



Microsoft Fabric

This presentation is the property of Microsoft and is intended for informational and educational purposes only. You may use, copy, and distribute this presentation for your personal, non-commercial purposes. You may not modify, alter, or create derivative works from this presentation without the prior written consent of Microsoft. You may not use this presentation to misrepresent, defame, or disparage Microsoft or its products, services, or affiliates. You may not use this presentation to endorse or promote any other products, services, or organizations without the prior written consent of Microsoft.

By using this presentation, you agree to abide by these terms. If you do not agree, you must not use this presentation. Microsoft reserves the right to change these terms and conditions at any time without notice. Microsoft disclaims any and all warranties, express or implied, relating to this presentation, including but not limited to the accuracy, completeness, timeliness, or suitability of the information contained herein. Microsoft is not liable for any damages, losses, or liabilities arising from your use of or reliance on this presentation.

Please review the terms of use posted in the content library.

#FABCONSQLCON2026

**FABCON**

Microsoft Fabric  
COMMUNITY CONFERENCE

**SQLCON**

Microsoft SQL  
COMMUNITY CONFERENCE

**ATLANTA** MARCH 16 - 20, 2026



# Taking Direct Lake to the Next Level

Phil Seamark

Power BI PM

**in** /seamark

**X** @PhilSeamark

Christian Wade

Power BI PM

**in** /christianwade1

**X** @\_christianwade

# Agenda

---

1

Updates

2

Fundamentals & lessons learned

3

Q&A



Fabric IQ

## Unified Data

Structured, Unstructured,  
Real-Time, Graph





# Fabric IQ

## Business Intelligence

Semantic Models



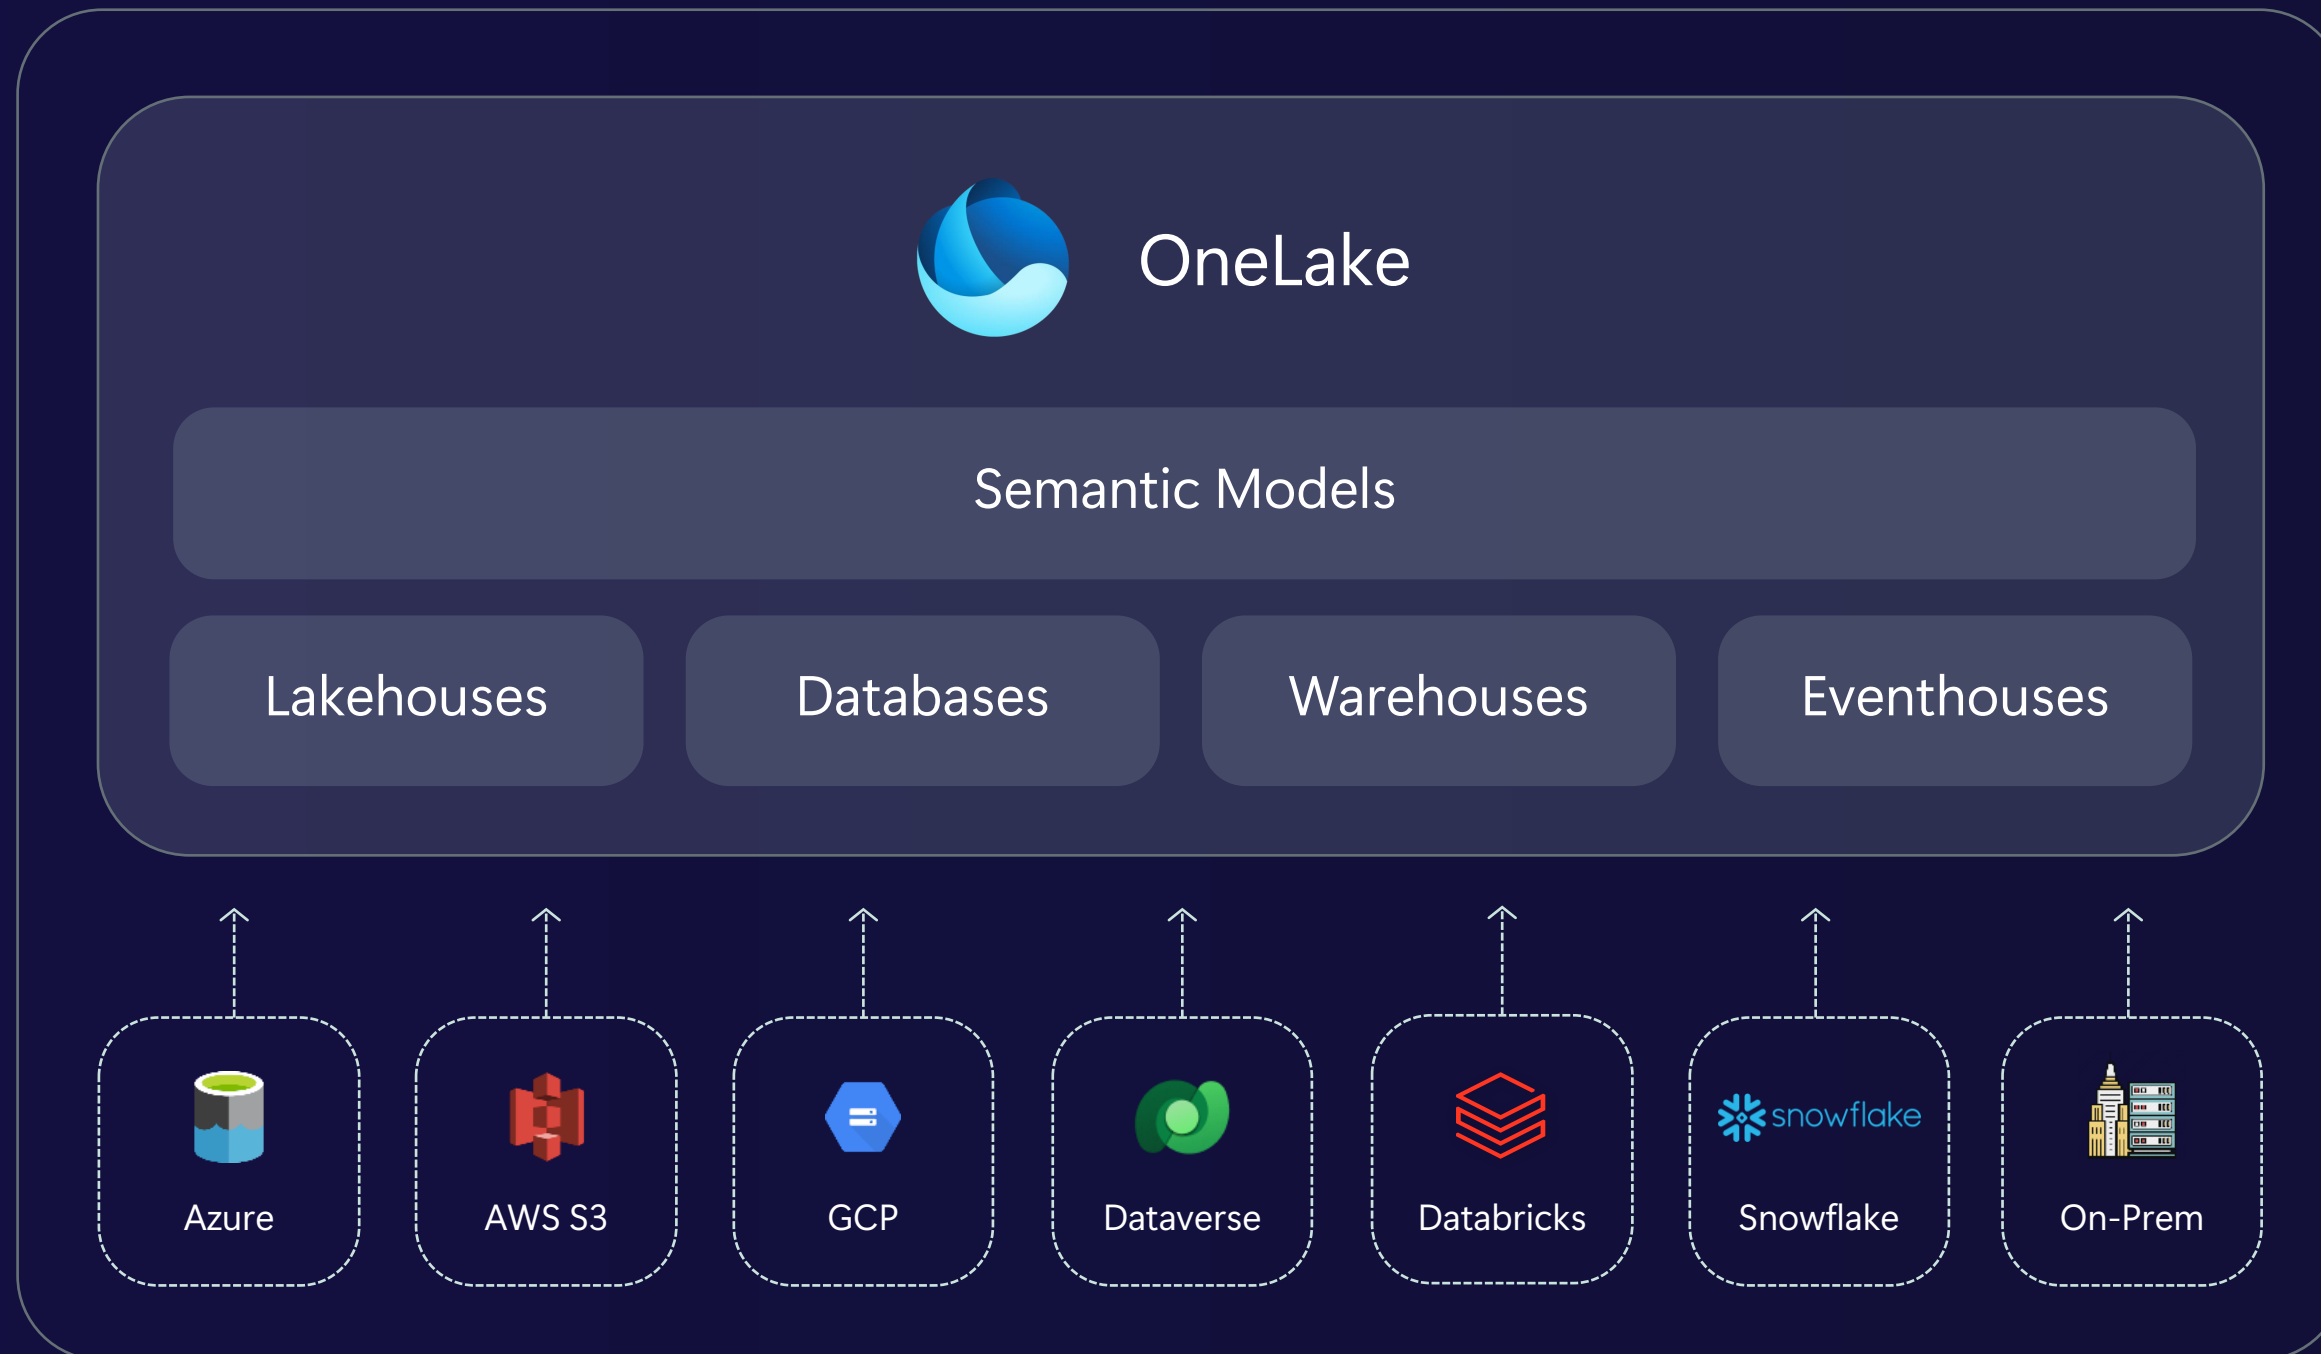
## Unified Data

Structured, Unstructured,  
Real-Time, Graph



# Semantic Models

Grounding AI in trusted business knowledge



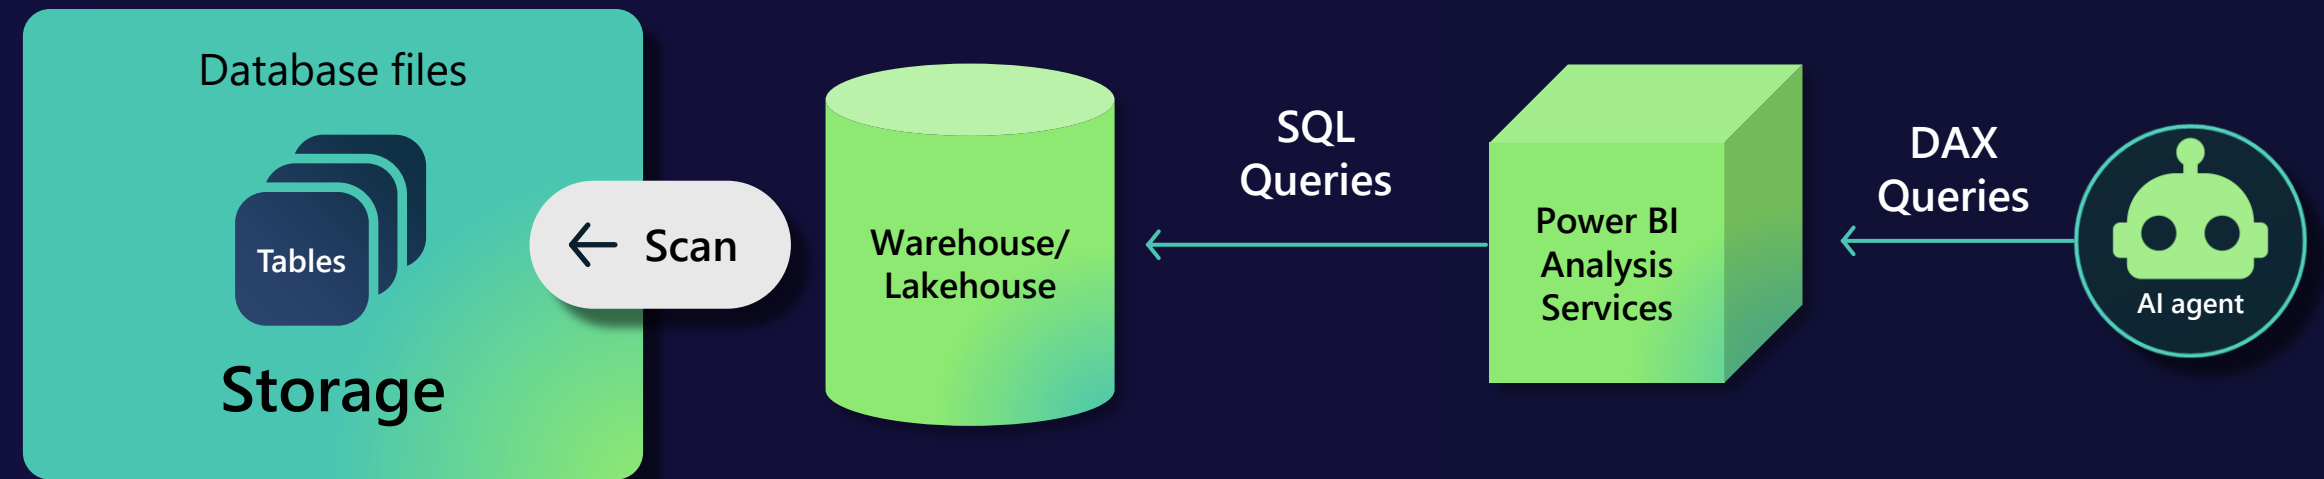
Expertly curated repositories of business data and metrics

Industry rush toward unified semantic layers as **key AI enabler**

**35M+** users regularly using semantic models

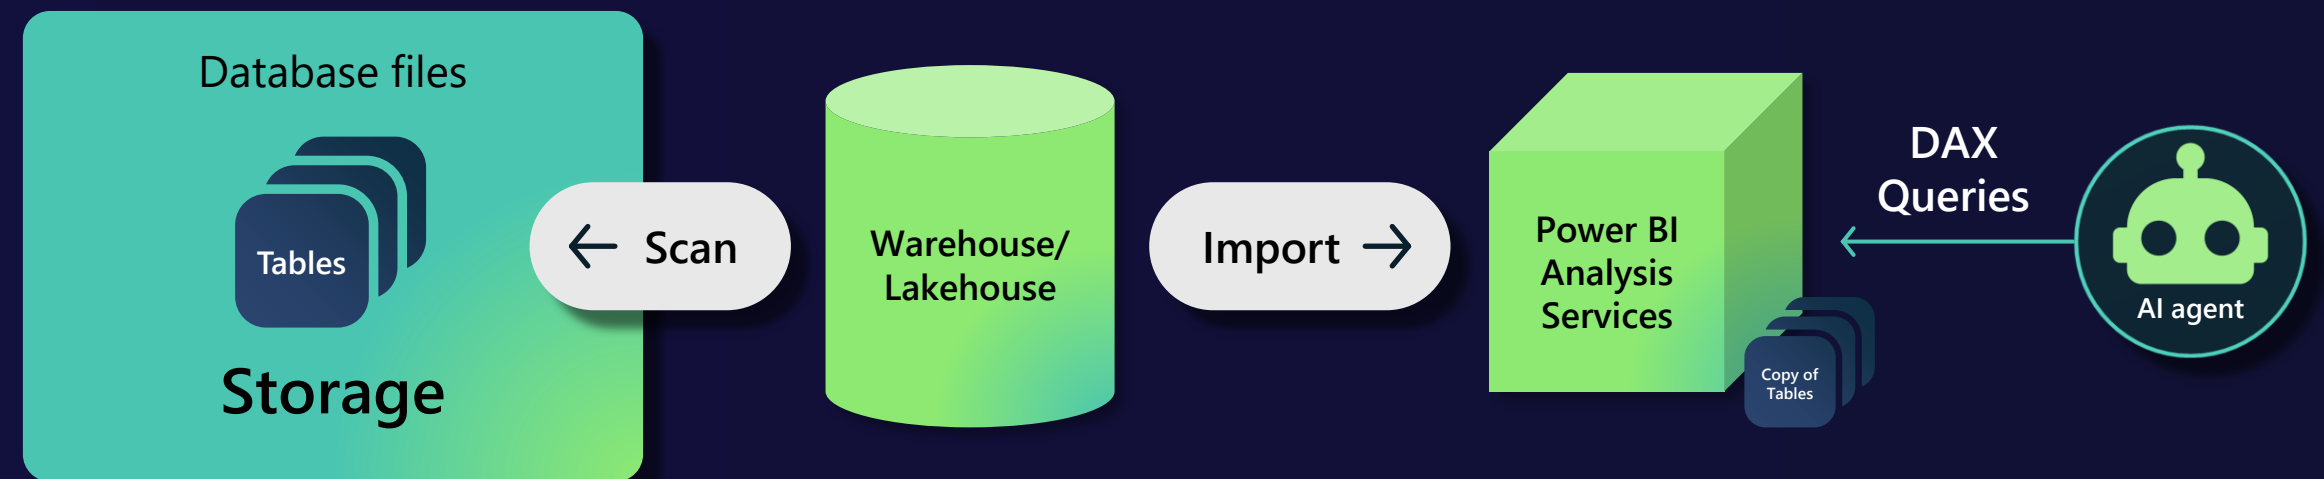
# “Direct Query Mode”

