

# Microsoft Fabric

## What's the fuzz about?

17-10-2024 | Webinar | S. van den Beemt



WEBINAR

## MS Fabric – What's the fuzz about

- De webinar wordt opgenomen
- Slides en opname worden achteraf gedeeld via de mail
- Q&A bewaren we tot het eind
- Inzoomen?
- Graag je microfoon uitgeschakeld houden
- Camera's aan: optioneel, liefst wel tijdens Q&A
- Eet smakelijk!



# Hello world!



**Steven van den Beemt**  
*Cloud Architect | ISO MCT*  
*@ICT Group*



# Agenda



- └ Introduction
- └ MS Fabric / OneLake architecture
- └ MS Fabric security
- └ Direct Lake

# Introduction

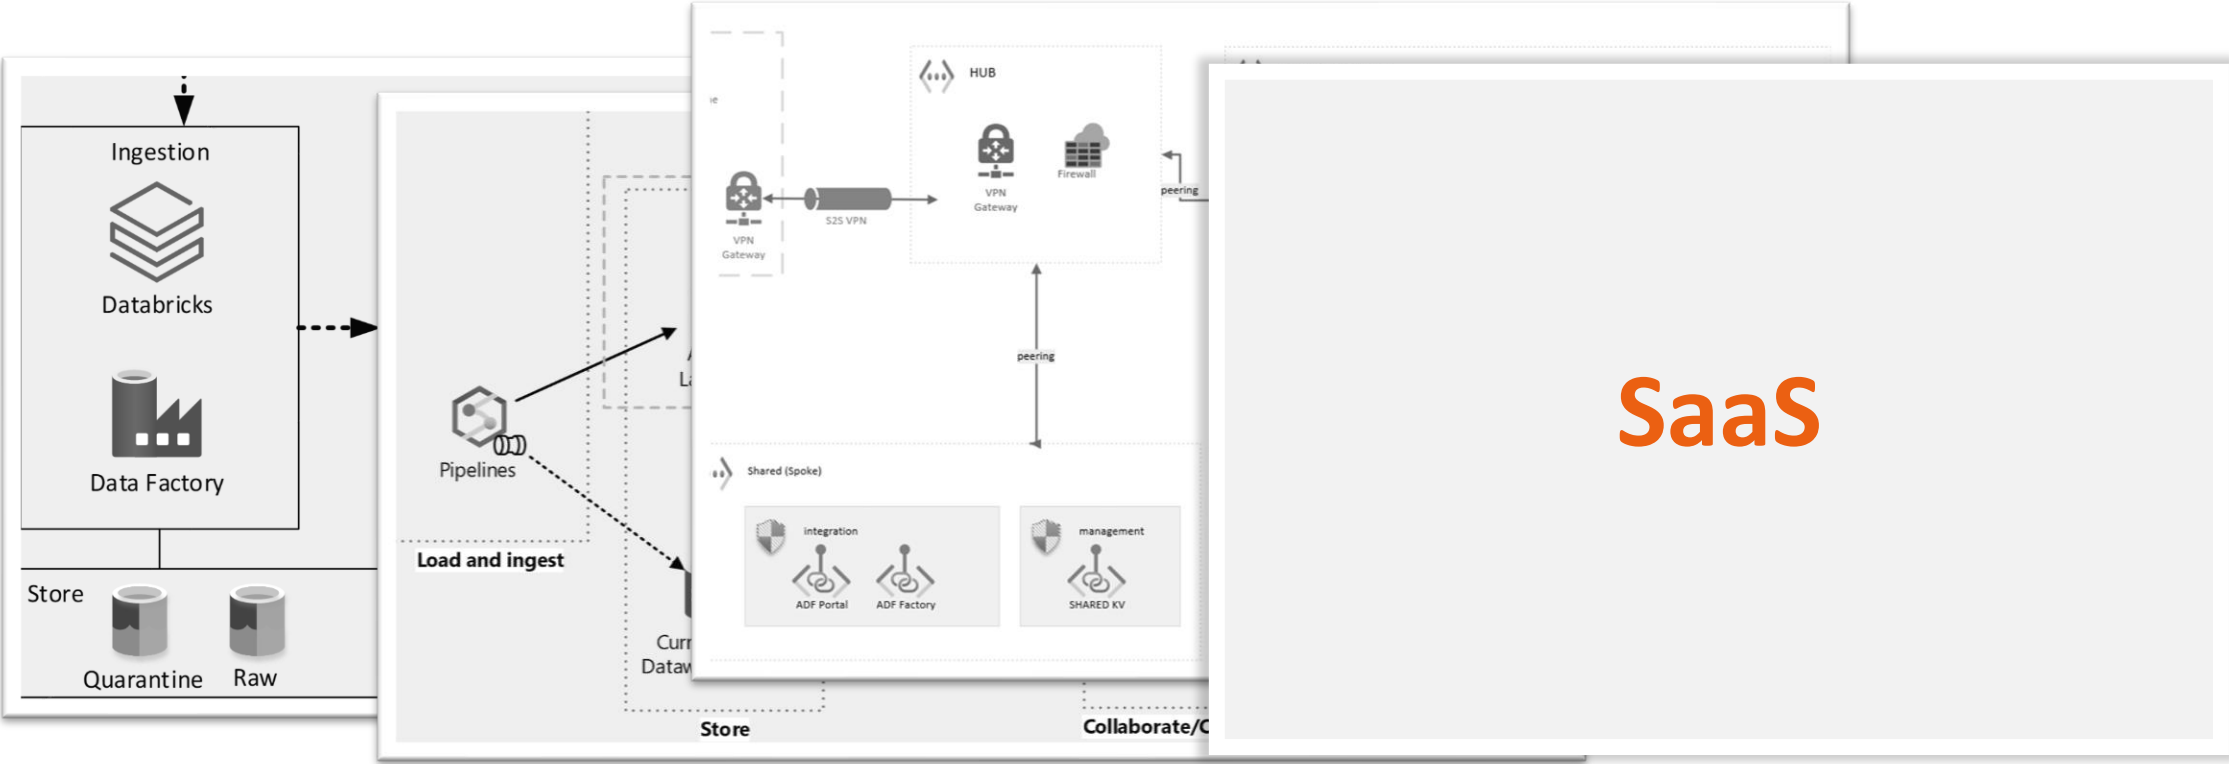
# Introduction

2014

2015

2020

2023




# Introduction

2014




2015



2020



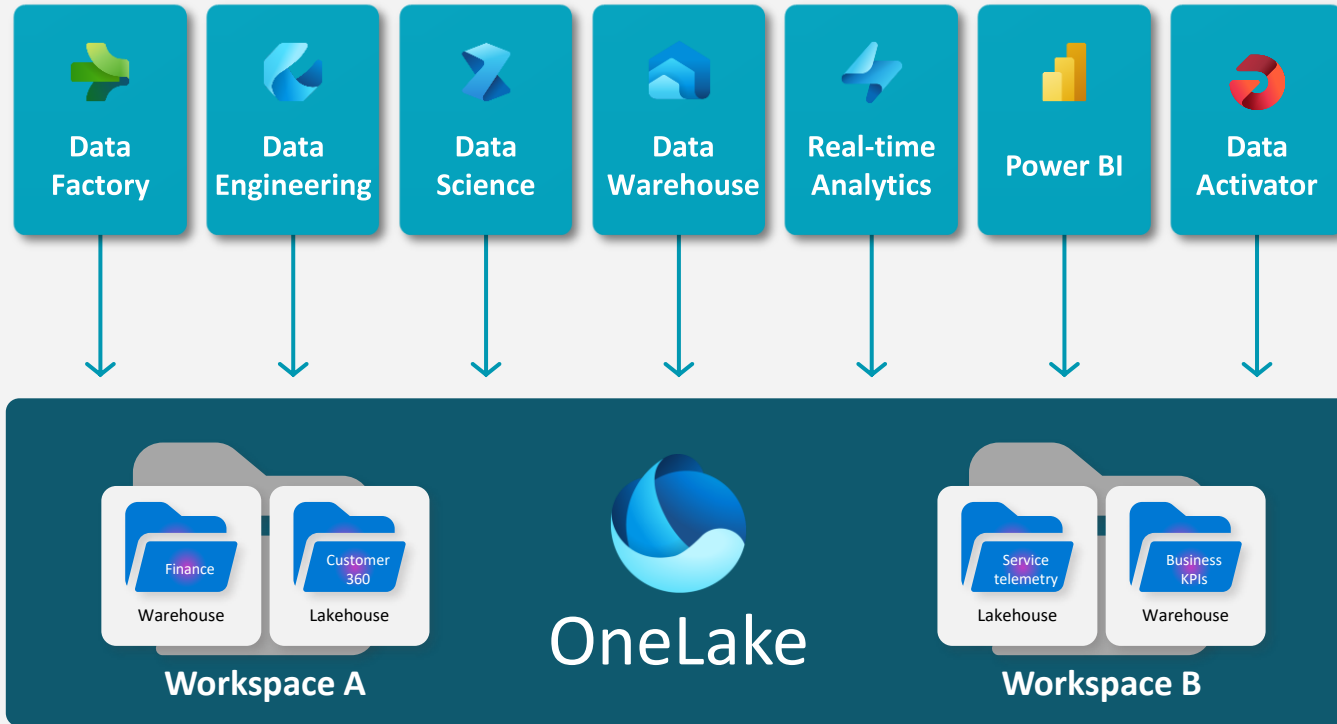
2023



# MS Fabric Architecture



# Understand Fabric architecture



Intelligent data foundation

## Single SaaS lake

Provisioned automatically with the tenant.

All workloads automatically store data in OneLake workspace folders.

All data is organized in a hierarchical namespace.

Data in OneLake is automatically indexed for discovery, MIP labels, lineage, PII scans, sharing, governance and compliance.



T-SQL

Spark

KQL

Analysis Services

**Lakehouse 1**

- Files
- Tables

**Warehouse 1**

- Compute
- Storage

domain

**Workspace 1**

**Workspace 2**

**Workspace 3**

Capacity

Capacity

**OneLake**

**Delta Log (meta data)**

**Delta Parquet files**

**ADLS2**

**A P I**



Classification: Training (R5)

