

CETERIS AG



An Introduction to  
Azure Delta Lake using Databricks

## Competences

data and AI-solutions based on Microsoft technologies  
development and operation of cloud/on-prem platforms  
corporate information design (IBCS)

## Partners

itacs GmbH  
Microsoft Goldpartner  
Microsoft Power BI Partner  
Azure Marketplace und Microsoft AppSource Publisher  
Certified Cubeware Partner  
Certified Zebra BI Partner  
GAPTEQ Partner  
graphomate Partner

## Team

small but capable team of experienced data analytics consultants  
self-organized team without typical hierarchies  
agile project approach with close customer coupling



Gold Data Analytics  
Gold Data Platform  
Gold Cloud Platform



GAPTEQ



## Tarek Salha

- Senior Consultant, at Ceteris AG since 2015
- Msc. Physics
- Topics:
  - Data Warehousing
  - Advanced Analytics
  - Data Lake Architectures
  - Definitely no specialist for visualization



## Thorsten Huss

- Msc. Business Informatics
- Started at Ceteris AG in 2013 as student employee, now Senior Consultant
- Topics:
  - Data Integration
  - ...but pretty much everything ETL, really.



Gold Data Analytics  
Gold Data Platform  
Gold Cloud Platform



GAPTEQ



CETERIS AG

**What is ...?**

5

CETERIS AG

**Delta Lake 101**

12

CETERIS AG

**Traveling in time**

18

CETERIS AG

**Creating and writing streams**

20

CETERIS AG

**Power BI Visualization  
on Delta Tables**

23

CETERIS AG

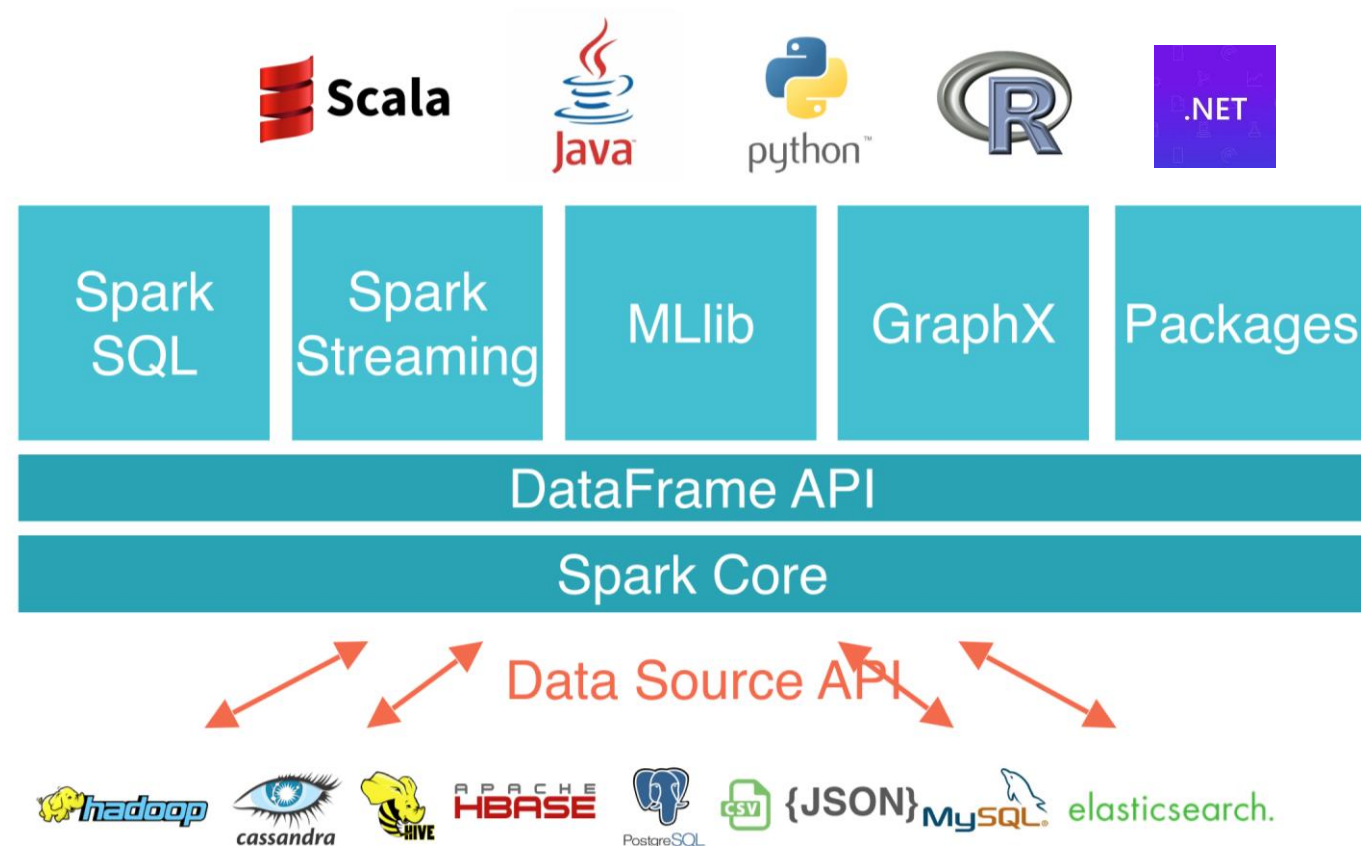
**Streaming DWH Demo**

26

**What is ...?**

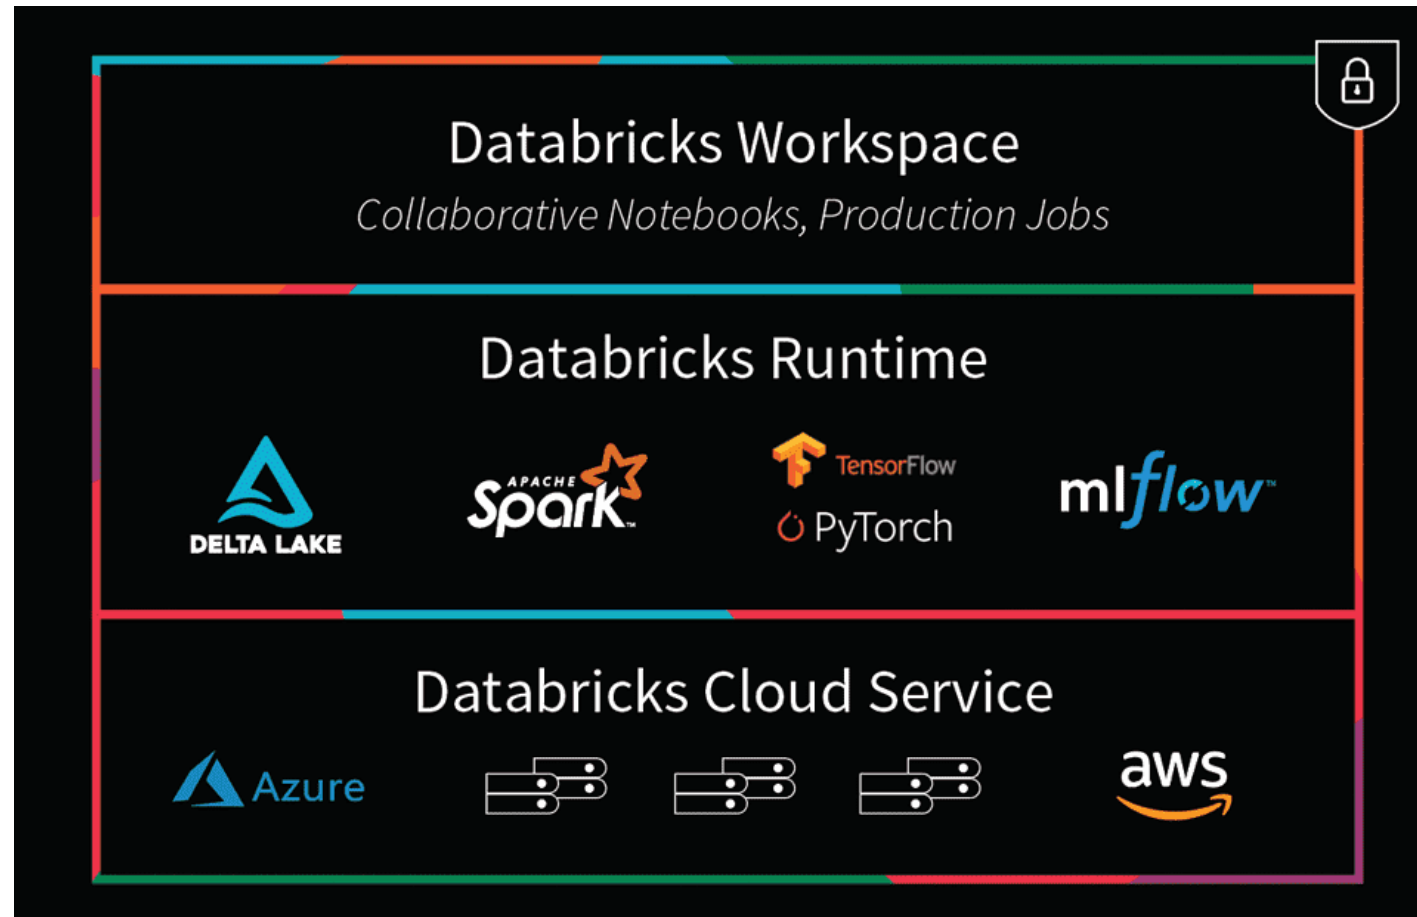
# What is Apache Spark?