Slow, but real time



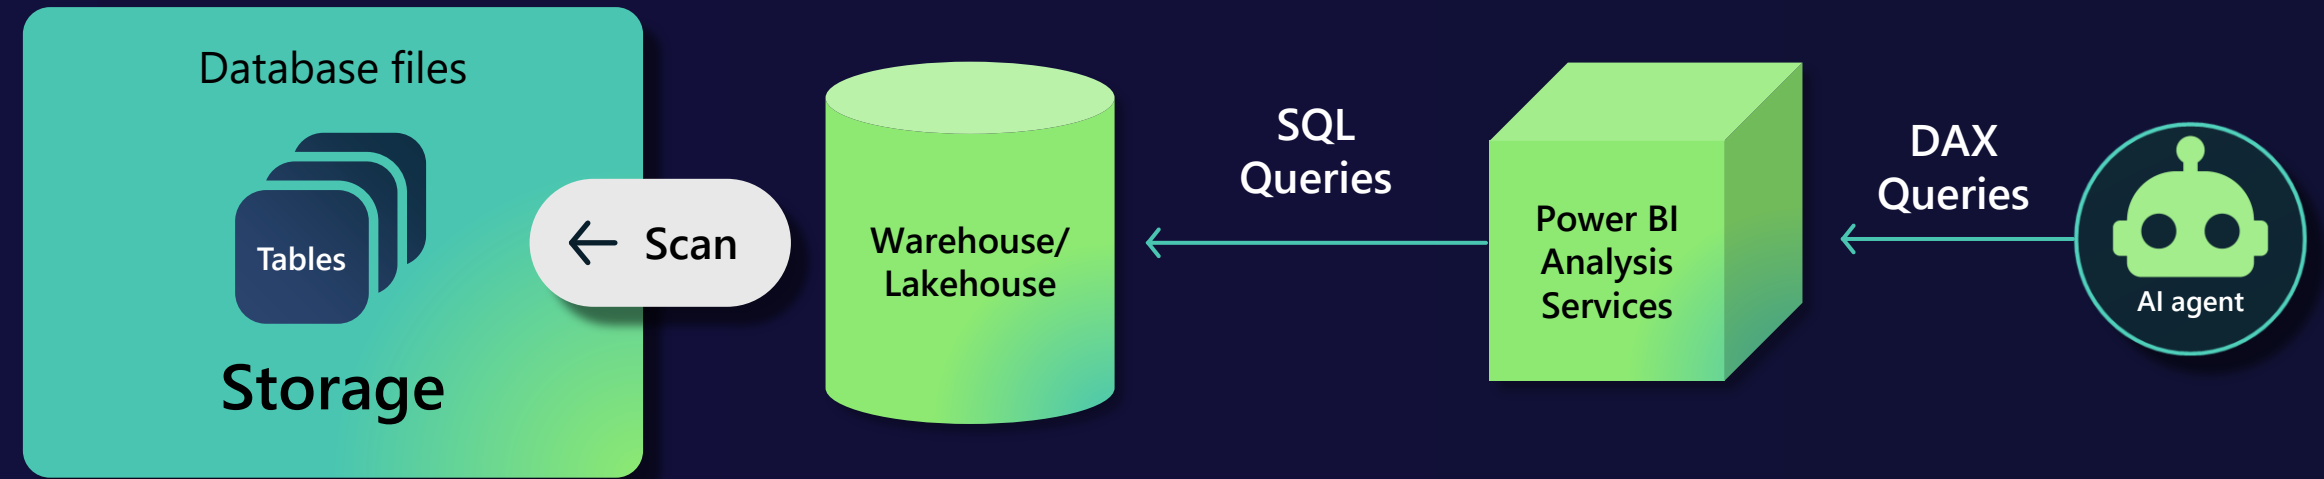
# “Import Mode”

Fast, but duplicative



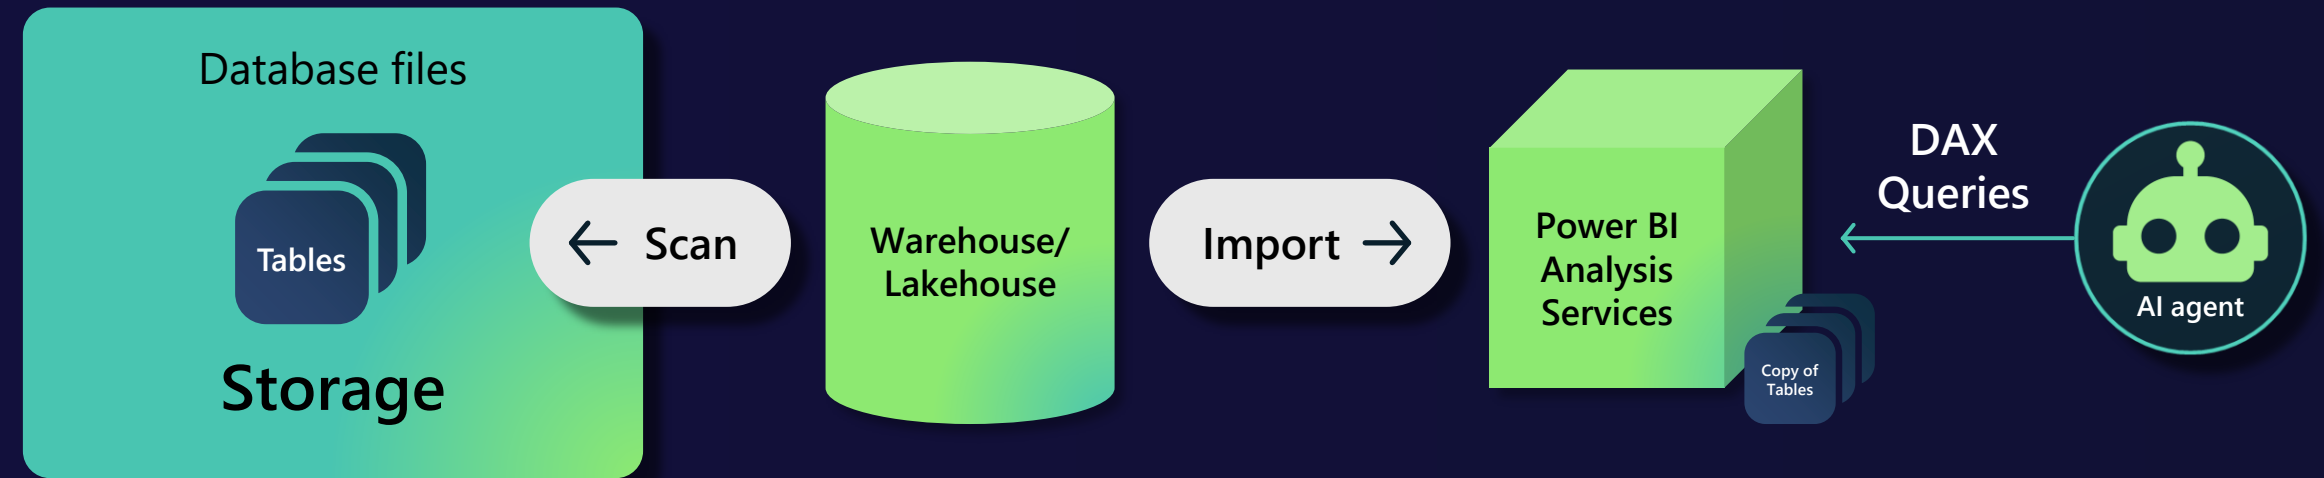
# “Direct Query Mode”

Slow, but real time



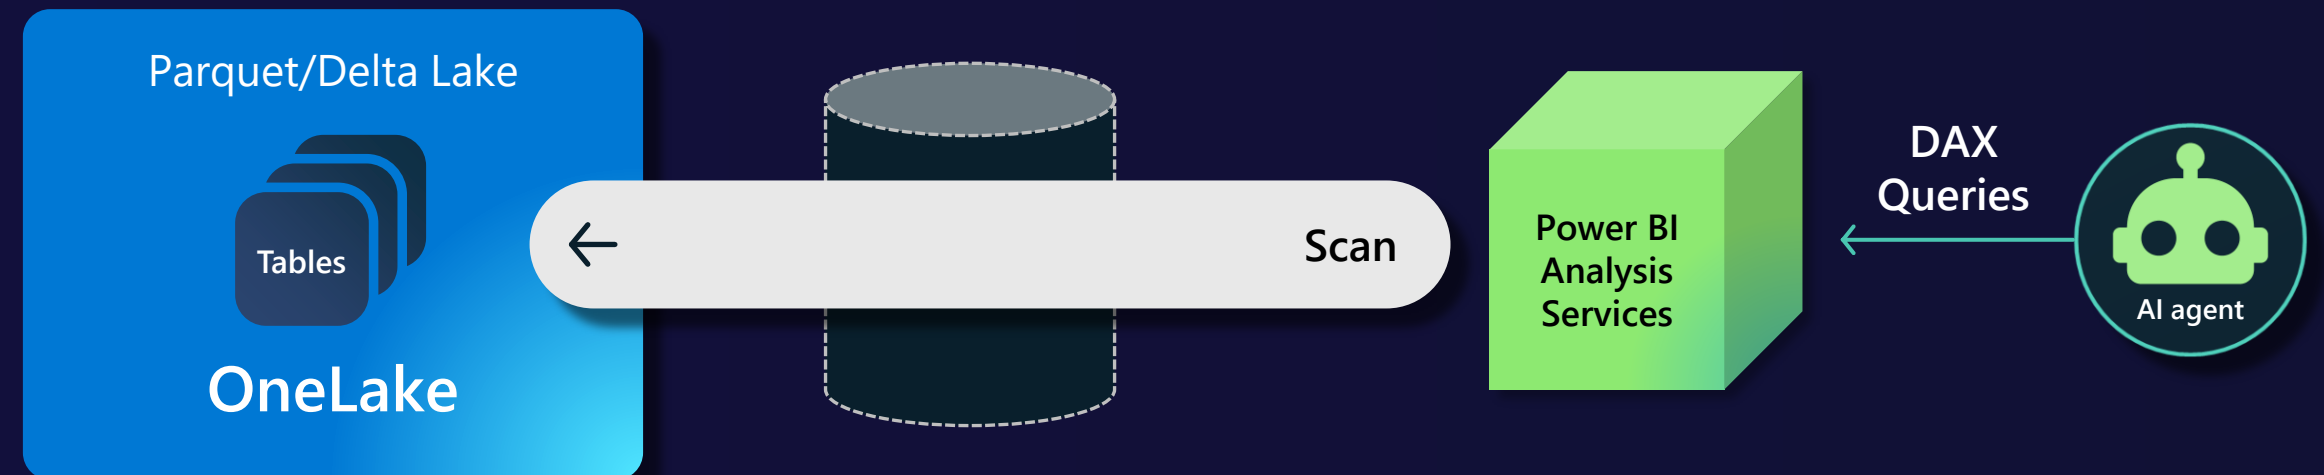
# “Import Mode”