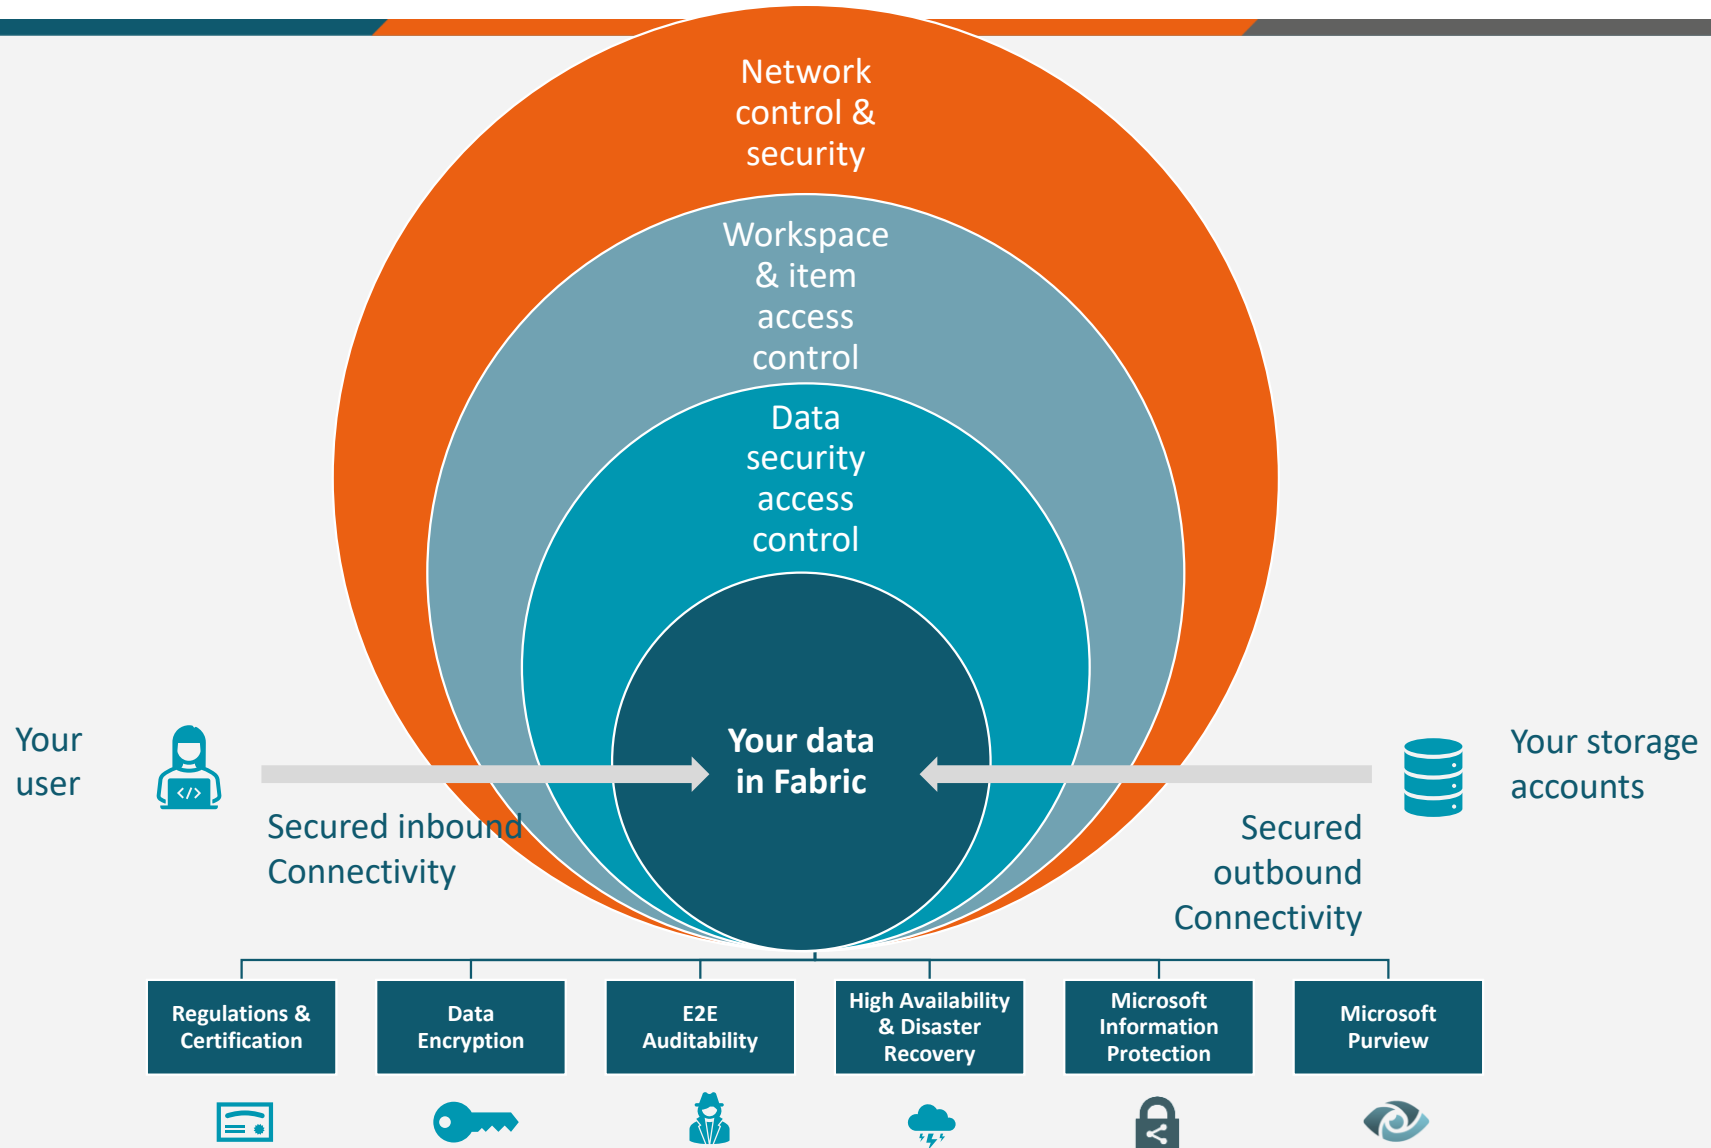
1 : 1



# Demo

# Security in MS Fabric

# Manage Fabric security

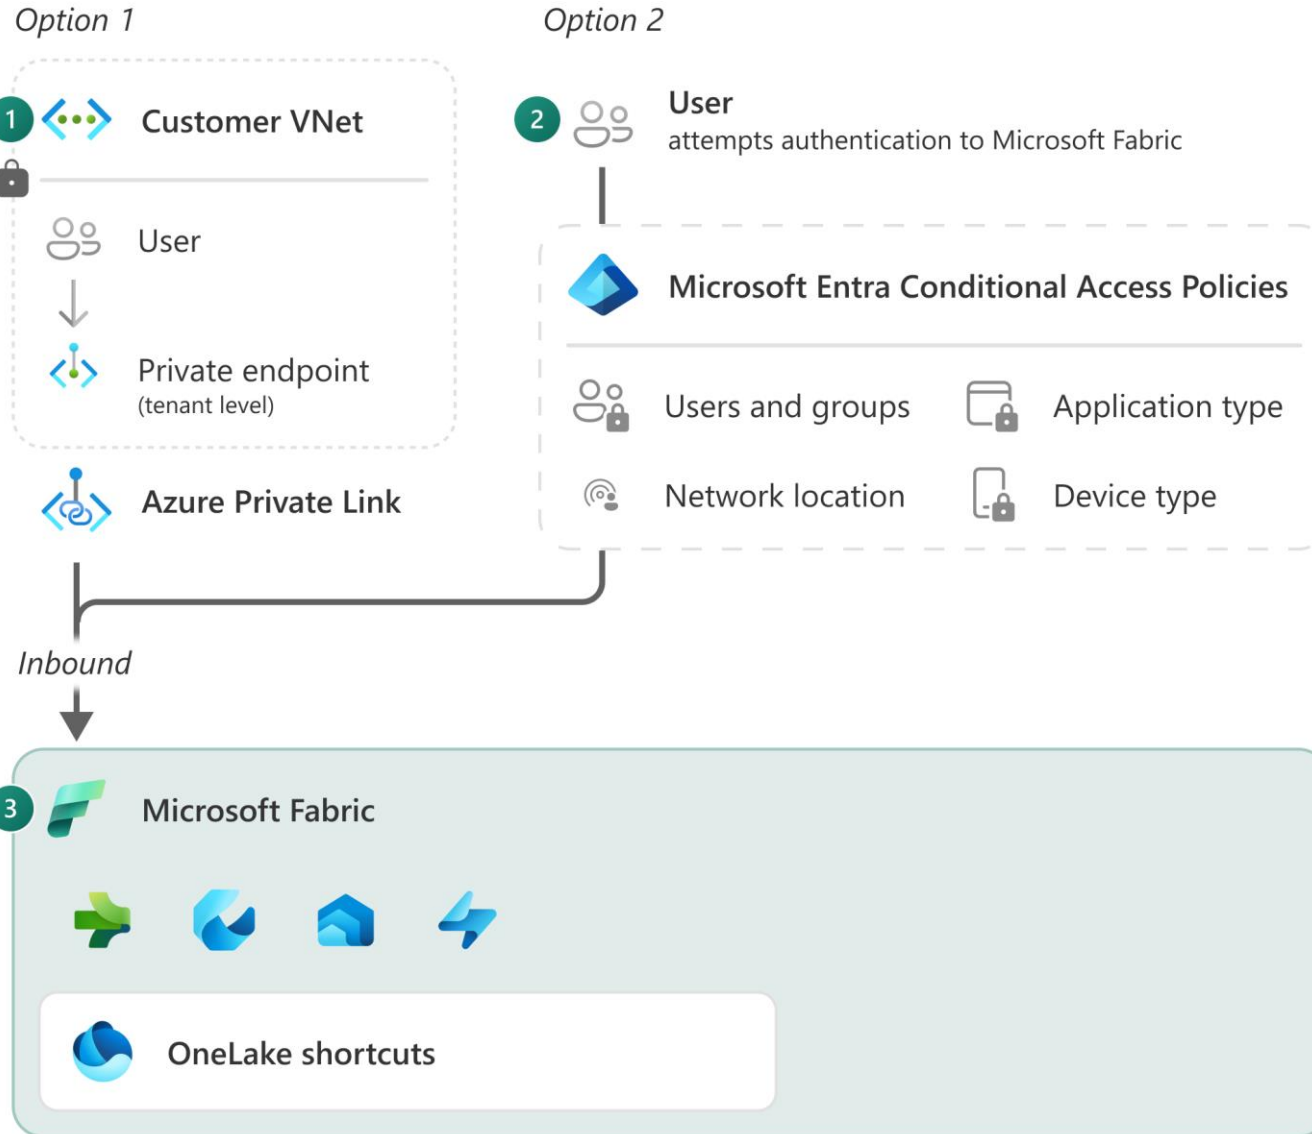




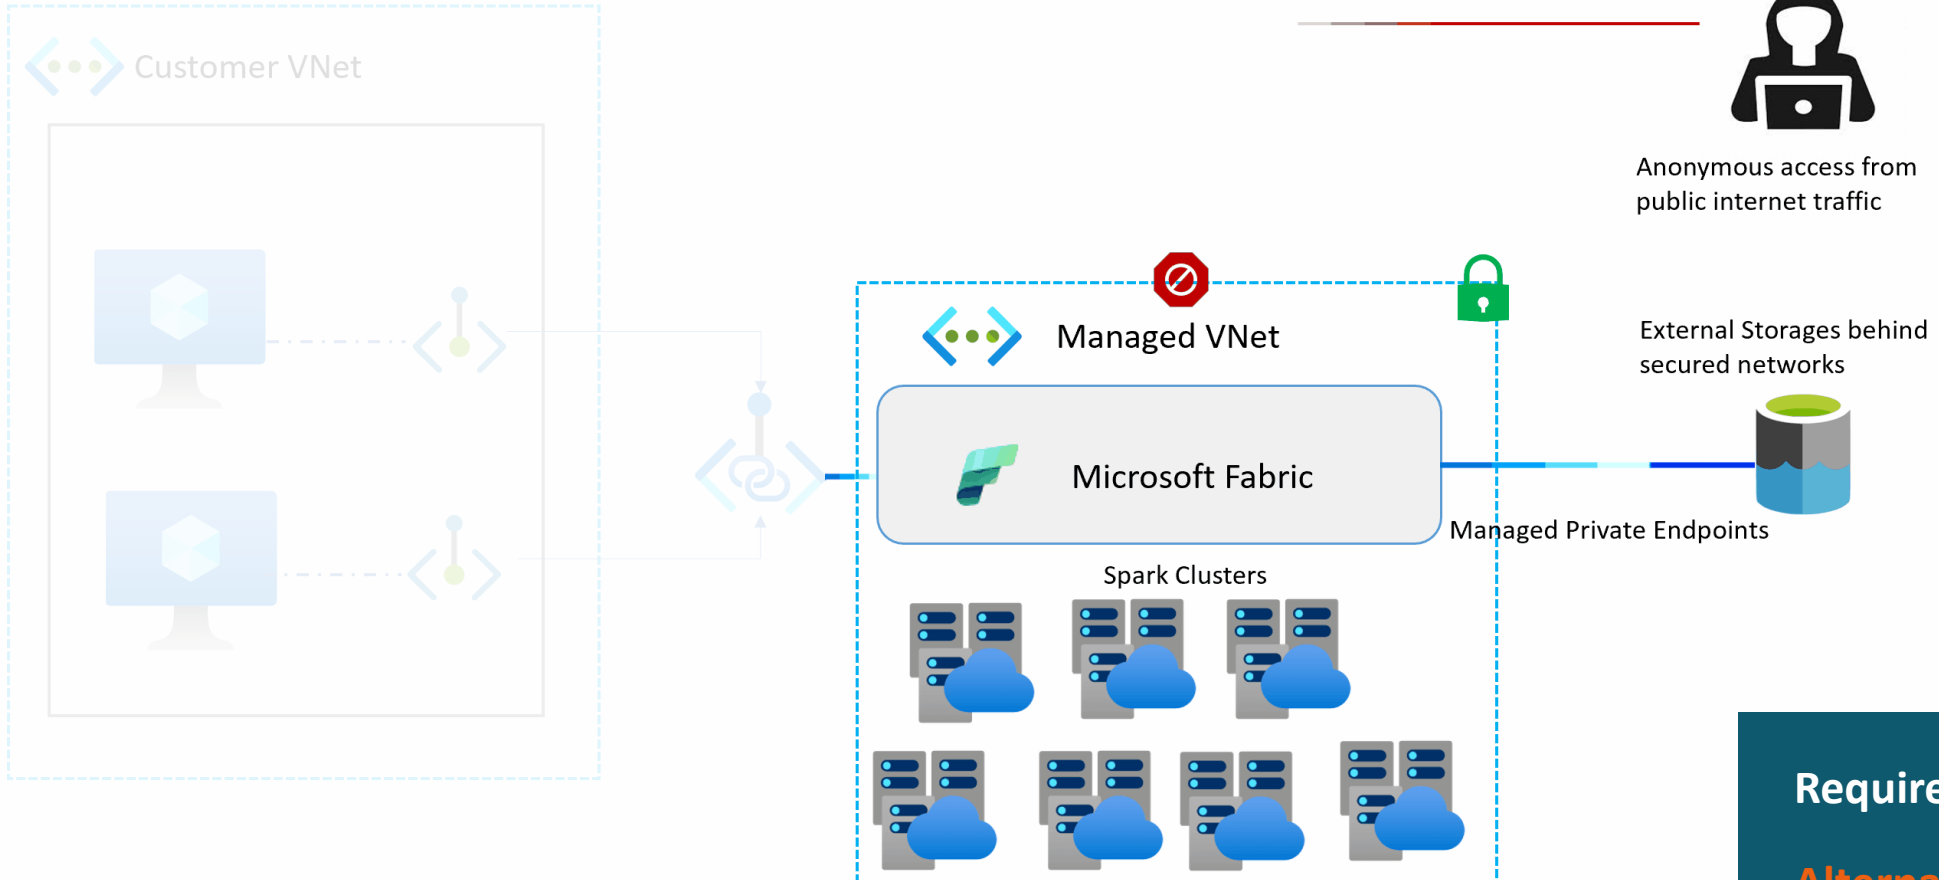
# Network Control - Inbound



Requires > F64 license



# Network Control - Outbound

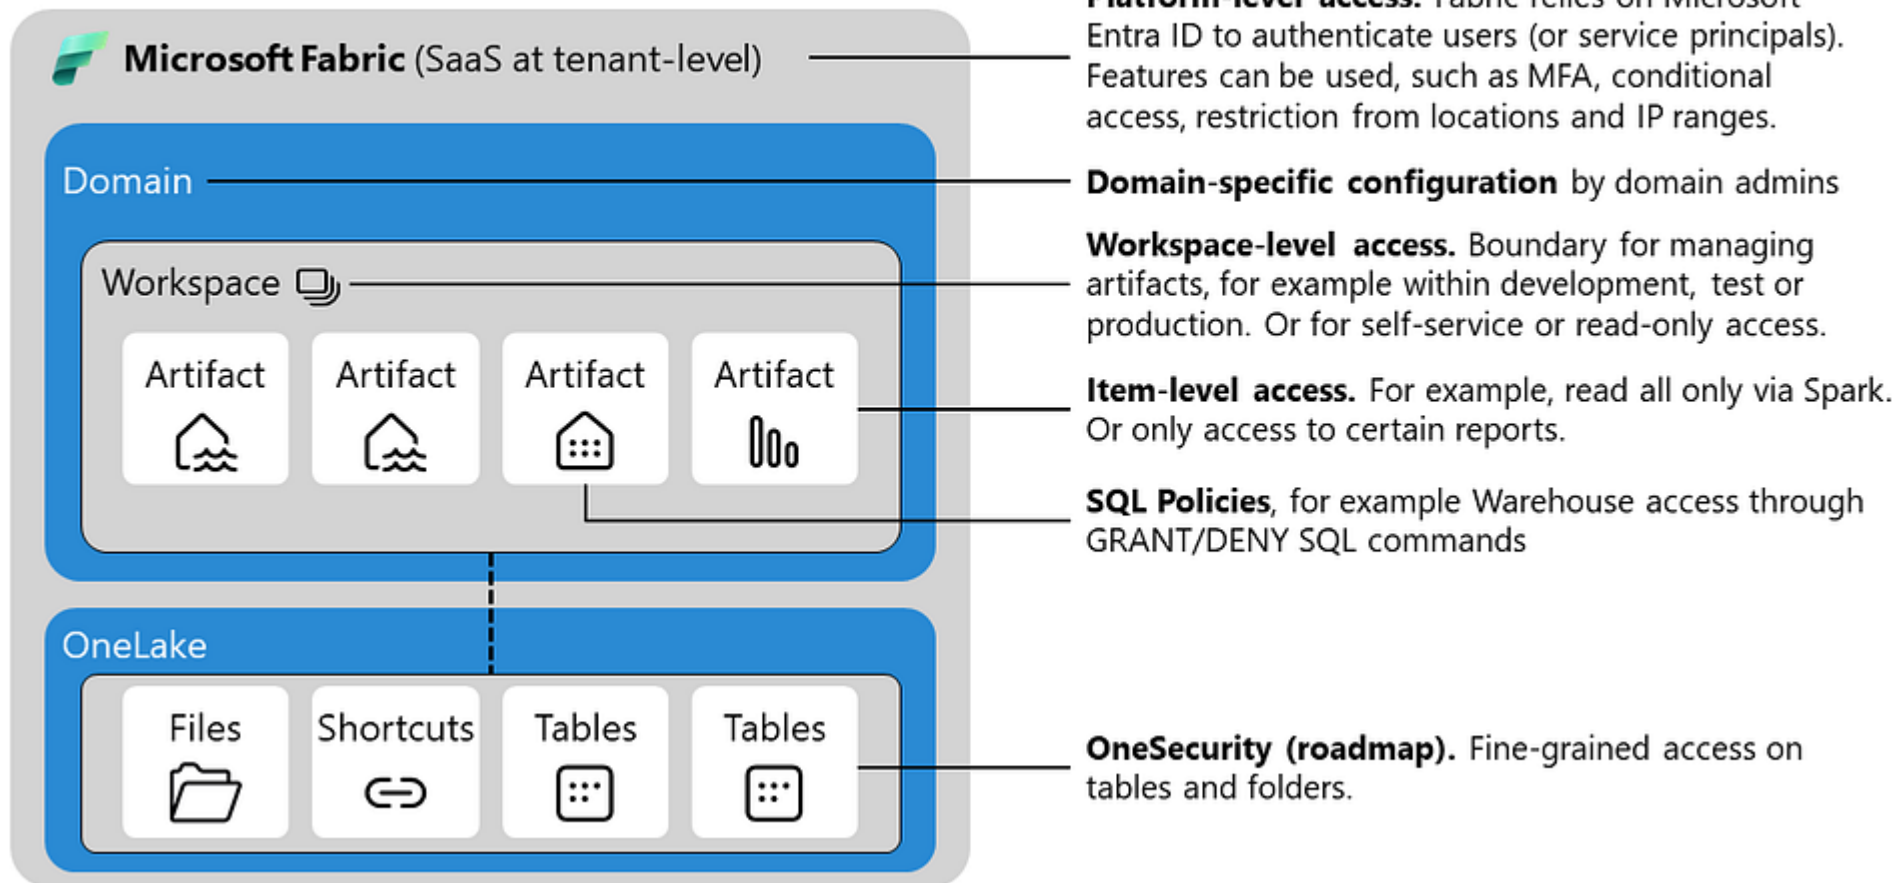


**Requires > F64 license**  
**Alternative to PE:**  
Trusted workspace access

# Workspace-, Item- & Data item security



## Microsoft Fabric: Multiple layers of security and access control




# Workspace-, Item- & Data item security



Tenant Level



Capacity Level




Workspace Level



Item Level



Object Level



# Workspace roles



## Granular se

Workspace roles and warehouse. However, constructs can be used

Microsoft Fabric data sensitive data from u features can provide

- Object-level se
- Column-level s
- Row-level secu
- Dynamic data r complete, such

### Grant people access

SecPocLH

People you share this Lakehouse with can open it and its SQL endpoint and read the default dataset. To allow them to read directly in the Lakehouse, grant additional permissions.

**Additional permissions**

- Read all SQL endpoint data ⓘ
- Read all Apache Spark ⓘ
- Build reports on the default dataset

**Notification Options**

- Notify recipients by email

ⓘ Depending on which additional permissions you select, recipients will have different access to the SQL endpoint, default dataset, and data in the lakehouse. For details, view lakehouse permissions documentation.