- Apache Spark is an analytics software framework, that combines cluster data processing and AI
- One of the most actively developed open source big data projects

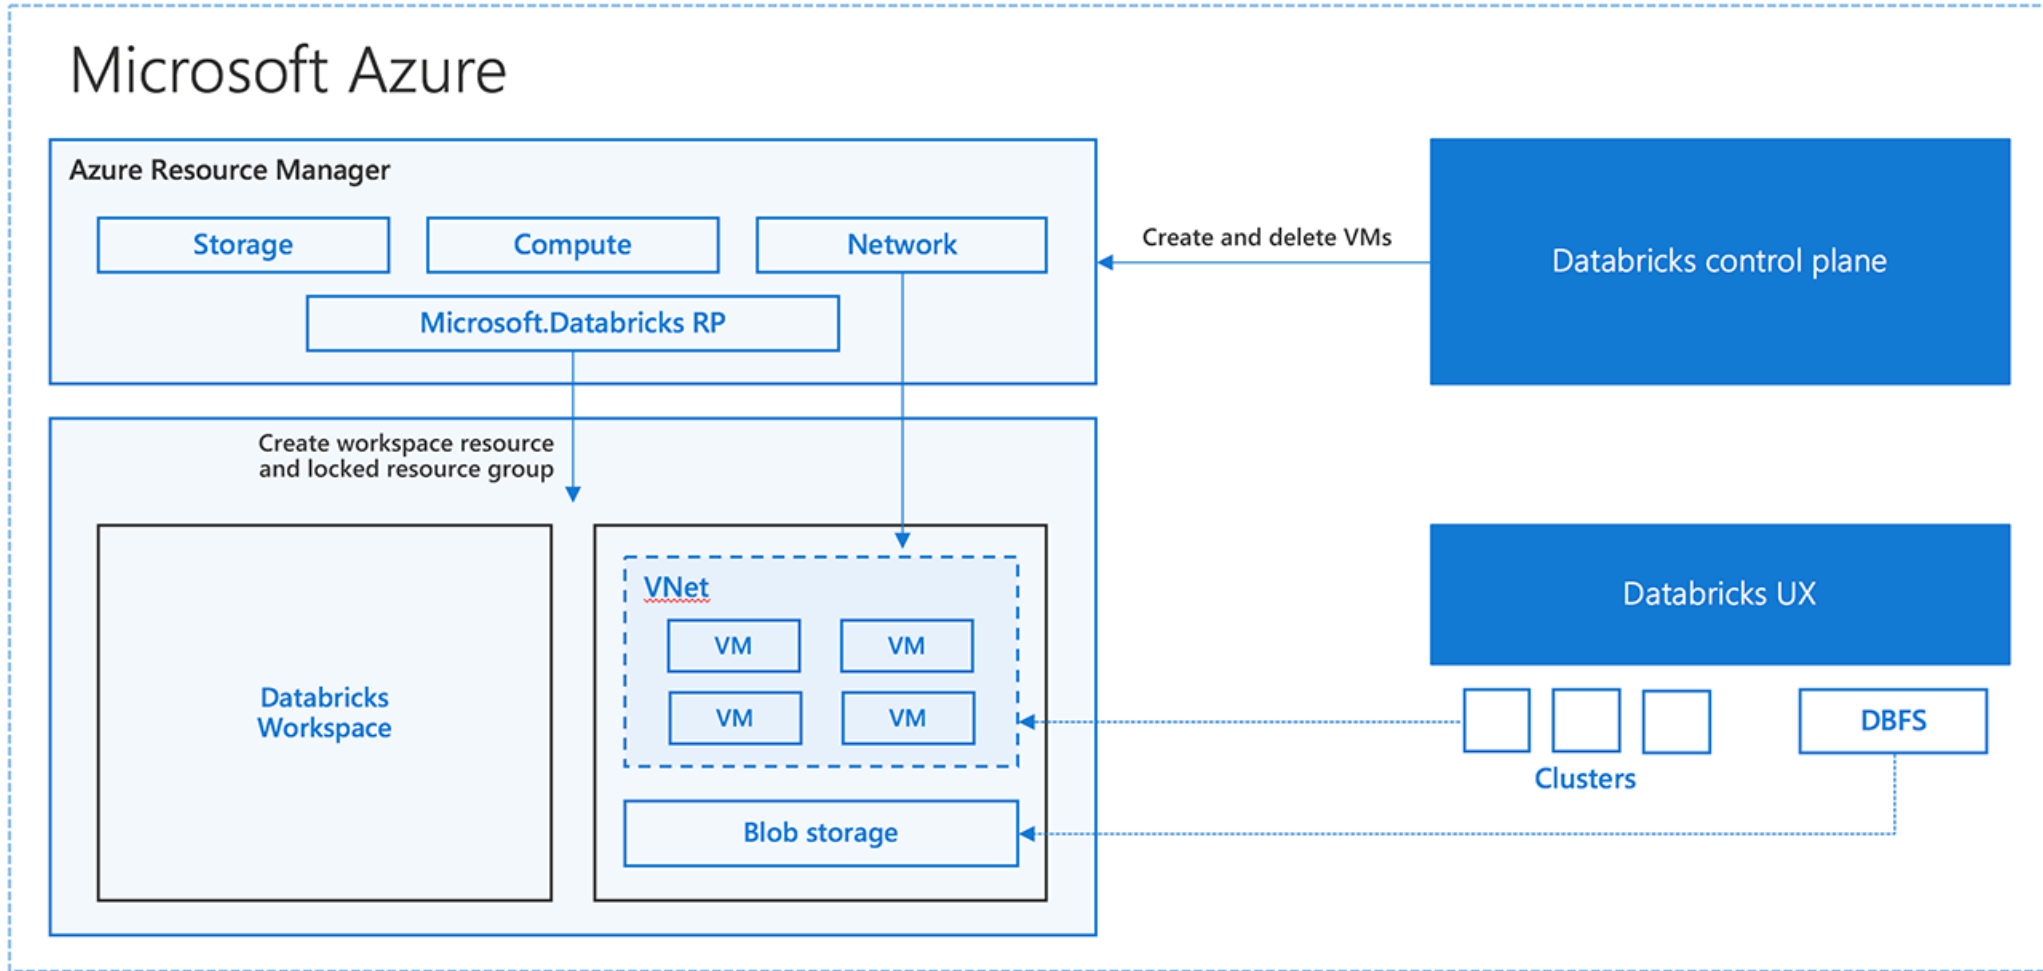