Fast, but duplicative



# “Direct Lake Mode”

Fast, real time, no duplication



# Direct Lake

## Direct Lake on OneLake (new)

- Compatible with OneLake Security

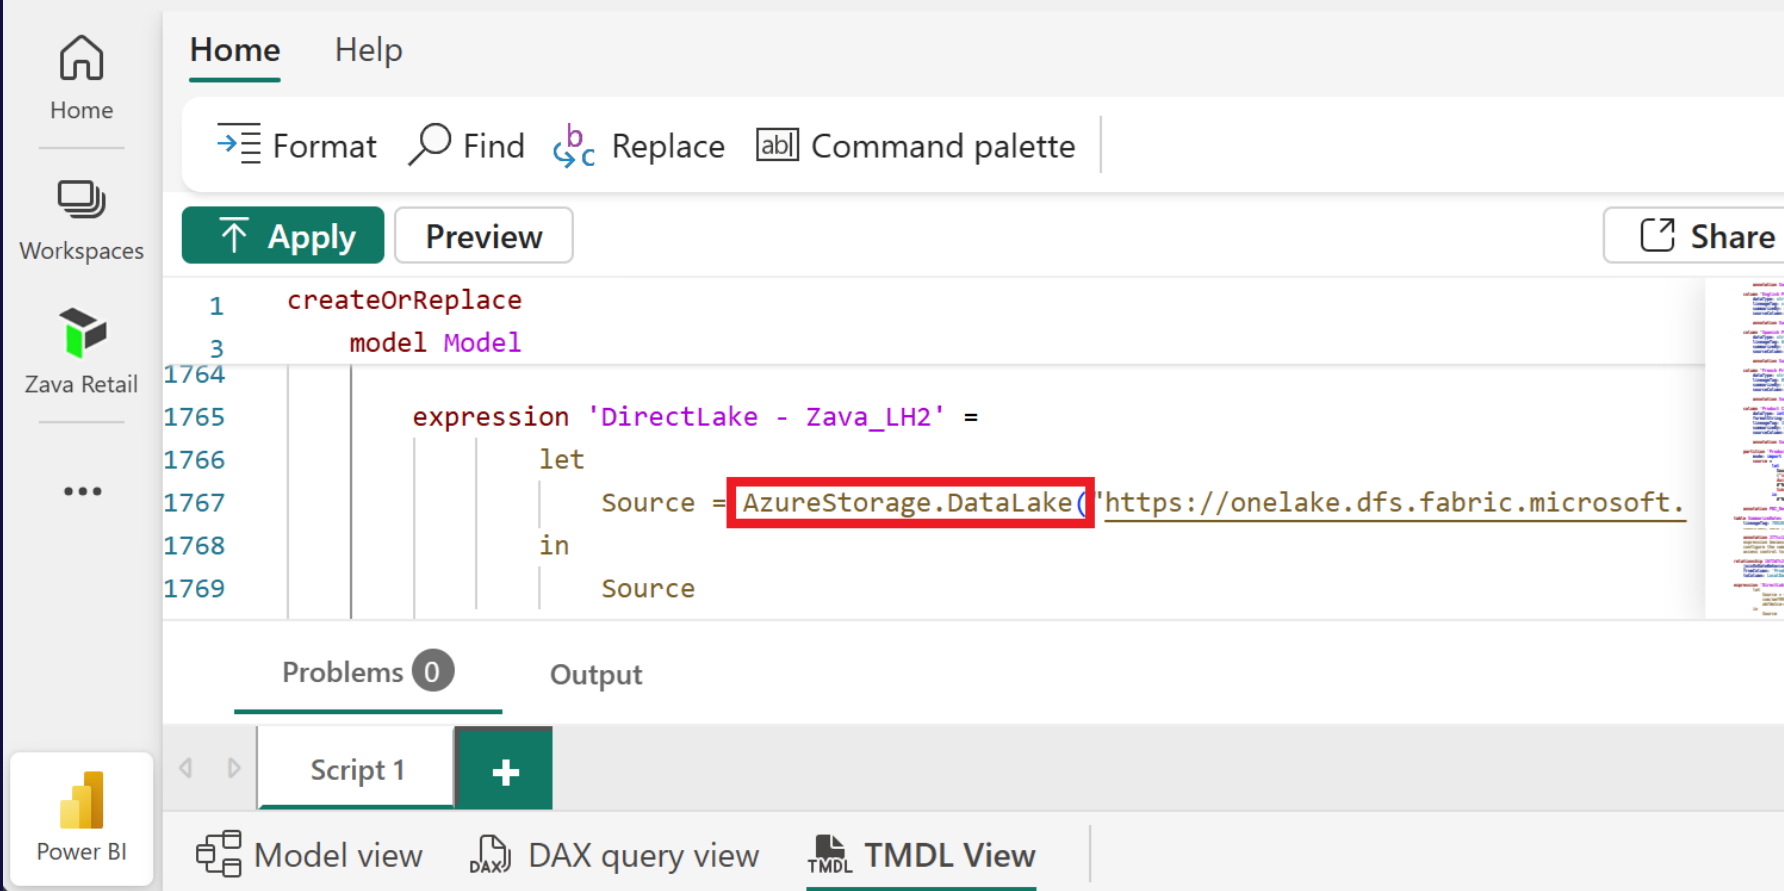
## Direct Lake on SQL

- Compatible with SQL Security  
Fallback to DirectQuery

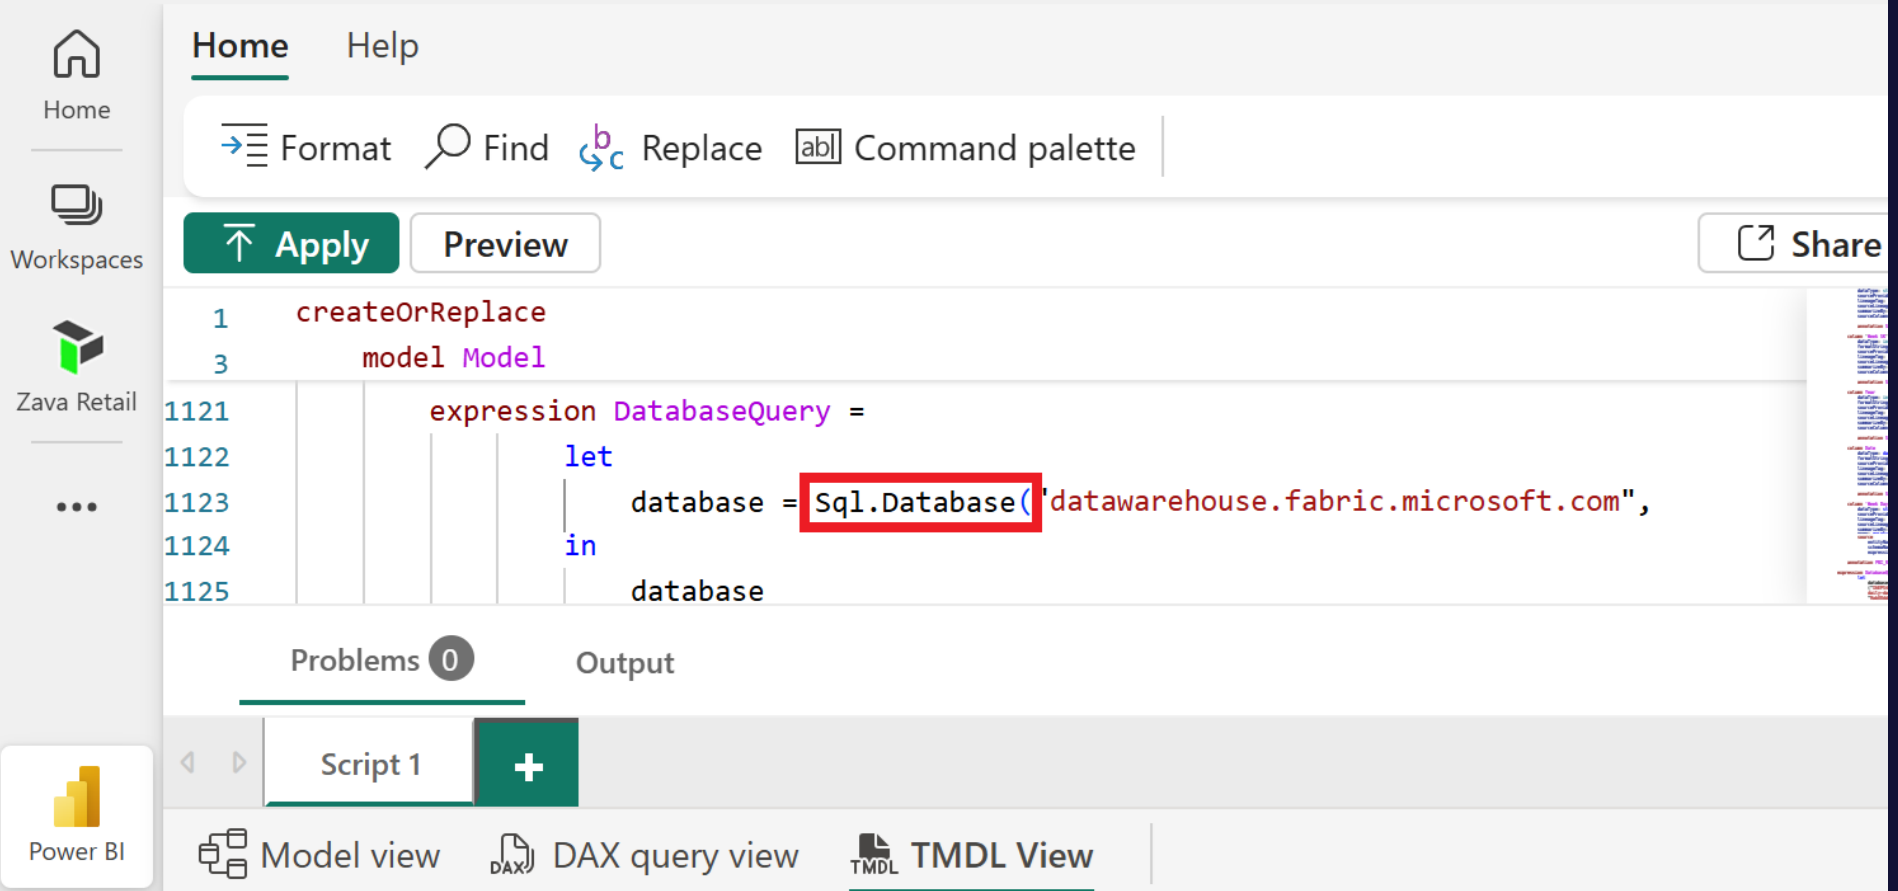
# Direct Lake

## Direct Lake on OneLake (new)

## Direct Lake on SQL



```
1 createOrReplace
3 model Model
1764
1765     expression 'DirectLake - Zava_LH2' =
1766         let
1767             Source = AzureStorage.DataLake('https://onelake.dfs.fabric.microsoft.com',
1768             in
1769             Source
```



```
1 createOrReplace
3 model Model
1121     expression DatabaseQuery =
1122         let
1123             database = Sql.Database('datawarehouse.fabric.microsoft.com',
1124             in
1125             database
```

# Demo

Direct Lake on OneLake composite model

**Explorer**

+ Warehouses

- Zava\_LH3
  - Schemas
    - dbo
      - Tables
        - Channel\_Sa
        - Customer
        - Date**
        - Outdoor\_Si
        - Product
        - Product\_Ca
        - Product\_Su
      - Views
      - Functions
      - Stored Proce...
    - INFORMATION\_...
    - queryinsights
    - sys
    - Security
  - Queries
    - My queries
    - Shared queries

Date

Data preview - Date

ID	Date_ID	Day_Name	Month_Name	Month_Of_Y...	Quarter_Na...
1	20120325	03-25-2012	March 2012	201203	Q1 2012
2	20200216	02-16-2020	February 2020	202002	Q1 2020
3	20200112	01-12-2020	January 2020	202001	Q1 2020
4	20200329	03-29-2020	March 2020	202003	Q1 2020
5	20190127	01-27-2019	January 2019	201901	Q1 2019
6	20190217	02-17-2019	February 2019	201902	Q1 2019
7	20110313	03-13-2011	March 2011	201103	Q1 2011
8	20010128	01-28-2001	January 2001	200101	Q1 2001
9	20010121	01-21-2001	January 2001	200101	Q1 2001
10	20160124	01-24-2016	January 2016	201601	Q1 2016
11	20140209	02-09-2014	February 2014	201402	Q1 2014
12	20210221	02-21-2021	February 2021	202102	Q1 2021
13	20050306	03-06-2005	March 2005	200503	Q1 2005
14	20120219	02-19-2012	February 2012	201202	Q1 2012
15	20120205	02-05-2012	February 2012	201202	Q1 2012
16	20140216	02-16-2014	February 2014	201402	Q1 2014
17	20160306	03-06-2016	March 2016	201603	Q1 2016
18	20160228	02-28-2016	February 2016	201602	Q1 2016
19	20210214	02-14-2021	February 2021	202102	Q1 2021
20	20140119	01-19-2014	January 2014	201401	Q1 2014
21	20140302	03-02-2014	March 2014	201403	Q1 2014
22	20220313	03-13-2022	March 2022	202203	Q1 2022
23	20020324	03-24-2002	March 2002	200203	Q1 2002
24	20020317	03-17-2002	March 2002	200203	Q1 2002
25	20030223	02-23-2003	February 2003	200302	Q1 2003
26	20220109	01-09-2022	January 2022	202201	Q1 2022
27	20140316	03-16-2014	March 2014	201403	Q1 2014
28	20230129	01-29-2023	January 2023	202301	Q1 2023
29	20040222	02-22-2004	February 2004	200402	Q1 2004
30	20040215	02-15-2004	February 2004	200402	Q1 2004
31	20120318	03-18-2012	March 2012	201203	Q1 2012
32	20120212	02-12-2012	February 2012	201202	Q1 2012
33	20030112	01-12-2003	January 2003	200301	Q1 2003

**Zava\_LH3**  
SQL analytics endpoint

Search

- About
- SQL endpoint
- Copilot
- SQL audit logs
- Data access mode (preview)**

**Data access mode (preview)**

Manage data access mode

Use OneLake security for tables (User's identity access mode)

Manage table permissions in OneLake security, which applies across Fabric. Manage all other object permissions using SQL security, which applies only to the SQL analytics endpoint. The user's identity will be used to access data. Permission control for each object type is shown in the following table.

