sealed (can be merged)

Extents are own by one-and-only-one node

Cached on node’s SSD / RAM (hot cache)

Extents are aligned with time, as data get ingested

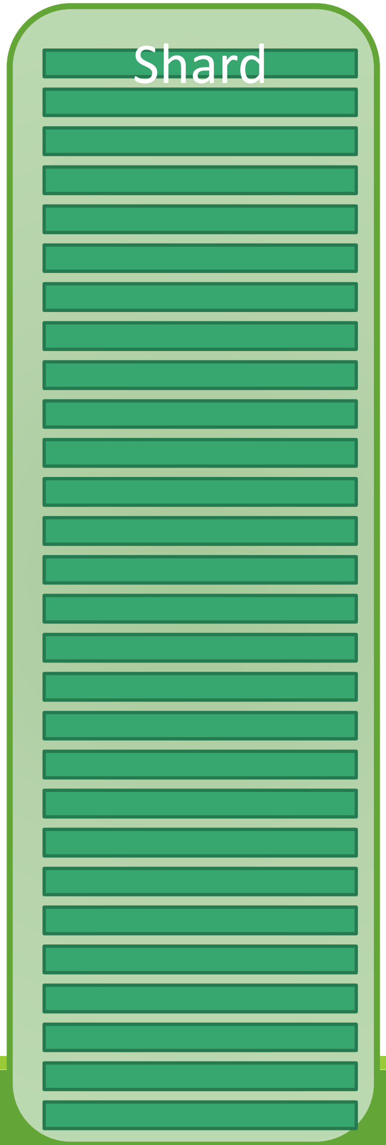

Real-time Analytics scenario

Azure Data Explorer	**Spark**
ADX ingest data in near real-time (seconds)	Spark ingests Data in near real-time (sub seconds)
Ingested data is indexed and can be queried ad hoc	Pre-determined aggregates are updated in near real-time
Real-time Analytics / Ad Hoc	Streaming / Real Time Transformation (ECP)