### Grant people access

SecPocWH

People you share this warehouse with can connect to it and read the default dataset. To allow them to read warehouse data, grant additional permissions.

**Additional permissions**

- Read all data using SQL
- Read all data using Apache Spark
- Build reports on the default dataset

**Notification Options**

- Notify recipients by email

ⓘ To define granular object-level security (OLS) for specific objects in the warehouse, use GRANT and DENY statements in T-SQL.



# Workspace-, Item- & Data item security



Tenant Level



Capacity Level



Workspace Level



Item Level



Object Level



Data Level ?



# What about the data in OneLake?



## OneLake security model

Estimated release timeline: Q1 2025

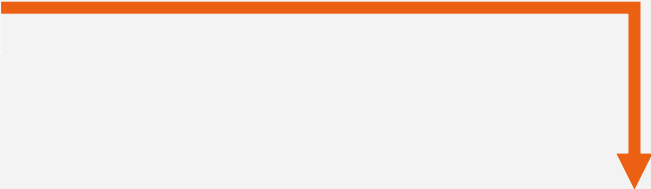
Release Type: Public preview

Managing data security across multiple analytical engines and copies of data is challenging. OneLake and Fabric simplify this by enabling the use of a single data copy across multiple analytical engines without any data movement or duplication. Taking the "one copy" concept further, OneLake is also enhancing security with a finer-grain model, allowing for table and folder access in addition to row and column level security. These security definitions live with the data and travel across shortcuts to wherever the data is used. Security defined at OneLake is universally enforced no matter which analytical engine is used to access the data.

# OneLake Data Access Roles



 Manage OneLake data access (preview)




### Turn on data access roles (preview)

You are about to turn on OneLake data access roles (preview) for this lakehouse. Using this capability, workspace admins, members, and contributors are permitted to change Onelake data access permissions at the folder level. [Learn more](#)

**ⓘ** After OneLake data access roles is applied to GoldLakeHouse only users with Write permission can share the data in this item using External data sharing (preview). Any existing external data shares may stop working. [Learn more](#)

[Continue](#) [Cancel](#)

# How are permissions evaluated?



### Note

\*Since Workspace Admin, Member and Contributor Roles automatically grant Write permissions to OneLake, it overrides any OneLake RBAC Read permissions.

Expand table

Workspace Role	Does OneLake apply RBAC Read permissions?
Admin, Contributor, Member	No, OneLake Security will ignore any OneLake RBC Read permissions
Viewer	Yes, if defined, OneLake RBAC Read permissions will be applied

# Demo


Power BI Home Search

Introducing the Power BI app in Teams Collaborate with your team members on data and take action. Select Open in Teams to get started [Learn more](#) Open in Teams