Object type	Set permissions with
Table	OneLake security
View	SQL security
Stored Procedure	SQL security
Function	SQL security

Use SQL security for all objects (Delegated identity access mode)

Manage all object permissions in SQL security. A delegated identity (usually the item owner's identity) will be used to access data. Permission control for each object type is shown in the following table.

Apply

The user's identity is being used to access data.

Home Security Management Help

View data access mode (preview) Manage OneLake data access Manage SQL security

**Explorer**

- Warehouses
- Zava\_LH3
  - Schemas
    - dbo
      - Tables
        - Channel\_Se
        - Customer
        - Date**
        - Outdoor\_Si
        - Product
        - Product\_Ca
        - Product\_Su
      - Views
      - Functions
      - Stored Proce...
      - INFORMATION\_...
      - queryinsights
      - sys
    - Security
  - Queries
    - My queries
    - Shared queries

**Data preview - Date**

ID	Date_ID	Day_Name	Month_Name	Month_Of_Y...	Quarter_Na...	Qua
1	20120325	03-25-2012	March 2012	201203	Q1 2012	1
2	20200216	02-16-2020	February 2020	202002	Q1 2020	1
3	20200112	01-12-2020	January 2020	202001	Q1 2020	1
4	20200329	03-29-2020	March 2020	202003	Q1 2020	1
5	20190127	01-27-2019	January 2019	201901	Q1 2019	1
6	20190217	02-17-2019	February 2019	201902	Q1 2019	1
7	20110313	03-13-2011	March 2011	201103	Q1 2011	1
8	20010128	01-28-2001	January 2001	200101	Q1 2001	1
9	20010121	01-21-2001	January 2001	200101	Q1 2001	1
10	20160124	01-24-2016	January 2016	201601	Q1 2016	1
11	20140209	02-09-2014	February 2014	201402	Q1 2014	1
12	20210221	02-21-2021	February 2021	202102	Q1 2021	1
13	20050306	03-06-2005	March 2005	200503	Q1 2005	1
14	20120219	02-19-2012	February 2012	201202	Q1 2012	1
15	20120205	02-05-2012	February 2012	201202	Q1 2012	1
16	20140216	02-16-2014	February 2014	201402	Q1 2014	1
17	20160306	03-06-2016	March 2016	201603	Q1 2016	1
18	20160228	02-28-2016	February 2016	201602	Q1 2016	1
19	20210214	02-14-2021	February 2021	202102	Q1 2021	1
20	20140119	01-19-2014	January 2014	201401	Q1 2014	1
21	20140302	03-02-2014	March 2014	201403	Q1 2014	1
22	20220313	03-13-2022	March 2022	202203	Q1 2022	1
23	20020324	03-24-2002	March 2002	200203	Q1 2002	1
24	20020317	03-17-2002	March 2002	200203	Q1 2002	1
25	20030223	02-23-2003	February 2003	200302	Q1 2003	1
26	20220109	01-09-2022	January 2022	202201	Q1 2022	1
27	20140316	03-16-2014	March 2014	201403	Q1 2014	1
28	20230129	01-29-2023	January 2023	202301	Q1 2023	1
29	20040222	02-22-2004	February 2004	200402	Q1 2004	1
30	20040215	02-15-2004	February 2004	200402	Q1 2004	1
31	20120318	03-18-2012	March 2012	201203	Q1 2012	1
32	20120212	02-12-2012	February 2012	201202	Q1 2012	1
33	20030112	01-12-2003	January 2003	200301	Q1 2003	1

Copy SQL connection string Succeeded (1 sec 497 ms)

All Lakehouse roles > **DefaultReader**

**DefaultReader** Edit

Permissions: **Read** | Type: **Grant**

Data in role **Members in role**

+ Add members Search

Members	Type	Added Using
Kay Unkroth	User	Permission group-ReadAll
Neil Barkhina	User	Permission group-ReadAll
Christopher Jernigan	User	Permission group-ReadAll
Christian Wade	User	Permission group-ReadAll
Rui Romano	User	Permission group-ReadAll

Uses

H3

Schemas

- dbo
- Tables
  - Channel\_Sa
  - Customer
  - Date**
  - Outdoor\_Si
  - Product
  - Product\_Ca
  - Product\_Su
- Views
- Functions
- Stored Proce...