# What is Databricks?

- Databricks is a company (original creators of Apache Spark)
- They offer a fast, easy and secure PaaS service to perform Spark operations



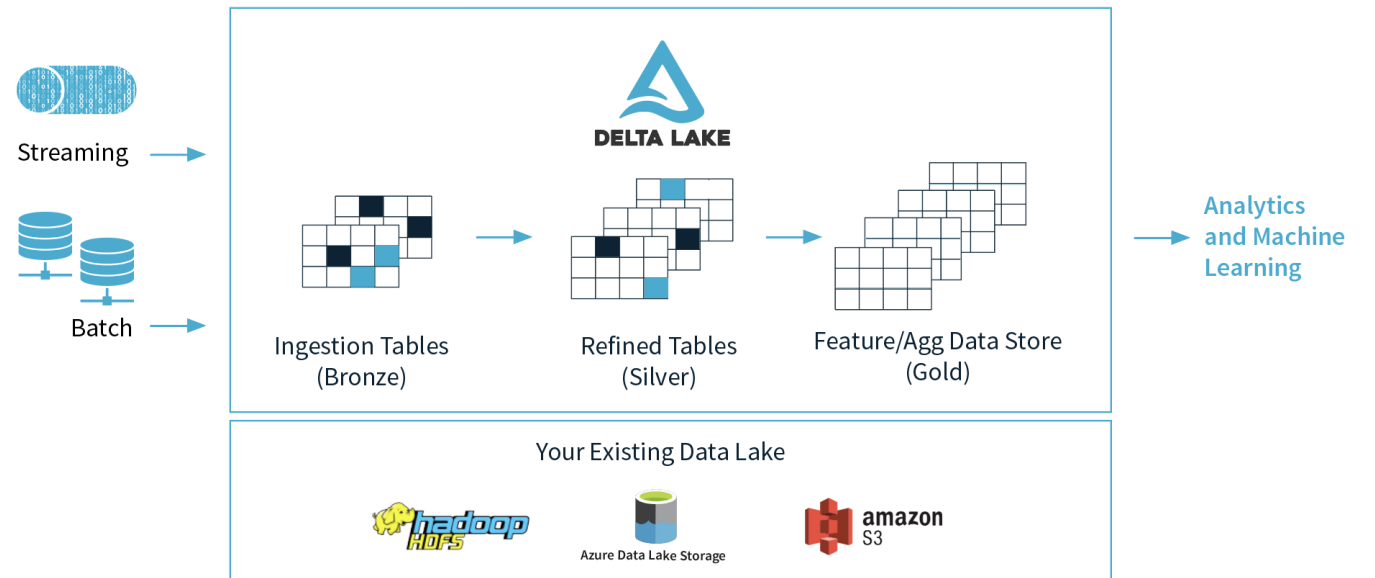
# How is Databricks working in Azure?