**+ New report** New items saved to: My workspace


**Recommended**

**My workspace**  
You frequently open this




**Open**

**Explore basic Power BI concepts**  
Getting started with Power BI




**Open**

**Explore the 100 most useful productivity tips**  
Explore this data story




**Open**

**Cancer statistics in the USA**  
Explore this data story



**Open**

**Intro—What is Power BI?**  
Getting started with Power BI



**Open**

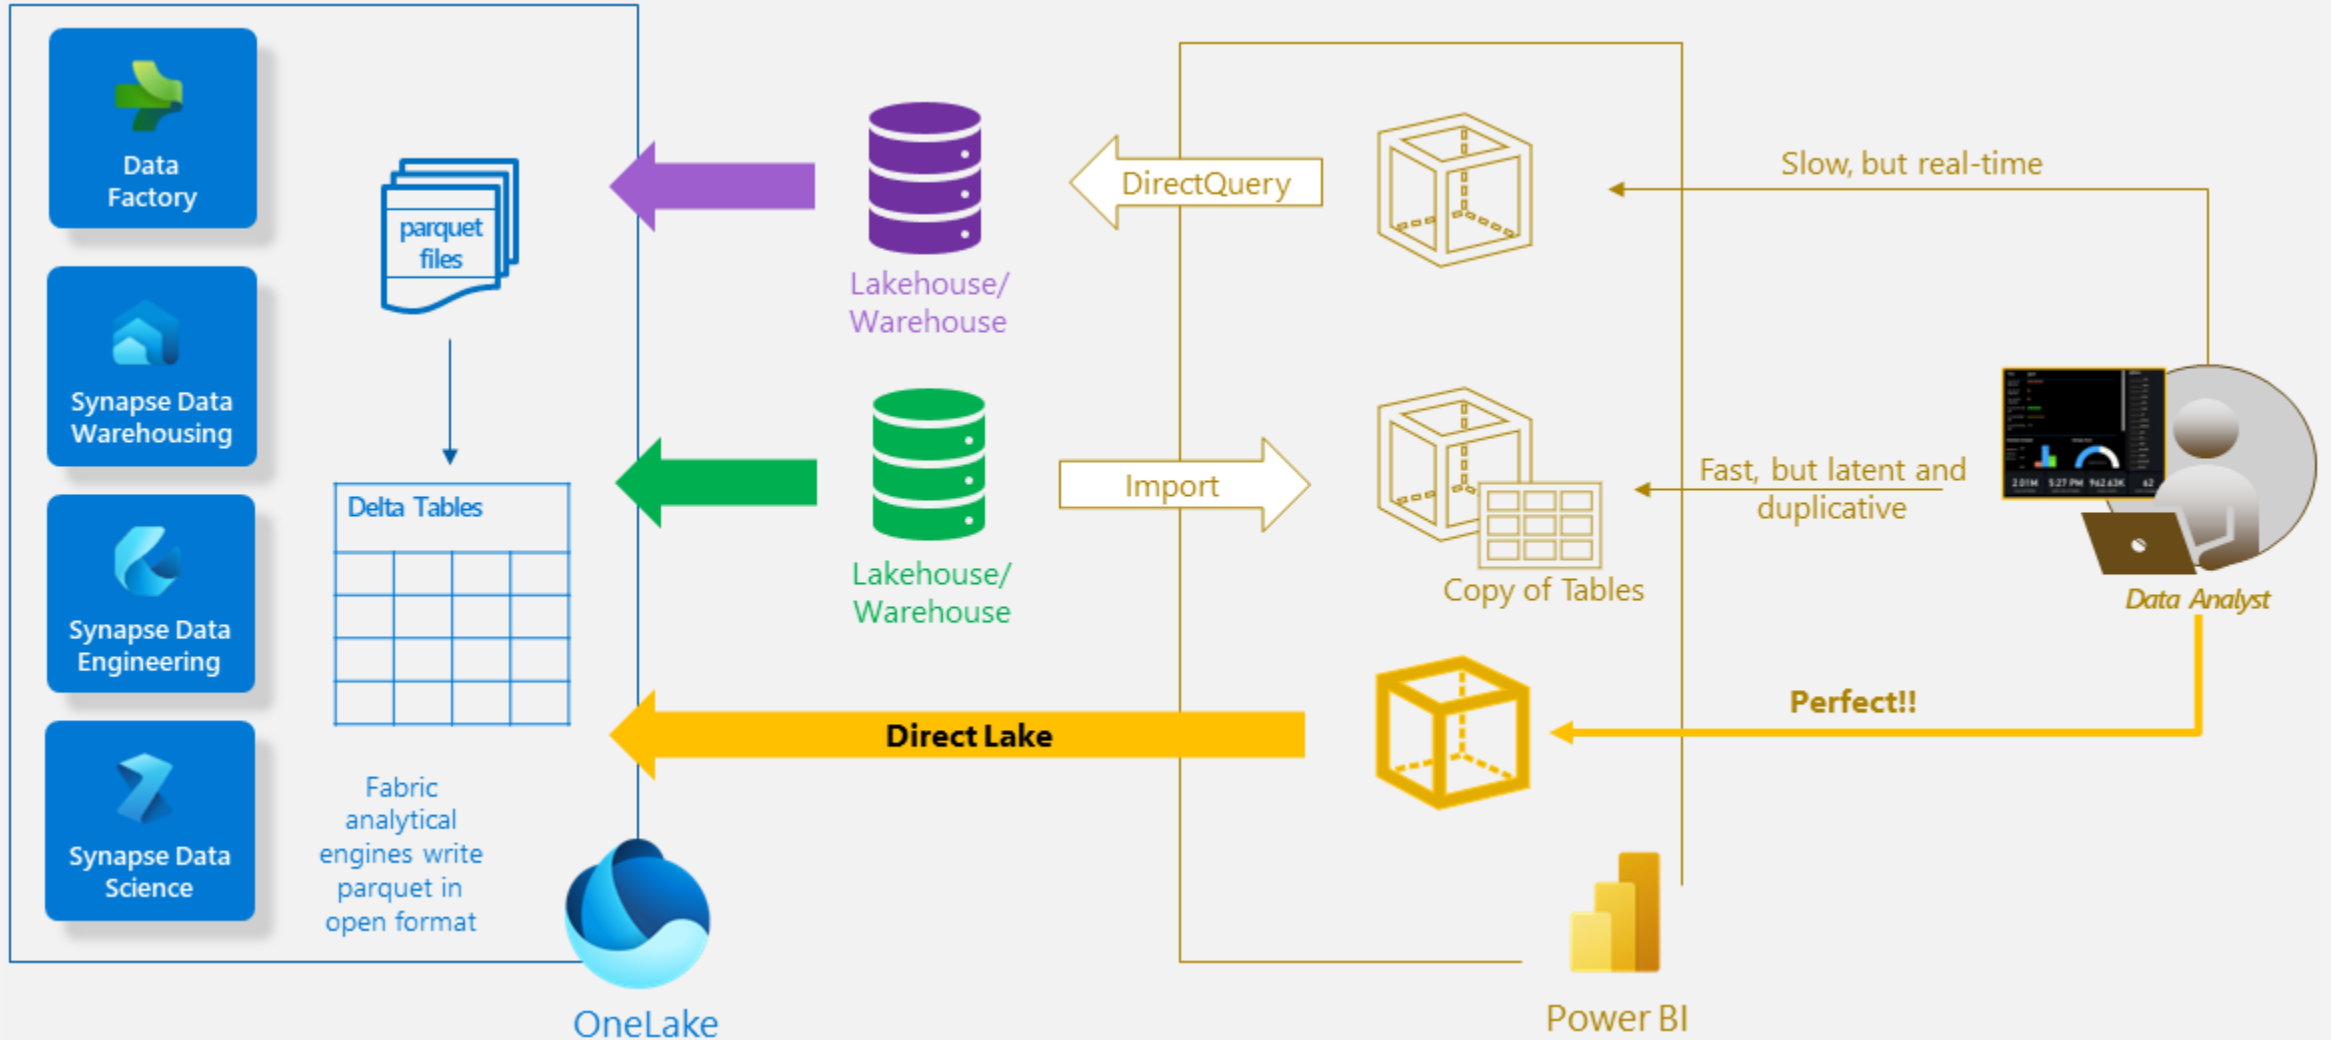
**Recent** Favorites My apps Filter by keyword Filter

Name	Type	Opened	Location	Endorsement	Sensitivity
My workspace	Workspace	13 days ago	Workspaces	—	—
BronzeLakeHouse	Lakehouse	13 days ago	—	—	—