INFORMATION\_...

queryinsights

sys

Security

s

My queries

Shared queries

Data preview - Date

123	Date_ID	ABC	Day_Name	ABC
1	20120325		03-25-2012	Ma
2	20200216		02-16-2020	Feb
3	20200112		01-12-2020	Jan
4	20200329		03-29-2020	Ma
5	20190127		01-27-2019	Jan
6	20190217		02-17-2019	Feb
7	20110313		03-13-2011	Ma
8	20010128		01-28-2001	Jan
9	20010121		01-21-2001	Jan
10	20160124		01-24-2016	Jan
11	20140209		02-09-2014	Feb
12	20210221		02-21-2021	Feb
13	20050306		03-06-2005	Ma
14	20120219		02-19-2012	Feb
15	20120205		02-05-2012	Feb
16	20140216		02-16-2014	Feb
17	20160306		03-06-2016	Ma
18	20160228		02-28-2016	Feb
19	20210214		02-14-2021	Feb
20	20140119		01-19-2014	Jan
21	20140302		03-02-2014	Ma
22	20220313		03-13-2022	Ma
23	20020324		03-24-2002	Ma
24	20020317		03-17-2002	Ma
25	20030223		02-23-2003	Feb
26	20220109		01-09-2022	Jan
27	20140316		03-16-2014	Ma
28	20230129		01-29-2023	Jan
29	20040222		02-22-2004	Feb
30	20040215		02-15-2004	Feb
31	20120318		03-18-2012	Ma

Direct Lake semantic model name ⓘ

Zava\_SM

Workspace ⓘ

Zava Retail

Storage mode

Direct Lake on OneLake ⓘ  Direct Lake on SQL ⓘ

Select or deselect tables for the semantic model.

Search

SQL views cannot be selected for Direct Lake on OneLake, but tables based on SQL views can be added to the semantic model later using other storage modes.

Select all

- dbo
  - Channel\_Sales
  - Customer
  - Date
  - Outdoor\_Sales
  - Product
  - Product\_Category
  - Product\_Subcategory

Learn more about Direct Lake ↗

Confirm Cancel

Showing 1000 rows

123	Year	Date	ABC	Week_Day	0/1	Weekend_FL...
2012	2012	2012-03-25 00:0...		Sunday	1	
2020	2020	2020-02-16 00:0...		Sunday	1	
2020	2020	2020-01-12 00:0...		Sunday	1	
2020	2020	2020-03-29 00:0...		Sunday	1	
2019	2019	2019-01-27 00:0...		Sunday	1	
2019	2019	2019-02-17 00:0...		Sunday	1	
2011	2011	2011-03-13 00:0...		Sunday	1	
2001	2001	2001-01-28 00:0...		Sunday	1	
2001	2001	2001-01-21 00:0...		Sunday	1	
2016	2016	2016-01-24 00:0...		Sunday	1	
2014	2014	2014-02-09 00:0...		Sunday	1	
2021	2021	2021-02-21 00:0...		Sunday	1	
2005	2005	2005-03-06 00:0...		Sunday	1	
2012	2012	2012-02-19 00:0...		Sunday	1	
2012	2012	2012-02-05 00:0...		Sunday	1	
2014	2014	2014-02-16 00:0...		Sunday	1	
2016	2016	2016-03-06 00:0...		Sunday	1	
2016	2016	2016-02-28 00:0...		Sunday	1	
2021	2021	2021-02-14 00:0...		Sunday	1	
2014	2014	2014-01-19 00:0...		Sunday	1	
2014	2014	2014-03-02 00:0...		Sunday	1	
2022	2022	2022-03-13 00:0...		Sunday	1	
2002	2002	2002-03-24 00:0...		Sunday	1	
2002	2002	2002-03-17 00:0...		Sunday	1	
2003	2003	2003-02-23 00:0...		Sunday	1	
2022	2022	2022-01-09 00:0...		Sunday	1	
2014	2014	2014-03-16 00:0...		Sunday	1	
2023	2023	2023-01-29 00:0...		Sunday	1	
2004	2004	2004-02-22 00:0...		Sunday	1	
2004	2004	2004-02-15 00:0...		Sunday	1	
2012	2012	2012-03-18 00:0...		Sunday	1	

...

**Date** ...

- Date
- Σ Date\_ID
- Day\_Name
- LastRefreshDate
- Month\_Name
- Σ Month\_Of\_Year
- Quarter\_Name
- Σ Quarter\_Of\_Year
- Week\_Day
- Collapse ^

**Outdoor\_Sales** ...

- Carrier\_Tracking\_Number
- Σ Currency\_Key
- Σ Customer\_Key
- Customer\_PO\_Number
- Σ Discount\_Amount
- Due\_Date
- Σ Due\_Date\_Key
- Σ Extended\_Amount
- Σ Freight
- Collapse ^

**Product** ...

- Arabic\_Description
- Chinese\_Description
- Class
- Color
- Σ Days\_To\_Manufacture
- Σ Dealer\_Price
- End\_Date
- English\_Description
- English Product Name
- Collapse ^

### Properties

General

**Name**  
Date

**Description**  
Enter a description

**Row label**  
Select a row label

**Is hidden**  
 No

Advanced

**Storage mode**  
Direct Lake

- Import
- Direct Lake

### Data

Tables Model

View and organize all of the items in your semantic model. [Learn how](#)