# What is Delta Lake?

- Delta Lake is an open-source storage layer that brings ACID transactions and other relational database features to Apache Spark (on top of it).
- It provides:
  - ACID transactions
  - Time travel
  - Open-source storage format
  - Streaming sources and sinks
  - Schema enforcement as well as evolution
  - Audit History
  - Update / delete commands

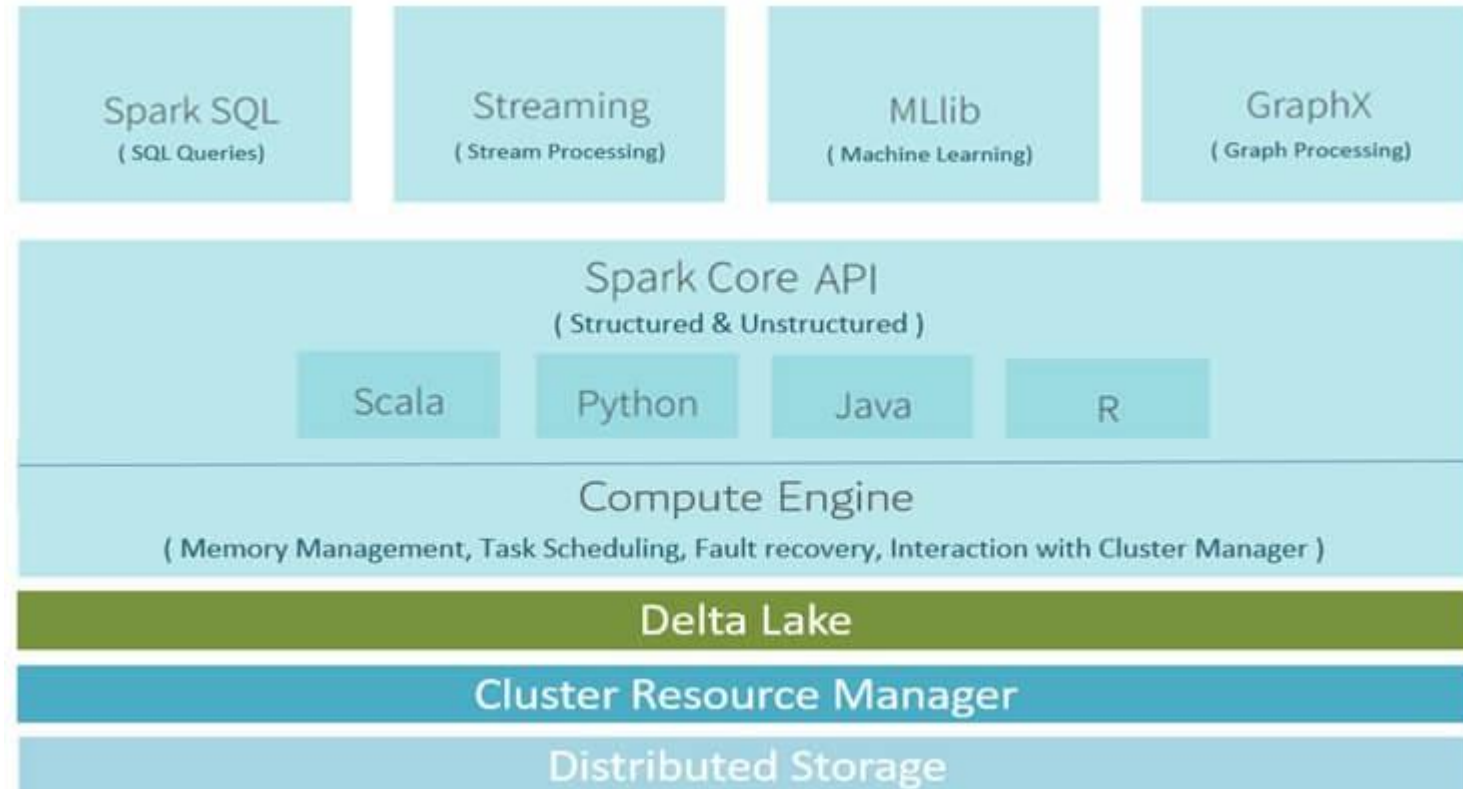


# How does Delta Lake ACID principle work?

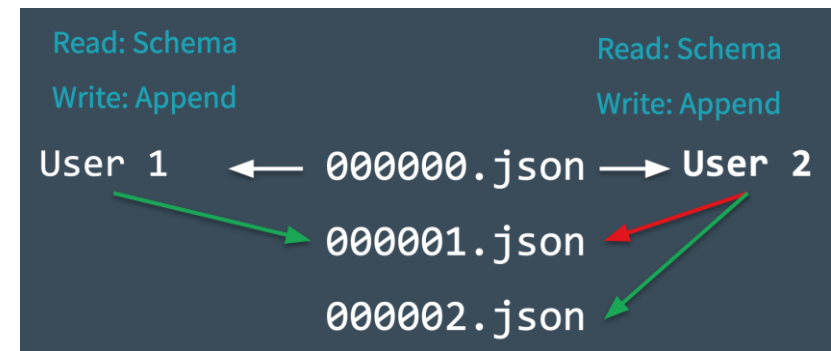
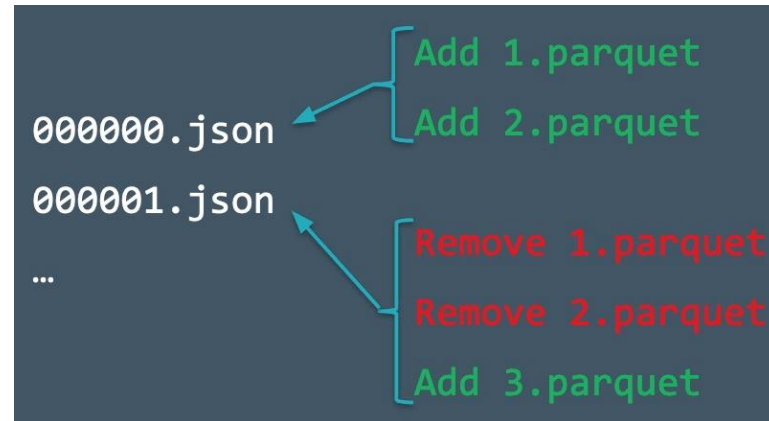
- Delta Lake guarantees atomicity and consistency via the so-called transaction log