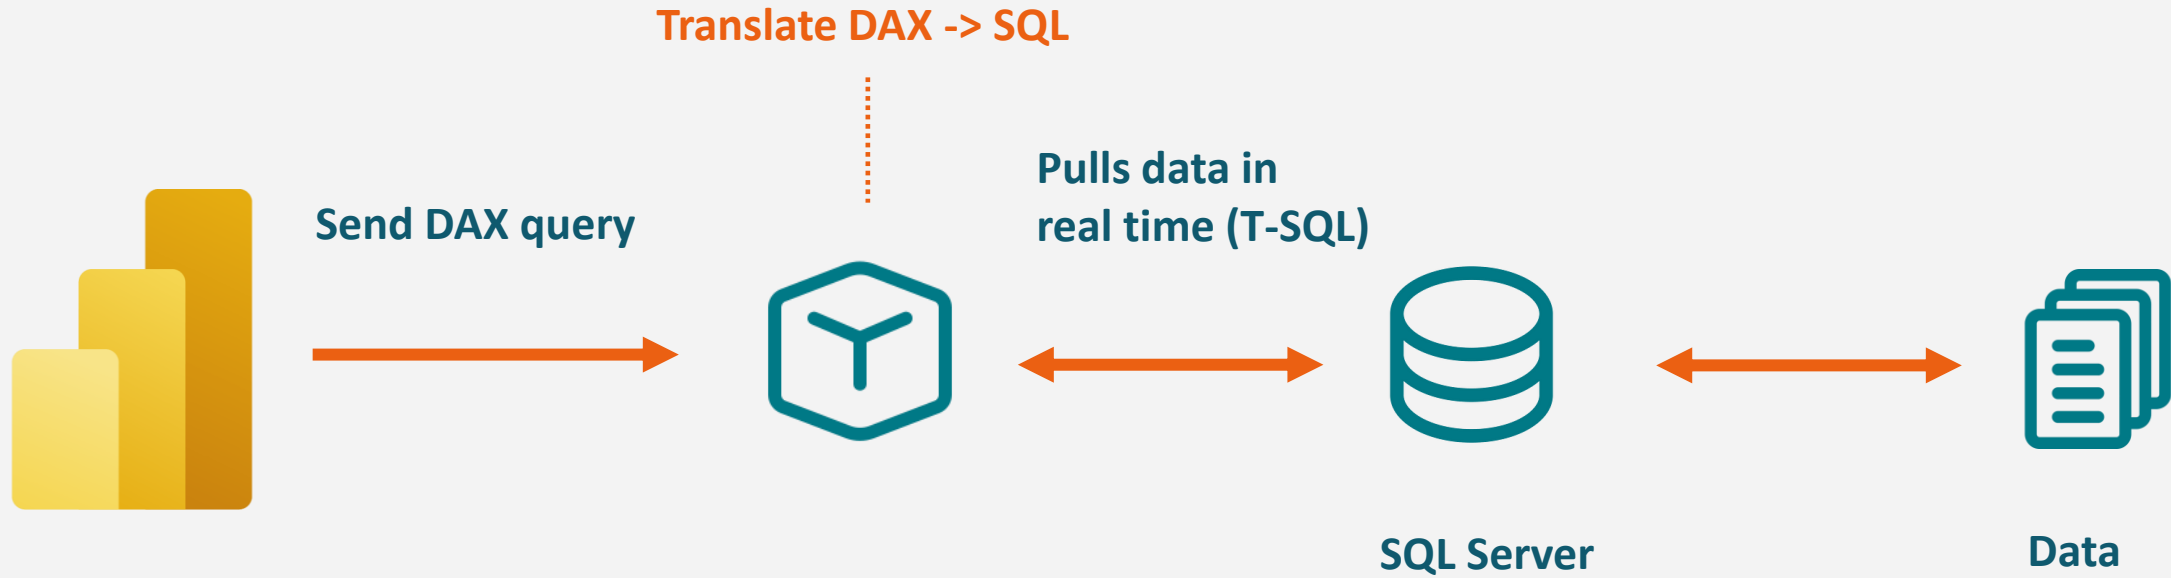


# Direct Lake

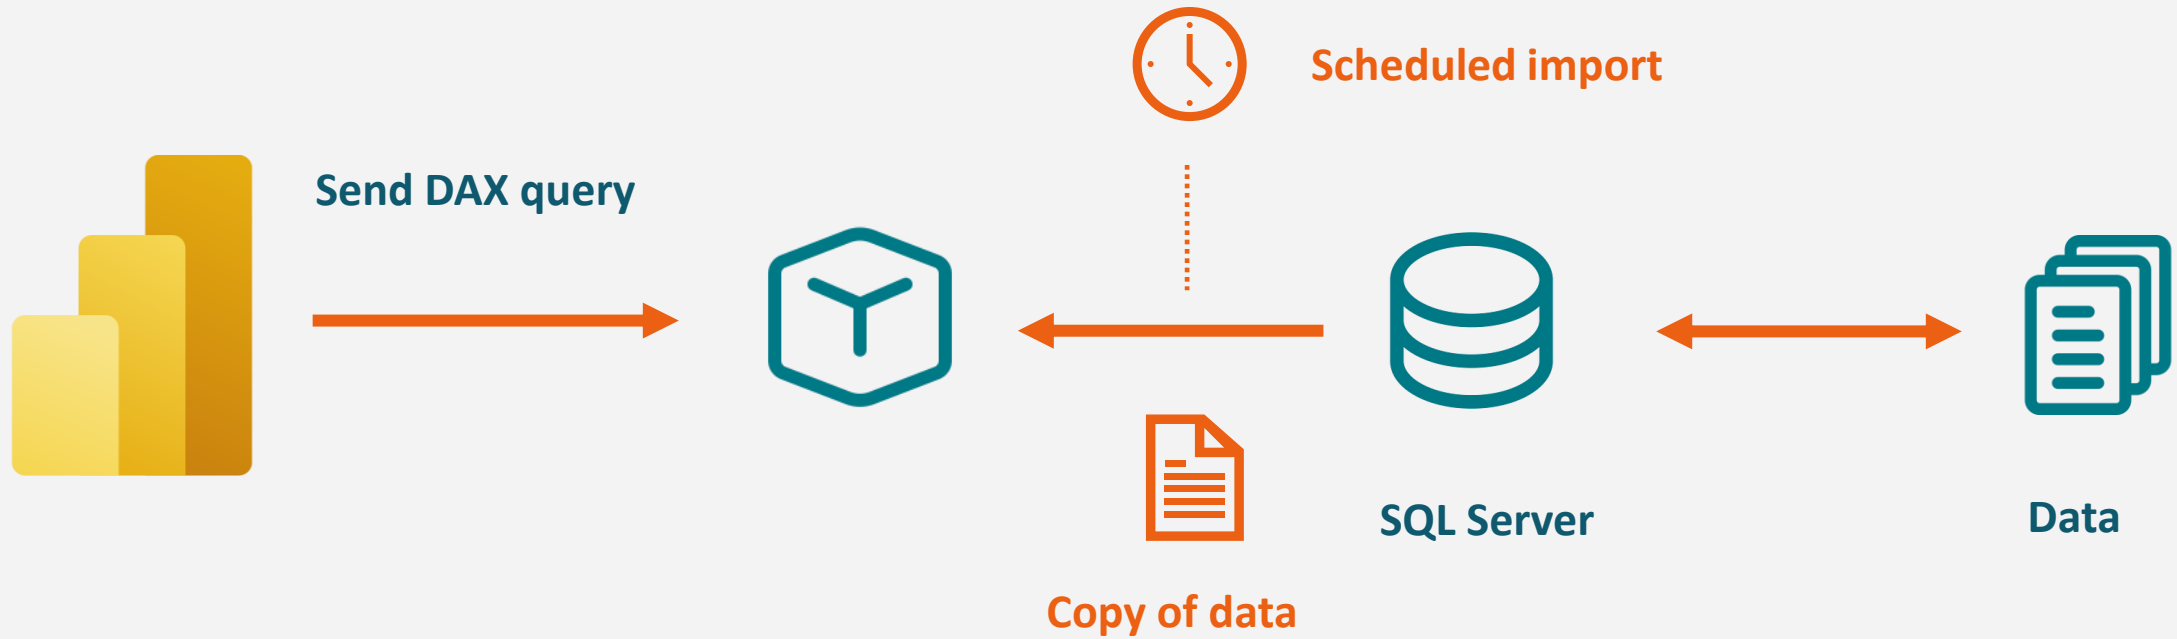
# Direct Lake

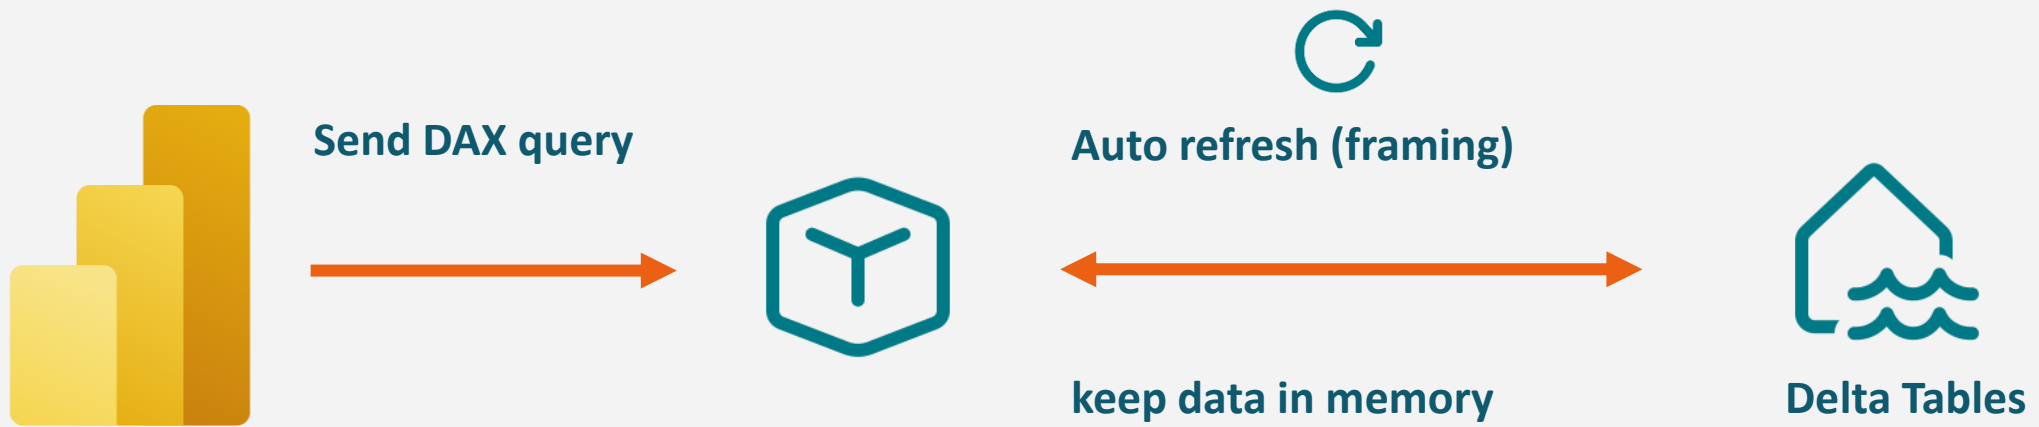


# DirectQuery

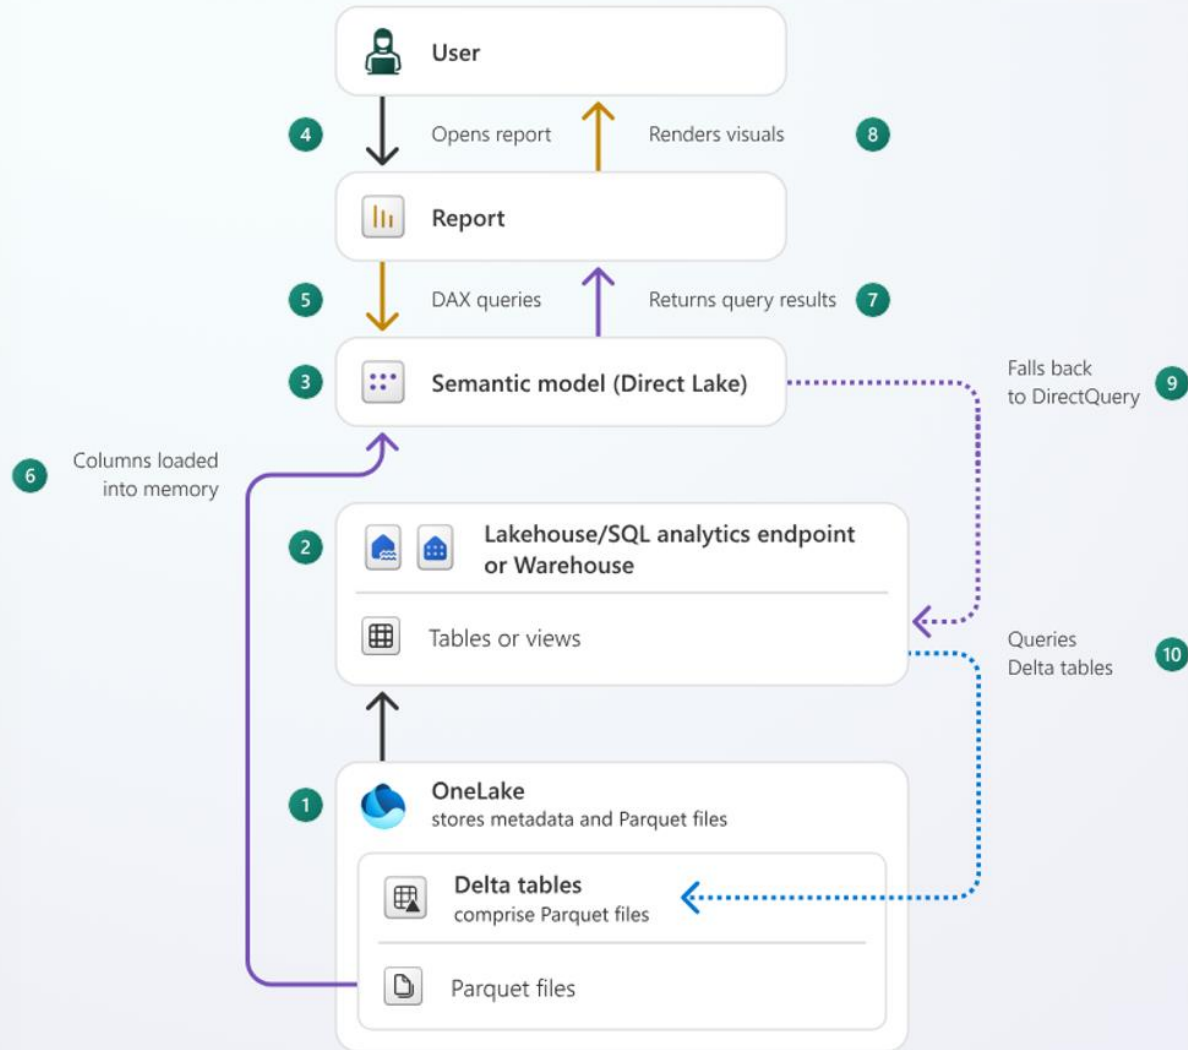


# Import





# Direct Lake semantic model



# Demo

The screenshot shows a Databricks workspace interface. At the top, there's a search bar and a notification: "Fabric Trial: 39 days left". The main area is a notebook with three cells. The first cell contains Python code for imports and has a status bar indicating it was executed in 10 seconds. The second cell contains a SQL query and has a status bar indicating it was executed in 2 seconds. The third cell contains a DataFrame creation command. The left sidebar shows the Explorer view with folders for Resources, Lakehouses, and Warehouses. The bottom status bar shows "Session ready" and "AutoSave: On".