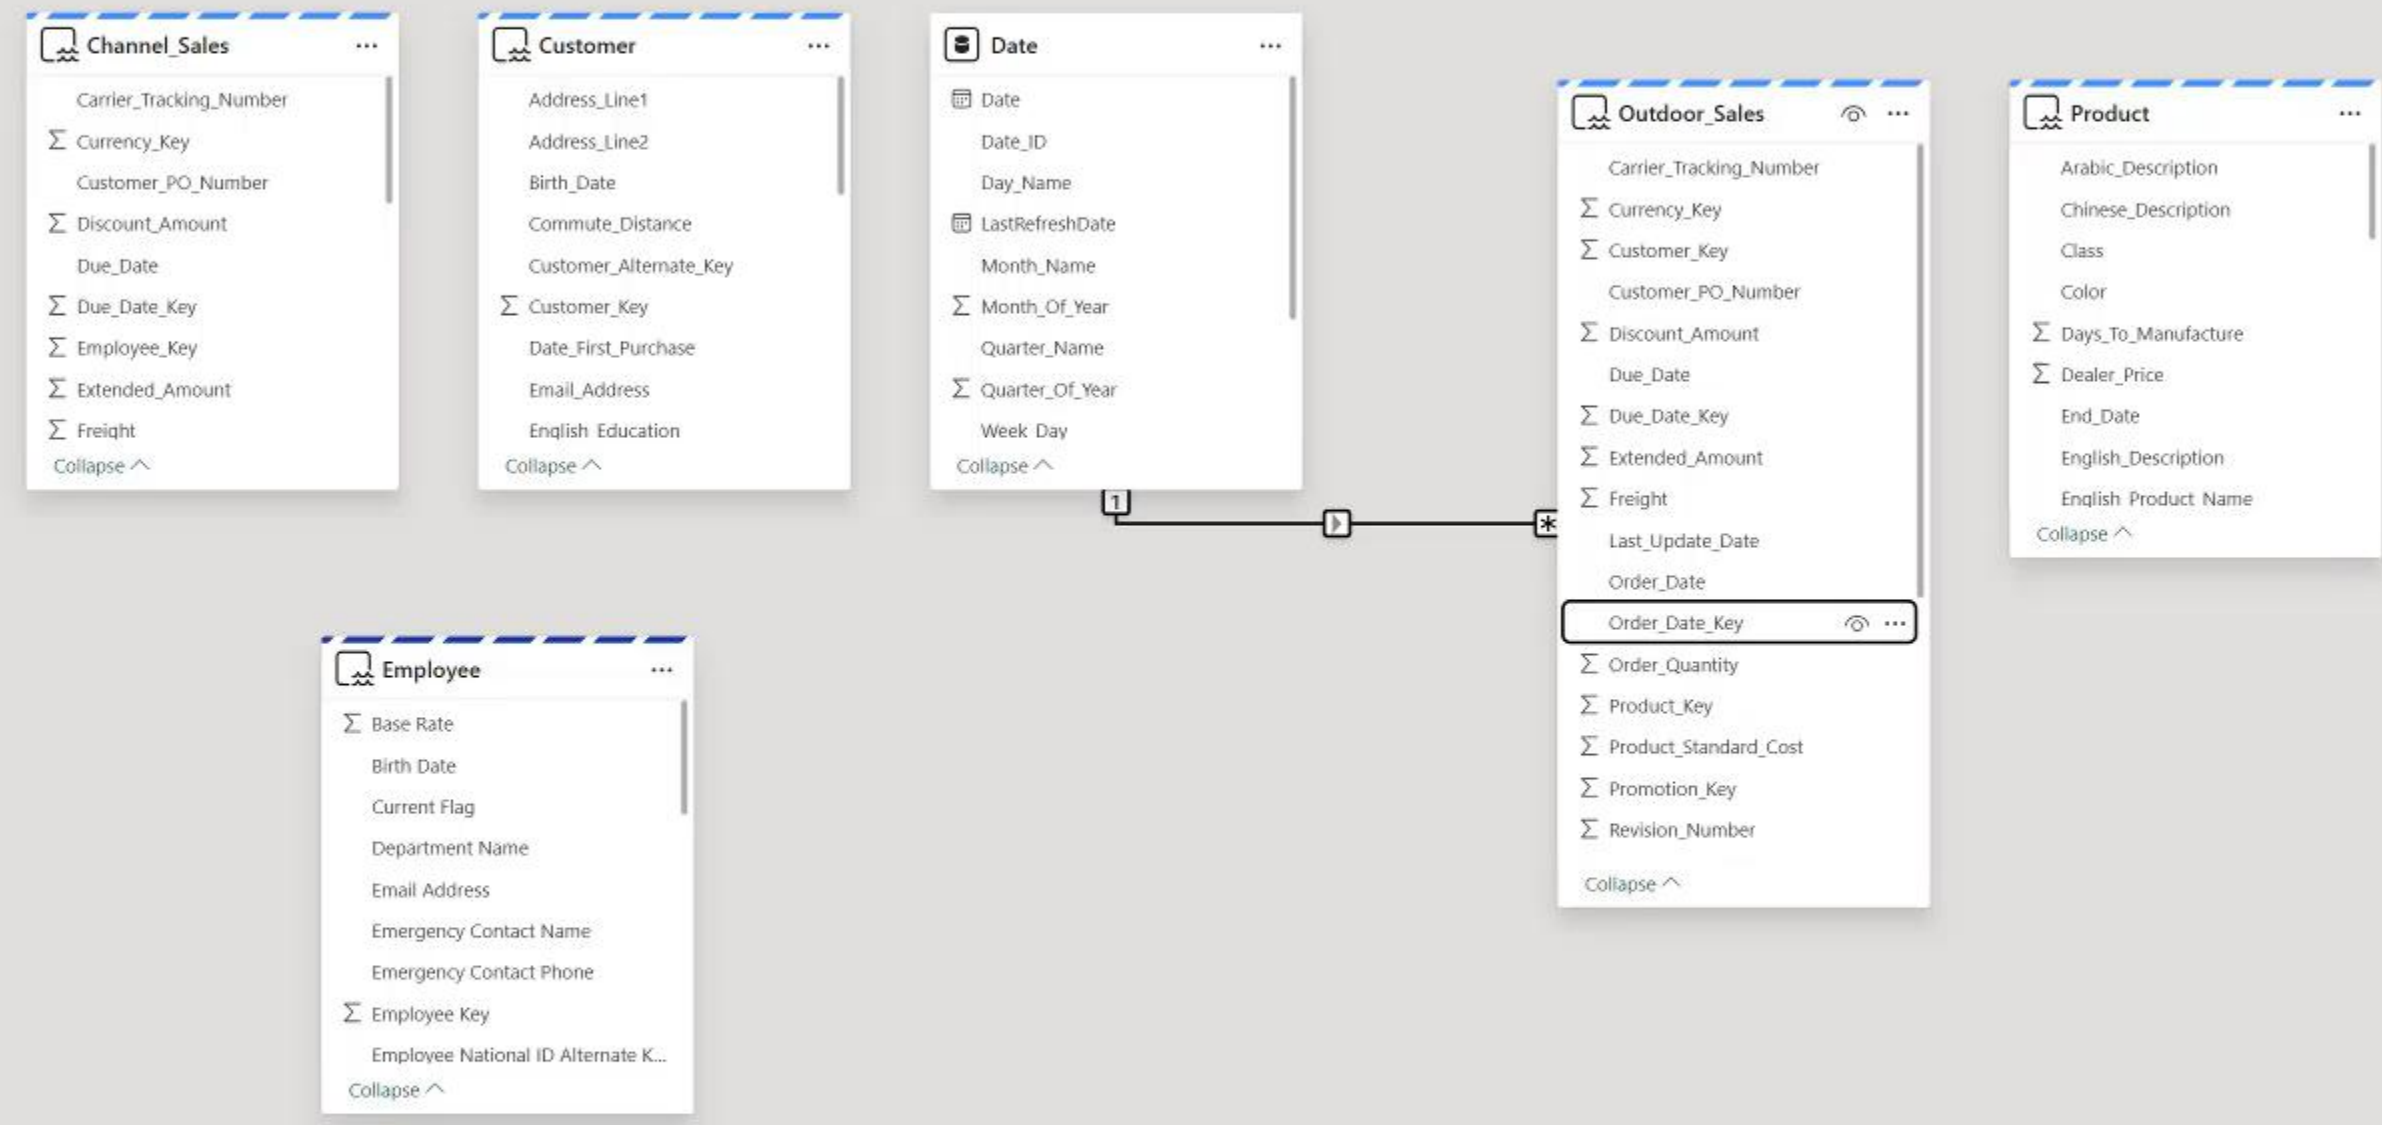
Search

- Semantic model
  - Calculation groups (0)
  - Cultures (0)
  - Expressions (2)
  - Functions (0)
  - Measures (0)
  - Perspectives (0)
  - Relationships (0)
  - Roles (0)
  - Tables (6)
    - Channel\_Sales
    - Customer
    - Date**
    - Employee
    - Outdoor\_Sales
    - Product

- Home
- Copilot
- Create
- OneLake catalog
- Apps
- Databases
- Monitor
- Learn
- Real-Time
- Workloads
- Workspaces
- Zava Retail
- Zava\_SM
- Zava\_LH3
- Zava\_LH3
- Power BI

File Home Help

Get data OneLake catalog Transform Refresh data New measure New column New table Calculation group Calendar options New parameter Manage roles Manage permissions Manage relationships Edit tables Best practice analyzer Memory analyzer Community notebooks Generate Ontology New report Explore Lineage Analyze in Excel Prep data for Copilot AI



**Properties**

Relationship

Table: Outdoor\_Sales Column: Order\_Date\_Key

Cardinality: Many to one (\*:1)

Table: Date Column: Date\_ID

Make this relationship active:  Yes

Cross-filter direction: Single

Apply security filter in both directions:  No

Buttons: Apply changes, Open relationship editor

**Data**

Tables Model

View and organize all of the items in your semantic model. [Learn how](#)