*"If it's not recorded in the transaction log, it never happened."*

- It provides serializability as level of isolation
- Durability is automatically conserved due to all information being written directly to disk



# Transaction Log



# Delta Lake 101

Tables are just references and metadata

```
CREATE TABLE events (  
  date DATE,  
  eventId STRING,  
  eventType STRING,  
  data STRING)  
USING DELTA
```

```
df.write.format("delta").saveAsTable("events")
```

```
INSERT INTO events SELECT * FROM newEvents
```

```
INSERT OVERWRITE events SELECT * FROM newEvents
```

```
df.write  
  .format("delta")  
  .mode("overwrite")  
  .option("replaceWhere", "date >= '2017-01-01' AND date <= '2017-01-31'")  
  .save("/mnt/delta/events")
```

```
import io.delta.tables._

val deltaTable = DeltaTable.forPath(spark, "/data/events/")

deltaTable.updateExpr(          // predicate and update expressions using SQL for
  "eventType = 'clck'",
  Map("eventType" -> "'click'")

import org.apache.spark.sql.functions._
import spark.implicits._

deltaTable.update(              // predicate using Spark SQL functions and implicits
  col("eventType") === "clck",
  Map("eventType" -> lit("click")));
```

```
MERGE INTO events
USING updates
ON events.eventId = updates.eventId
WHEN MATCHED THEN
  UPDATE SET events.data = updates.data
WHEN NOT MATCHED
  THEN INSERT (date, eventId, data) VALUES (date, eventId, data)
```

```
DELETE FROM events WHERE date < '2017-01-01'
```

```
DELETE FROM delta.`/data/events/` WHERE date < '2017-01-01'
```

```
import io.delta.tables._

val deltaTable = DeltaTable.forPath(spark, "/data/events/")

deltaTable.delete("date < '2017-01-01'")

import org.apache.spark.sql.functions._
import spark.implicits._

deltaTable.delete(col("date") < "2017-01-01")
```



```
DROP TABLE IF EXISTS <example-table> // deletes the metadata and data  
CREATE TABLE <example-table> AS SELECT ...
```

```
DROP TABLE IF EXISTS <example-table> // deletes the metadata  
dbutils.fs.rm("<your-s3-path>", true) // deletes the data  
CREATE TABLE <example-table> USING org.apache.spark.sql.parquet
```

# Traveling in time

# „Time traveling? As if...” – „AS OF”!

- Go back to the state of a table at a specific timestamp or table version
- Scala/Python: `spark.read(...).option(„timestampAsOf”, “2020-07-02”).load(„myPath”)`
- SQL: `SELECT * FROM myTable VERSION AS OF 1`
  - View table versions and audit information with `DESCRIBE HISTORY` (or just use the UI)
- Use Cases: Rollbacks, time series analytics, pinned views,...



(Unfortunately, you can really just go to the past and back to the future, aka the present)

# Creating and writing streams

- Get a Databricks cluster up and running (and add any configs and libraries before you start it up)
- Before you stream anything to delta, configure your Gen2 storage and a mounting point
- Think about creating „external“ tables (i.e. not managed by Databricks) beforehand
- Prepare source configuration
  - File names/locations
  - EventHub endpoint
  - SQL Server jdbc drivers
  - ...

# ... and write it to a delta table

- Basic scala syntax:

```
insertDF.writeStream
  .format("delta")
  .outputMode("append")
  .option("checkpointLocation", "/mnt/MountFolder/TableName/_checkpoints/etl-from-json")
  .start("/mnt/MountFolder/TableName") //or table(„TableName„)
  .trigger(Trigger.Once)
  .partitionBy(„PartitionColumn1“,„PartitionColumn2“)
```

- Output options:

- append – default, appends rows to existing or newly created table
- complete – replace the entire table
- update – only writes rows that have changed since last trigger (only used with aggregations)

- Trigger options:

- Trigger.Once – triggers exactly once and then stops the stream (in theory...)
- Trigger.ProcessingTime("60 seconds") – triggers in given interval (can be anything from ms to days)
- Default: behaves as if ProcessingTime set to 0 ms, tries to fire queries as fast as possible

# Power BI Visualization on Delta Tables

# How to connect to Databricks?

1. Get a personal access token
2. Get your cluster's server hostname, port, and HTTP path
3. Construct the server address to use in in Power BI Desktop
  - a. Use the schema https://
  - b. Append the server hostname after the schema
  - c. Append the HTTP path after the server host name

→ <https://westeurope.azuredatabricks.net/sql/protocolv1/o/0123456789/0123-456789-sometext>

4. In Power BI use Spark connector and use
  - a. „token“ as username
  - b. personal access token as password

It supports Import AND DirectQuery models!