DirectLakeDemo | No label · Saved

Home Edit Run View

AutoSave On | Add code cell below | Clear all output(s) | Browse code snippets | Find and replace

Explorer

- + Data sources
- Resources  
Uploaded data and files
- Lakehouses  
1 item(s) added
- Warehouses  
0 item(s) added

Other people in your organization may have access to notebooks and Spark job definitions in this workspace. Carefully review this item before running it.

```
1 import sempy.fabric as fabric
2 import pyspark.sql.functions as F
```

- Command executed in 10 sec 162 ms by Steven van den Beemt on 4:30:52 PM, 10/03/24

PySpark (Python)

### Step 1: What's in our lakehouse?

```
1 %%sql
2
3 SELECT puYear, count(*)
4 FROM taxirides
5 GROUP BY puYear
```

- Command executed in 2 sec 336 ms by Steven van den Beemt on 4:19:35 PM, 10/03/24

Plaintext

### Step 2: Check what's in memory

```
1 df_dax = fabric.evaluate_dax(
2     "NewYorkTaxi",
3     """
4     SELECT
5         MEASURE_GROUP_NAME,
6         ATTRIBUTE_NAME,
7         DATATYPE,
8         DICTIONARY_SIZE,
9         DICTIONARY_TEMPERATURE,
10        DICTIONARY_LAST_ACCESSED
11     FROM $SYSTEM.DISCOVER STORAGE TABLE COLUMNS
```

Session ready | AutoSave: On

Selected Cell 12 of 20 cells



## └ What makes it fast?

- └ PowerBi queries Parquet files stored in OneLake directly
- └ Loaded in memory using the VertiPaq engine
- └ Data will be transcoded on-the-fly in VertiPaq format
- └ V-order storage

## └ Fallback to DirectQuery mode



# Wrap-up

# Key takeaways



- ↵ Outthink your workspace design upfront
- ↵ Make sure your capacity is deployed in Europe
- ↵ Authorization within MS Fabric is still a challenge
  - ↵ Hopefully, this will improve in Q1 - 2025
- ↵ DirectLake speeds up your reports when using large datasets





Q&A