Not the same scenario

ADX Ingestion Model

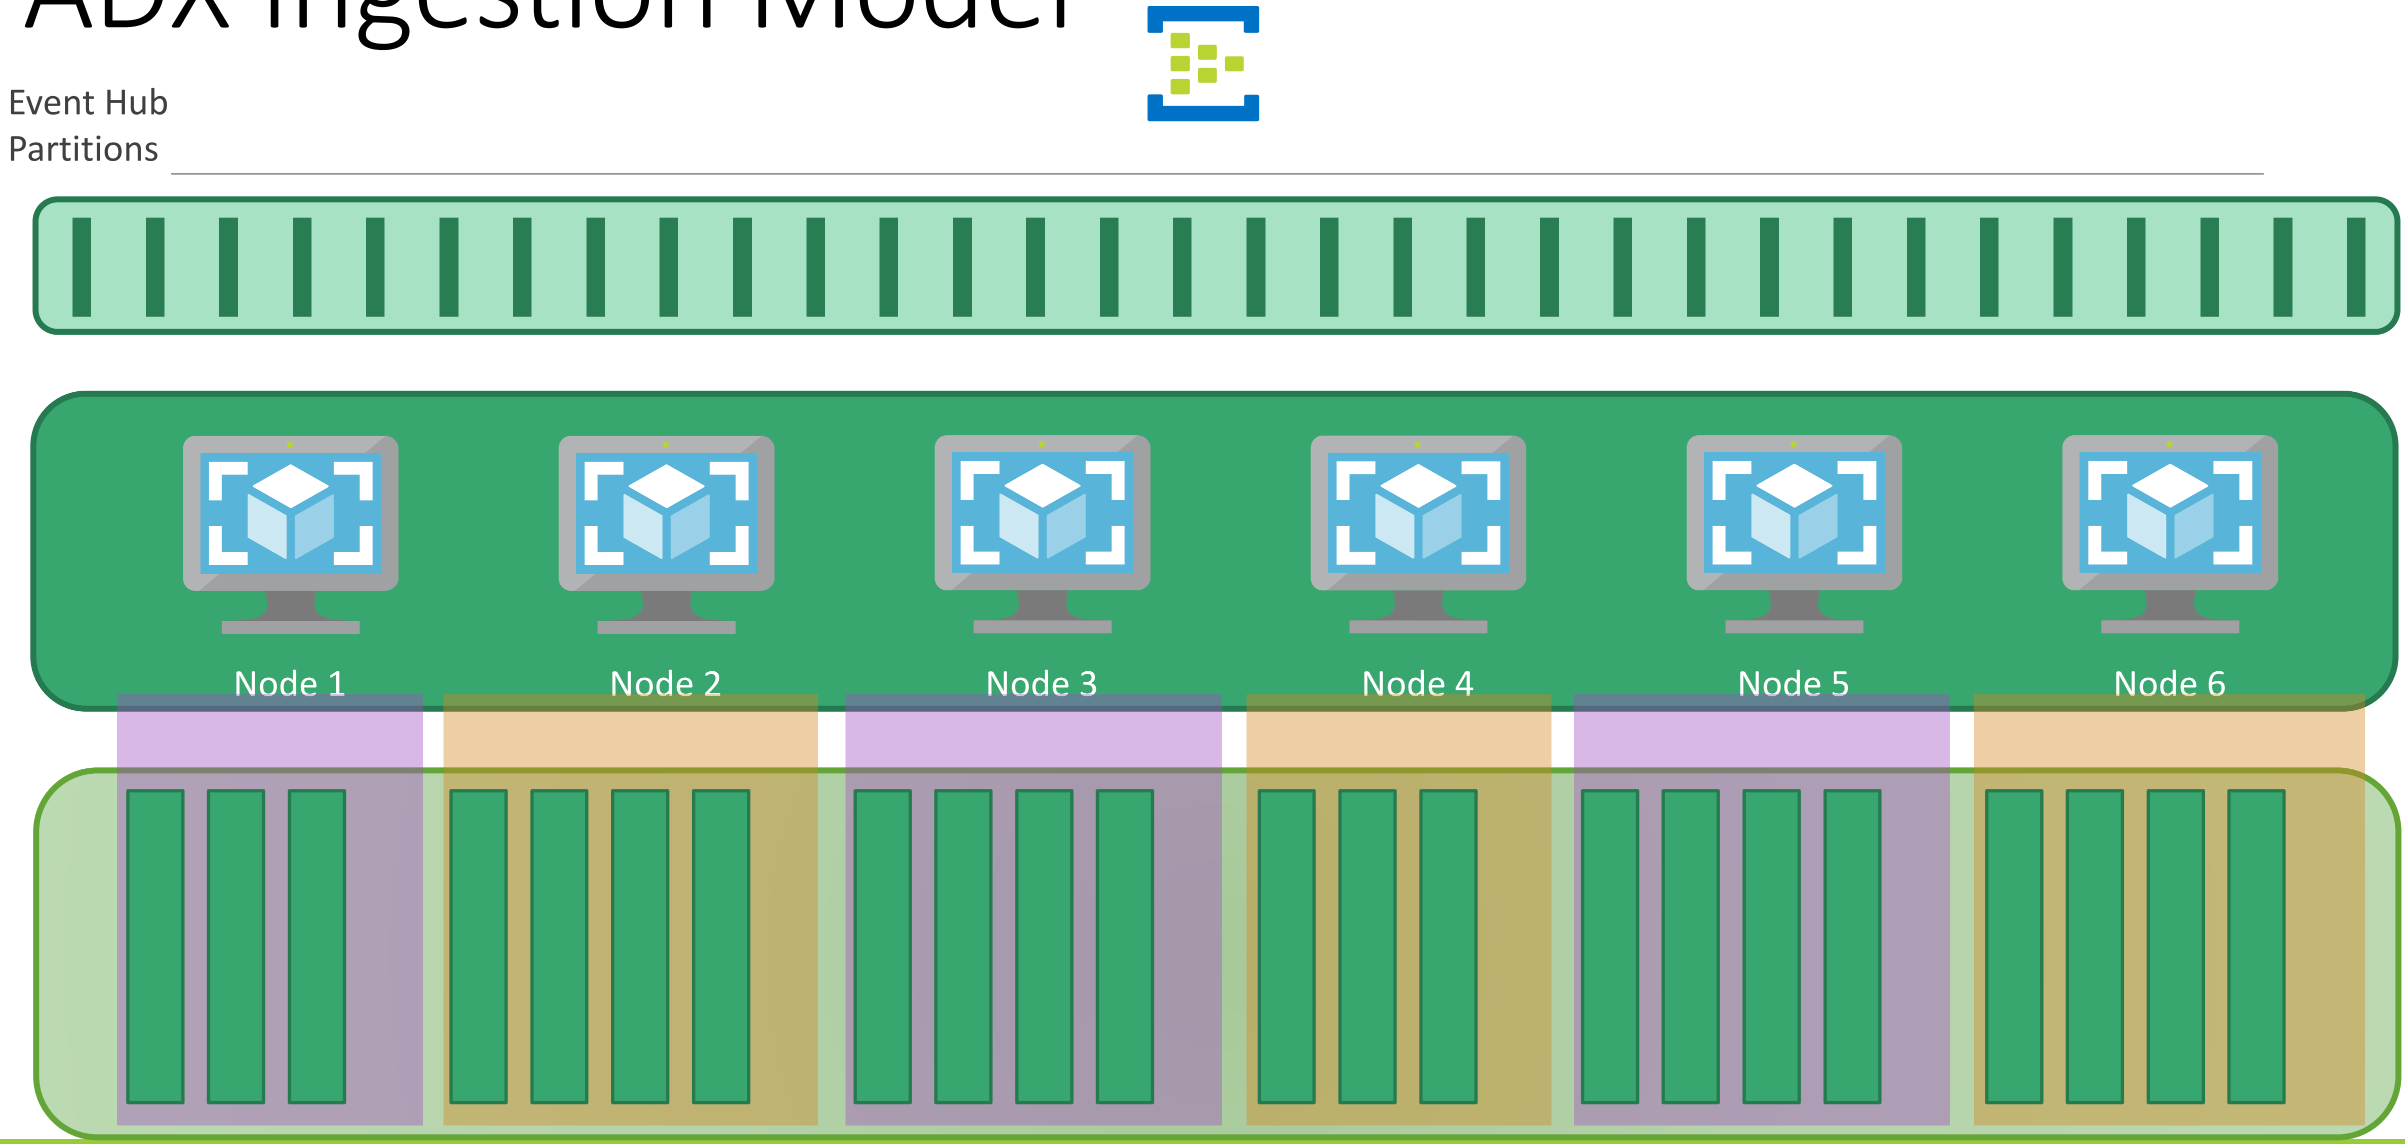

Observations (Real-time analytics)

Each node can ingest in parallel => True linear scaling

Even at a node level, there is no contention: shard is created in isolation

Only when the shard is “committed” is there coordination (i.e. serialization)

Typical data warehouse systems are transactional in nature

With ingestion of each row comes latches, version management, etc.

Observations (Time Series)

ADX has highly optimized built-in functions for Time Series

ADX has specific representation for Time Series allowing analysing multiple time series concurrently

Leverages first two scenarios

- Data is cached and indexed, hence access is fast
- Data can be available for query within a few seconds of creation

Analyzing data within time window is efficient (shard trimming)

The other side…
(ADX Shortcoming / when it shouldn’t be positioned)

Long running tasks

Data Engineering

- Very good at real time processing (update policies) not at batch (long running tasks)
- Probably wasteful to load batch data just to transform it

Can’t update / delete data

- Purge for GDPR scenarios (compute intensive)
- Extents can be swapped or deleted (bulk data movement)

No row level transaction (only at extent level)

Streaming

Training Machine Learning Model

- As with Azure Synapse, ADX can run models but we can’t train
- Does clustering with few algorithms & linear regressions