The screenshot shows the 'Account Settings' page in Databricks, specifically the 'Access Tokens' tab. It displays a table of generated tokens with columns for Token ID, Comment, Creation, and Expiration. Below the table, there are tabs for 'Spark', 'Tags', 'Logging', 'Init Scripts', 'JDBC/ODBC', and 'Permissions'. The 'JDBC/ODBC' tab is active, showing fields for 'Server Hostname', 'Port', 'Protocol', 'HTTP Path', and 'JDBC URL'. The 'JDBC URL' field contains a complex string: `jdbc:spark://[hostname]:443/default;transportMode=http;ssl=1;httpPath=sql/protocolv1/o/7064161269814046/0801-[token];AuthMech=3;UID=token;PWD=<personal-access-token>`

The screenshot shows the 'Spark' connector dialog box in Power BI. It displays the URL `https://[hostname]:443/sqp/protoc...` and the 'User name' field set to 'token'. The 'Password' field is masked with dots. There are 'Back', 'Connect', and 'Cancel' buttons at the bottom.



# How to connect to Databricks?

---



# Streaming DWH Demo

# What this demo will show (if there's enough time)

---

- Streaming from ~~EventHub~~ storage all the way to Synapse
- Joins in streams
- Watermarking
- How to actually write to tables in Synapse and why we ended up needing a classic blob storage
- Including user-defined functions
- Handling late-arriving data and SCD2

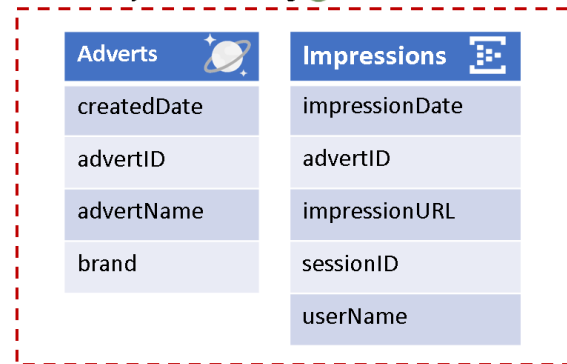
# Scenario and source

- Demo available by following instructions on blog by Nicholas Hurt:

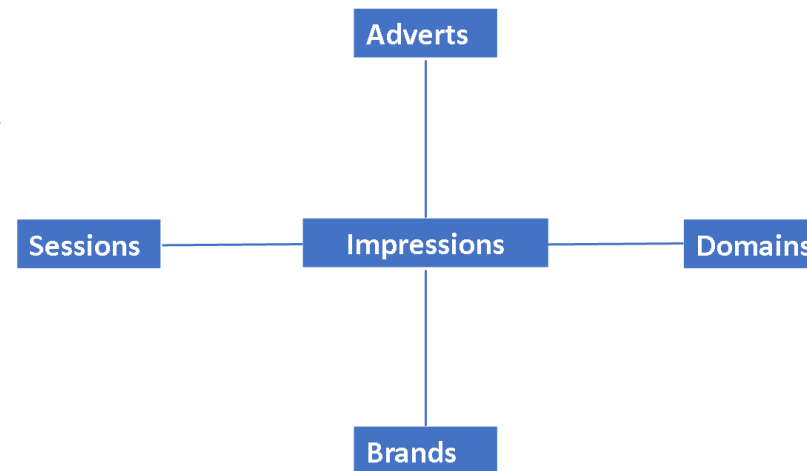



<https://medium.com/microsoftazure/an-introduction-to-streaming-etl-on-azure-databricks-using-structured-streaming-databricks-16b369d77e34>