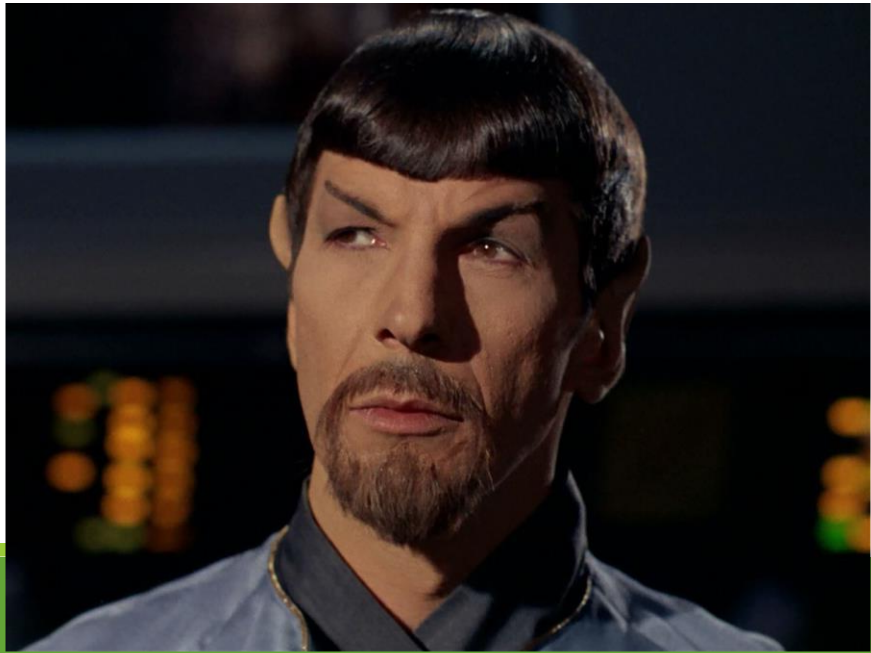

Positioning

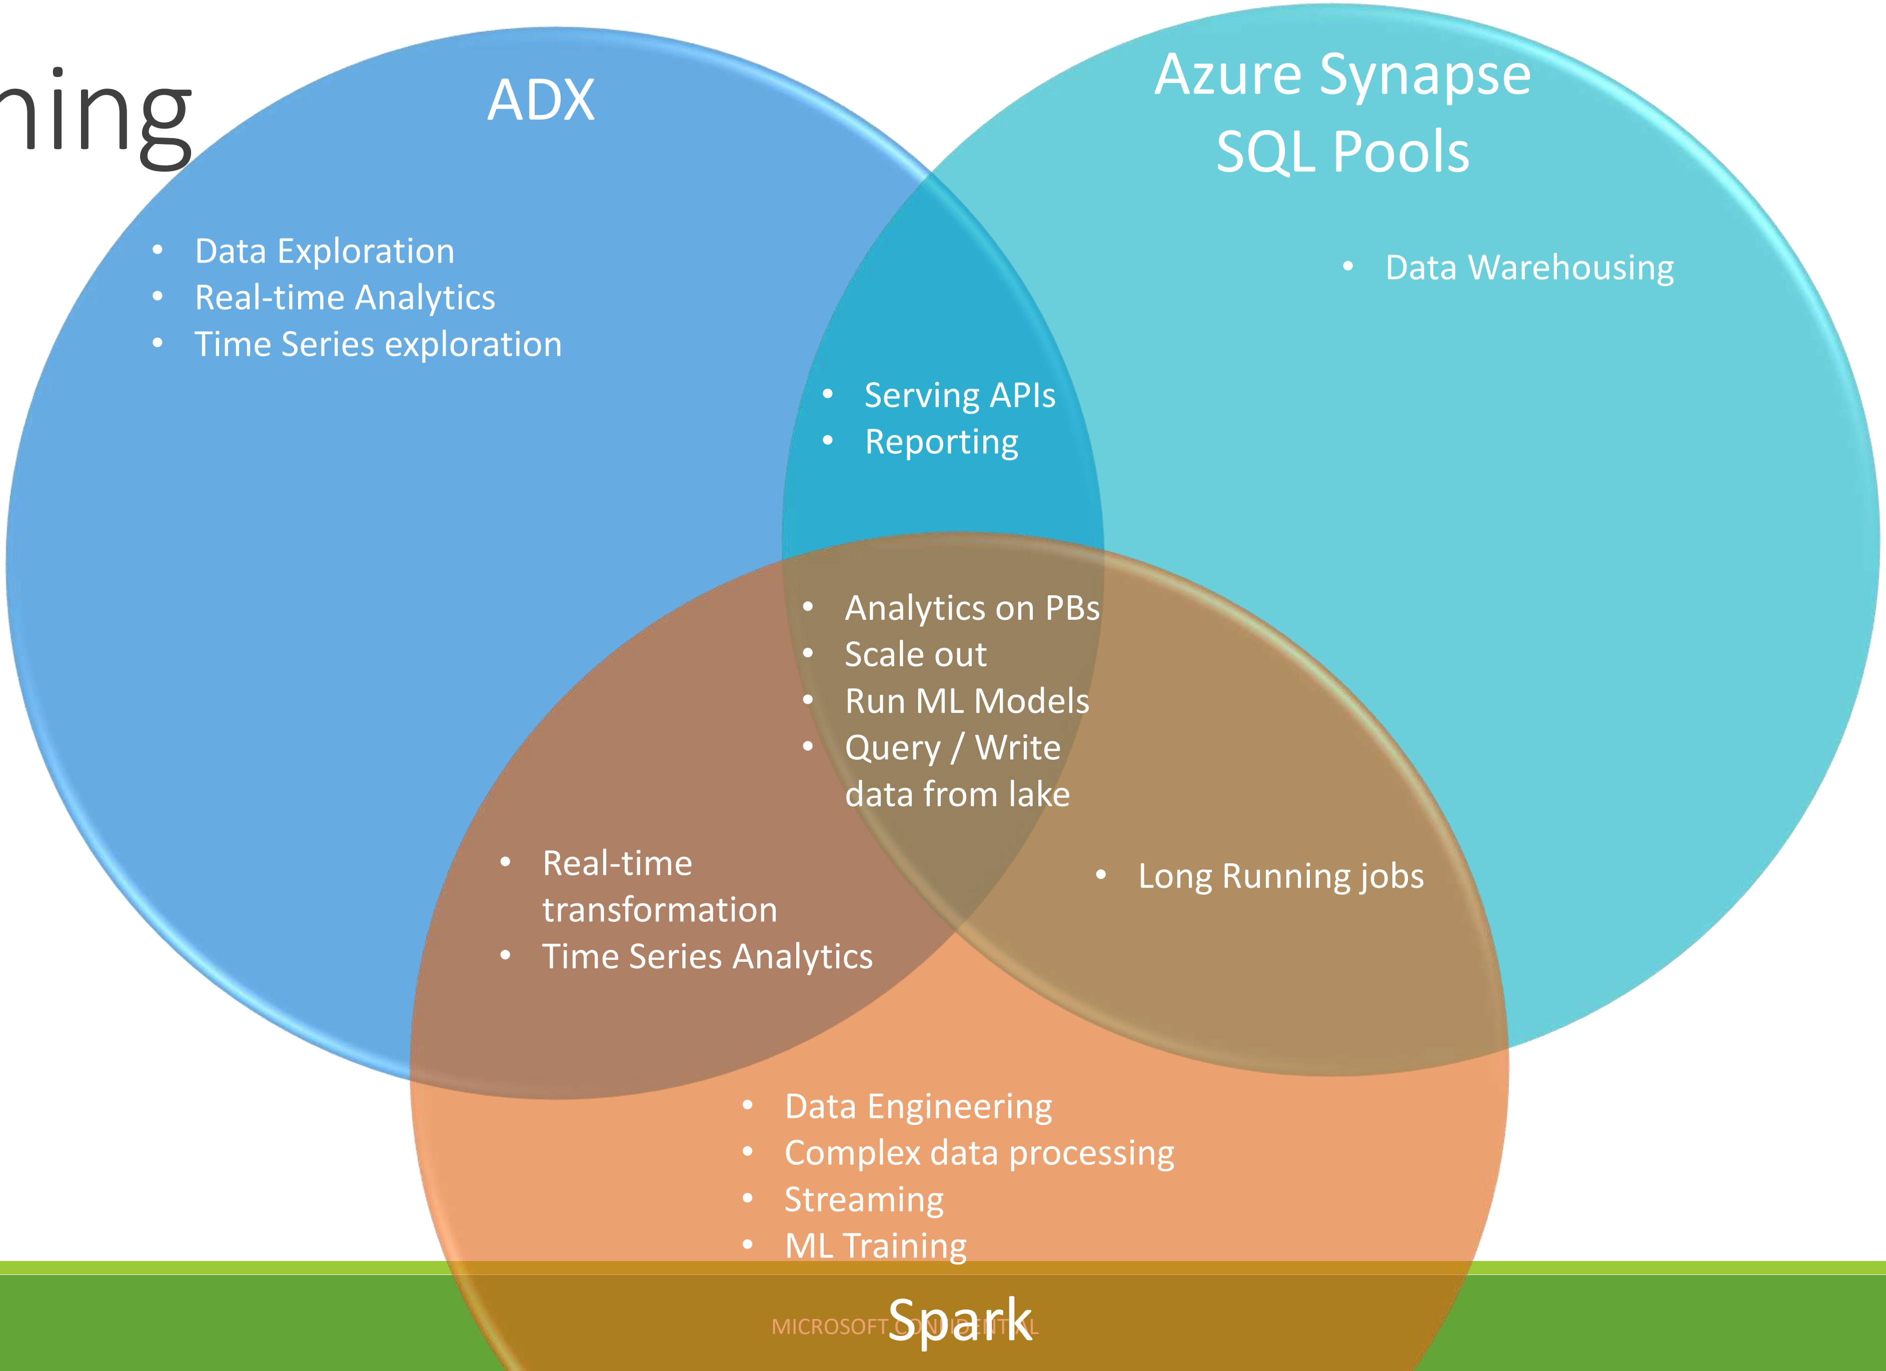

ADX

- Data Exploration
- Real-time Analytics
- Time Series exploration

Azure Synapse SQL Pools

- Data Warehousing

- Serving APIs
- Reporting

- Analytics on PBs
- Scale out
- Run ML Models
- Query / Write data from lake

- Real-time transformation
- Time Series Analytics

- Long Running jobs

- Data Engineering
- Complex data processing
- Streaming
- ML Training

Spark