- ...though of course we had to simplify and change it up a bit

Continuous feed - streaming 🌱



Slowly changing data - batch



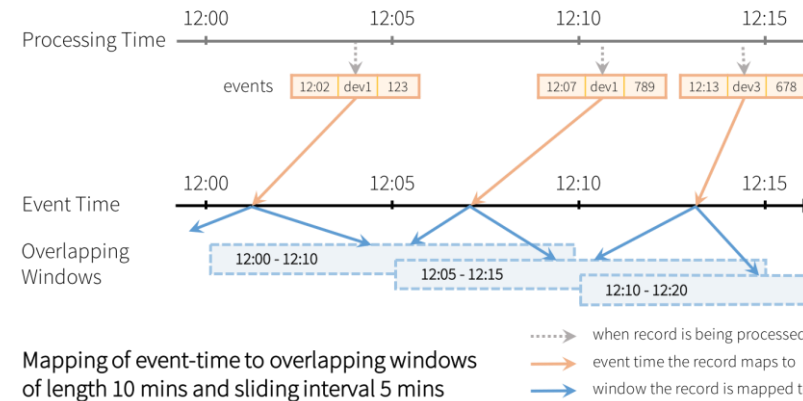
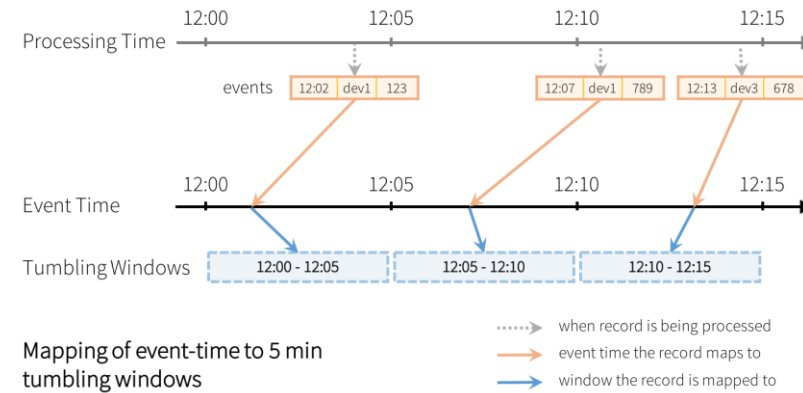
## ■ Define windows for aggregations

- Tumbling window

```
windowedAvgSignalDF = \  
  eventsDF \  
    .groupBy(window("eventTime", "5 minutes")) \  
    .count()
```

- overlapping window

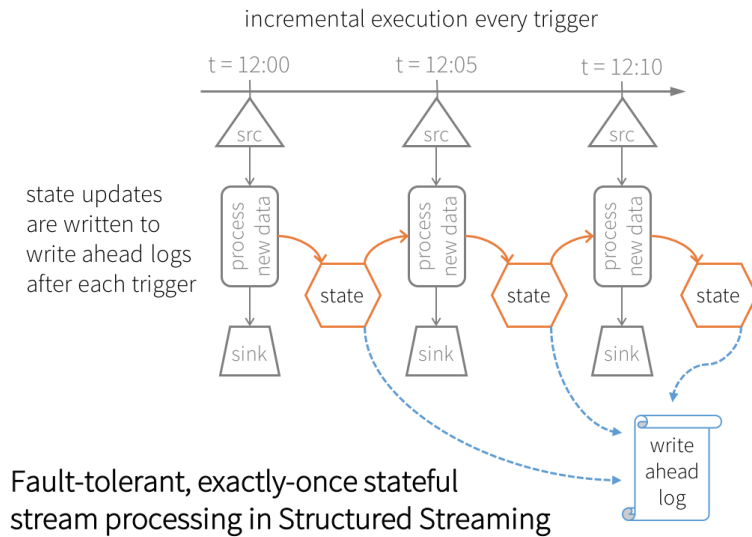
```
windowedAvgSignalDF = \  
  eventsDF \  
    .groupBy(window("eventTime", "10 minutes", "5 minutes")) \  
    .count()
```



Source:


<https://databricks.com/blog/2017/05/08/event-time-aggregation-watermarking-apache-sparks-structured-streaming.html>

## States:



## How to prevent inevitable memory overflow?

- `.withWatermark(„eventTime“, „10 minutes“)`



Source:

<https://databricks.com/blog/2017/05/08/event-time-aggregation-watermarking-apache-sparks-structured-streaming.html>

- Due to data continuously streaming, joins have to be reimagined a bit
- For every row of table a, there could always arrive a row in table b that matches
- For left joins, you cannot ever tell if there will be no match and result has to be NULL
- Solution: use watermarks to determine how long the join operation waits for matches!
  - Define watermark delays on inputs
  - Define event-time range conditions on join operator

```
// Define watermarks
val impressionsWithWatermark = impressions
  .select($"adId".as("impressionAdId"), $"impressionTime")
  .withWatermark("impressionTime", "10 seconds ") // max 10 seconds late

val clicksWithWatermark = clicks
  .select($"adId".as("clickAdId"), $"clickTime")
  .withWatermark("clickTime", "20 seconds") // max 20 seconds late

// Inner join with time range conditions
display(
  impressionsWithWatermark.join(
    clicksWithWatermark,
    expr("""
      clickAdId = impressionAdId AND
      clickTime >= impressionTime AND
      clickTime <= impressionTime + interval 1 minutes
      """)
  )
)
```

# Questions?

---

[thus@ceteris.ag](mailto:thus@ceteris.ag)

[tsalha@ceteris.ag](mailto:tsalha@ceteris.ag)

LinkedIn:

<https://www.linkedin.com/in/thorstenhuss/>

<https://www.linkedin.com/in/tarek-salha-2a39ab189/>

Web: [www.ceteris.ag](http://www.ceteris.ag)