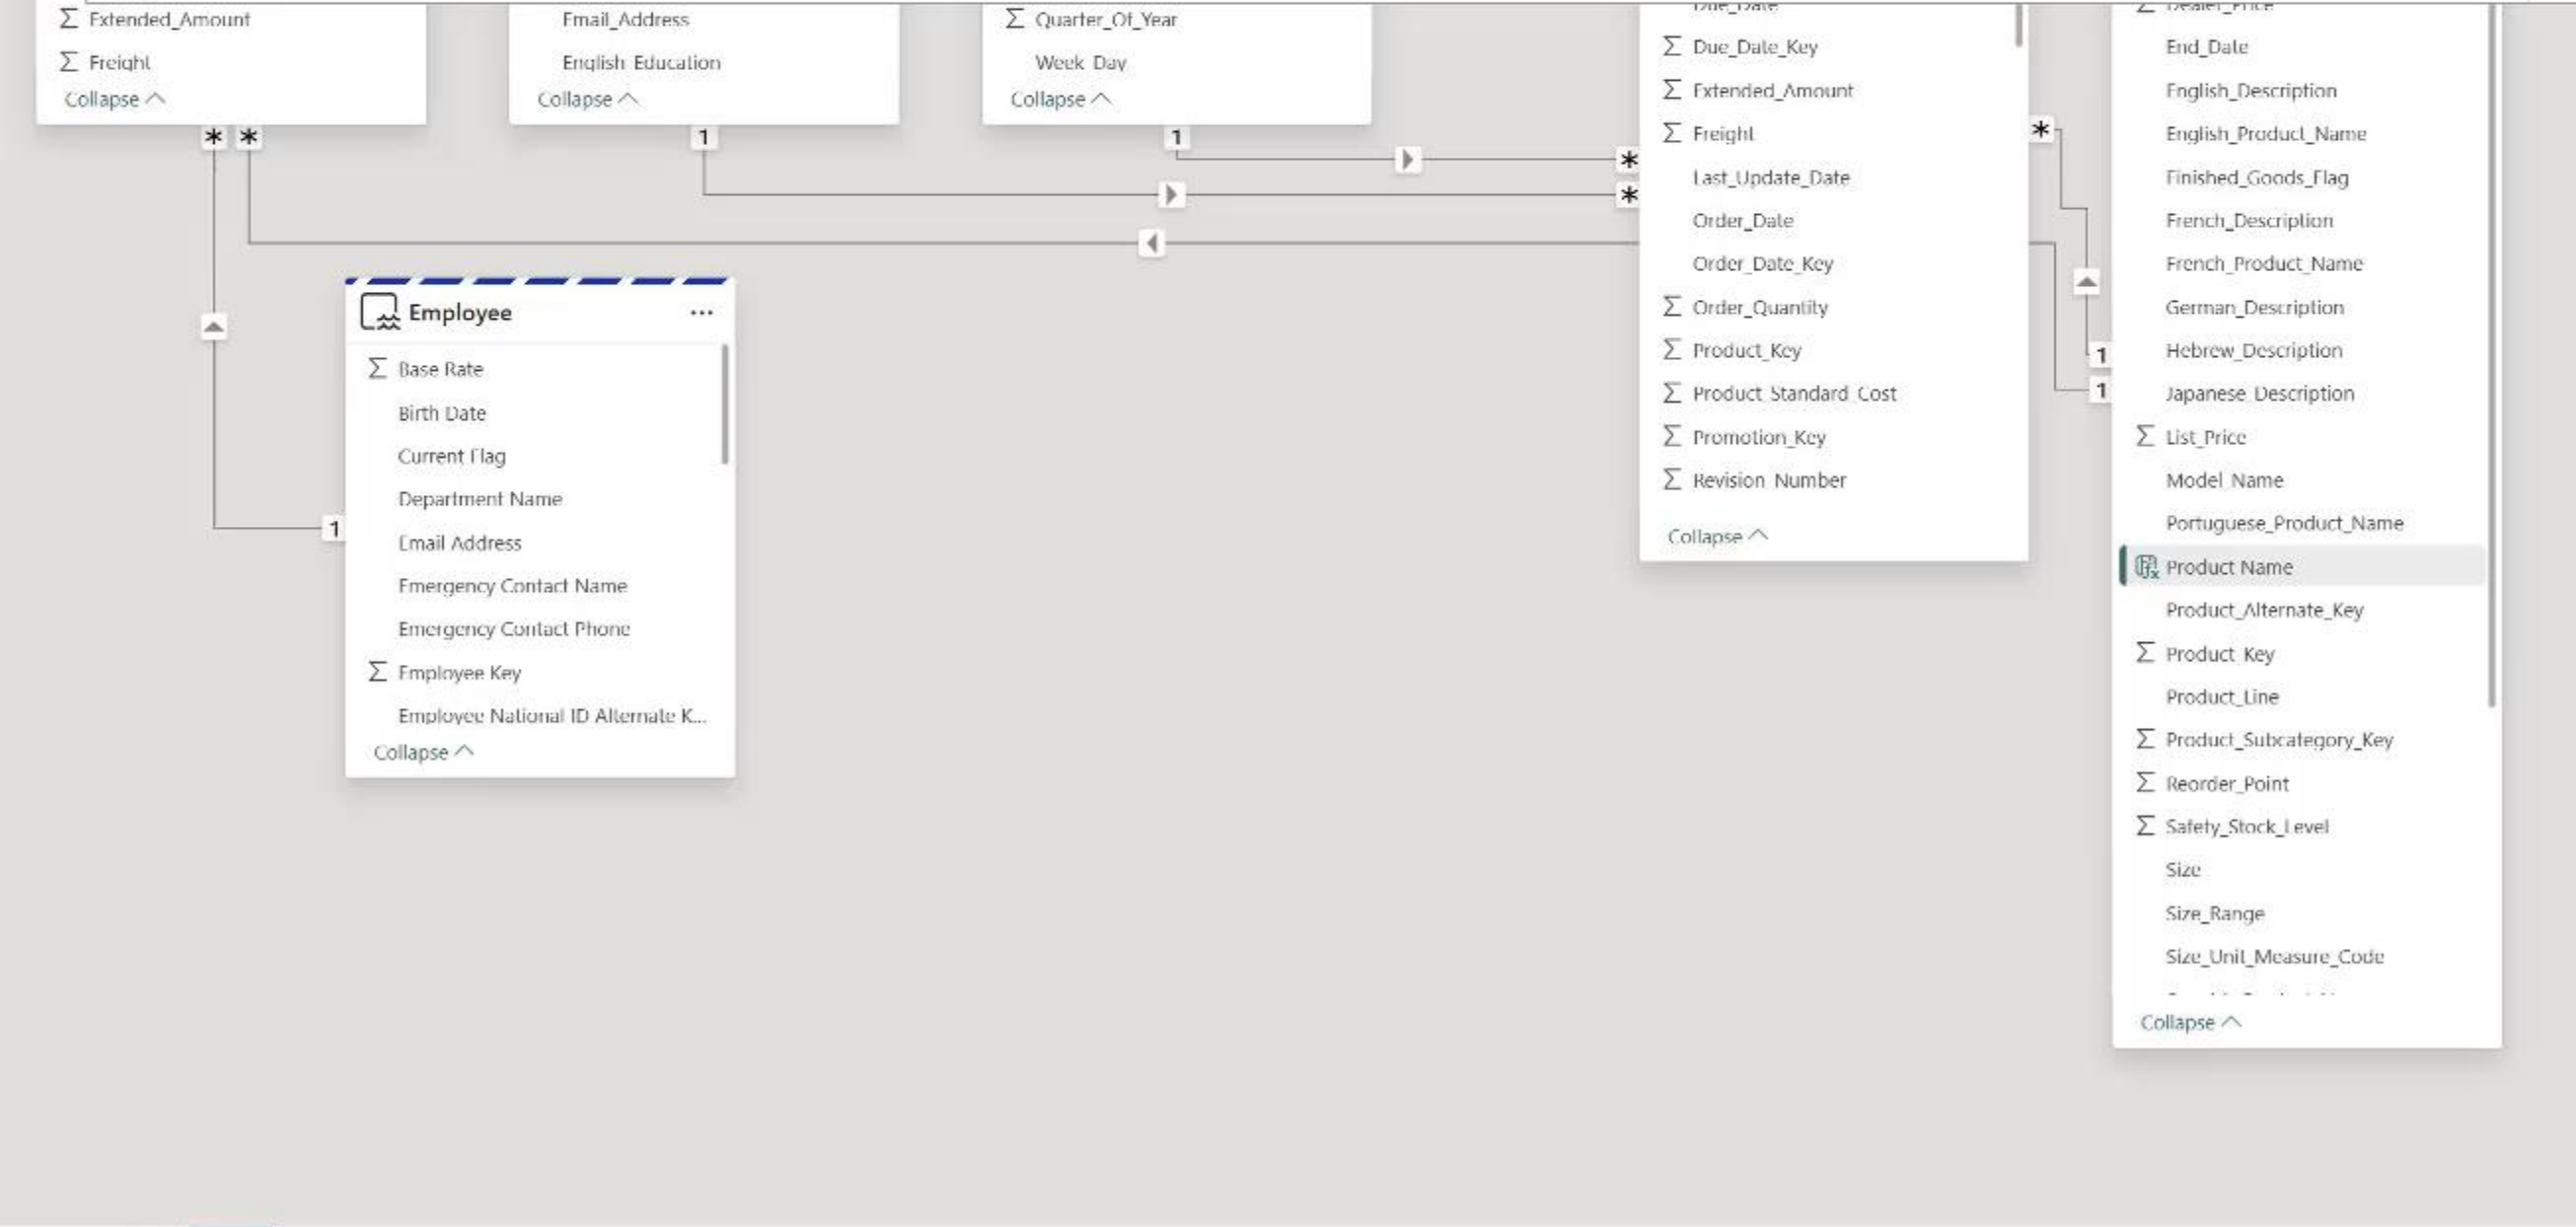
Search

Semantic model

- Calculation groups (0)
- Cultures (0)
- Expressions (2)
- Functions (0)
- Measures (0)
- Perspectives (0)
- Relationships (1)
  - Outdoor\_Sales[Order\_Date\_Key] <...
- Roles (0)
- Tables (6)
  - Channel\_Sales
  - Customer
  - Date
  - Employee
  - Outdoor\_Sales
  - Product

```

1 Product Name =
2 VAR Culture = USERCULTURE()
3 RETURN
4 SWITCH (
5     Culture,
6     "en-US", 'Product'[English_Product_Name],
7     "fr-FR", 'Product'[French_Product_Name],
8     "es-ES", 'Product'[Spanish_Product_Name],
9     "pt-BR", 'Product'[Portuguese_Product_Name],
10    'Product'[English_Product_Name] -- default
11 )
    
```



### Properties

Enter a description

Display folder: Enter the display folder

Is hidden:  No

### Formatting

Data type: Text

Format: Text

### Advanced

Sort by column: Product Name (Default)

Data category: Uncategorized

Summarize by: None

### Expression Context

User Context

Standard

User Context

### Data

Tables | **Model**

View and organize all of the items in your semantic model. [Learn how](#)

Search

- Σ Dealer\_Price
- End\_Date
- English\_Description
- English\_Product\_Name
- Finished\_Goods\_Flag
- French\_Description
- French\_Product\_Name
- German\_Description
- Hebrew\_Description
- Japanese\_Description
- Σ List\_Price
- Model\_Name
- Portuguese\_Product\_Name
- Product Name**
- Product\_Alternate\_Key
- Σ Product\_Key
- Product\_Line
- Σ Product\_Subcategory\_Key
- Σ Reorder\_Point
- Σ Safety\_Stock\_Level
- Size
- Size\_Range
- Size\_Unit\_Measure\_Code
- Spanish\_Product\_Name
- Σ Standard\_Cost
- Start\_Date
- Status

# Fundamentals & Lessons Learned

What we've learnt from enterprise Direct Lake  
deployments

# The Question Every Enterprise Asks

- You've built your Lakehouse
- You've got your delta tables
- Now you want a semantic model on top

What do you need to know to set yourself up for success?

# Three Performance Dimensions

- Cold cache paging — how long to page columns into the model?
- Model size — how big does it get? (determines your Fabric SKU)
- DAX query compute — how much compute to scan those columns?

Example: if your columns combined are ~70 GB,  
you can't use Fabric F64 — you'll need a much larger capacity

# Three Things That Influence All of These

- Cardinality
- V-Order
- Data Layout

Get these right and you set yourself up for success

Get them wrong and performance will suffer

# 1. Cardinality

The number of unique values in a column

Low cardinality (e.g. Country, Status, Category):

Pages in fast, small memory footprint

Even with billions of rows

High cardinality (e.g. Transaction ID, Timestamp, GUID):

Slower to page in, more memory consumed


More compute for DAX queries to scan

Fruit	ID
Apple	1
Banana	2
Cherry	3

Fruit	ID
Apple	17
Banana	9
Cherry	1

Fruit	ID
Apple	99
Banana	44
Cherry	55

PARQUET DICTIONARIES



MODEL DICTIONARY

Fruit	ID
Apple	1
Banana	2
Cherry	3
Date	4
Eggfruit	5
Feijoa	6

# Cardinality — What Can You Do?

- Review data types
  - avoid float/double when you don't need precision
  - e.g. 7.45 instead of 7.4523462734
- Consider splitting high-cardinality columns
  - e.g. DateTime split into separate Date and Time columns
- Remove columns you don't need in the model
- Use Delta Analyzer to identify the worst offenders before you build

## 2. V-Order

Microsoft-proprietary optimization for Parquet files

Uses the same compression algorithms as Power BI Import mode

V-Ordered Delta tables:

- Smaller on disk (less I/O for all workloads)

- Faster to page into Direct Lake

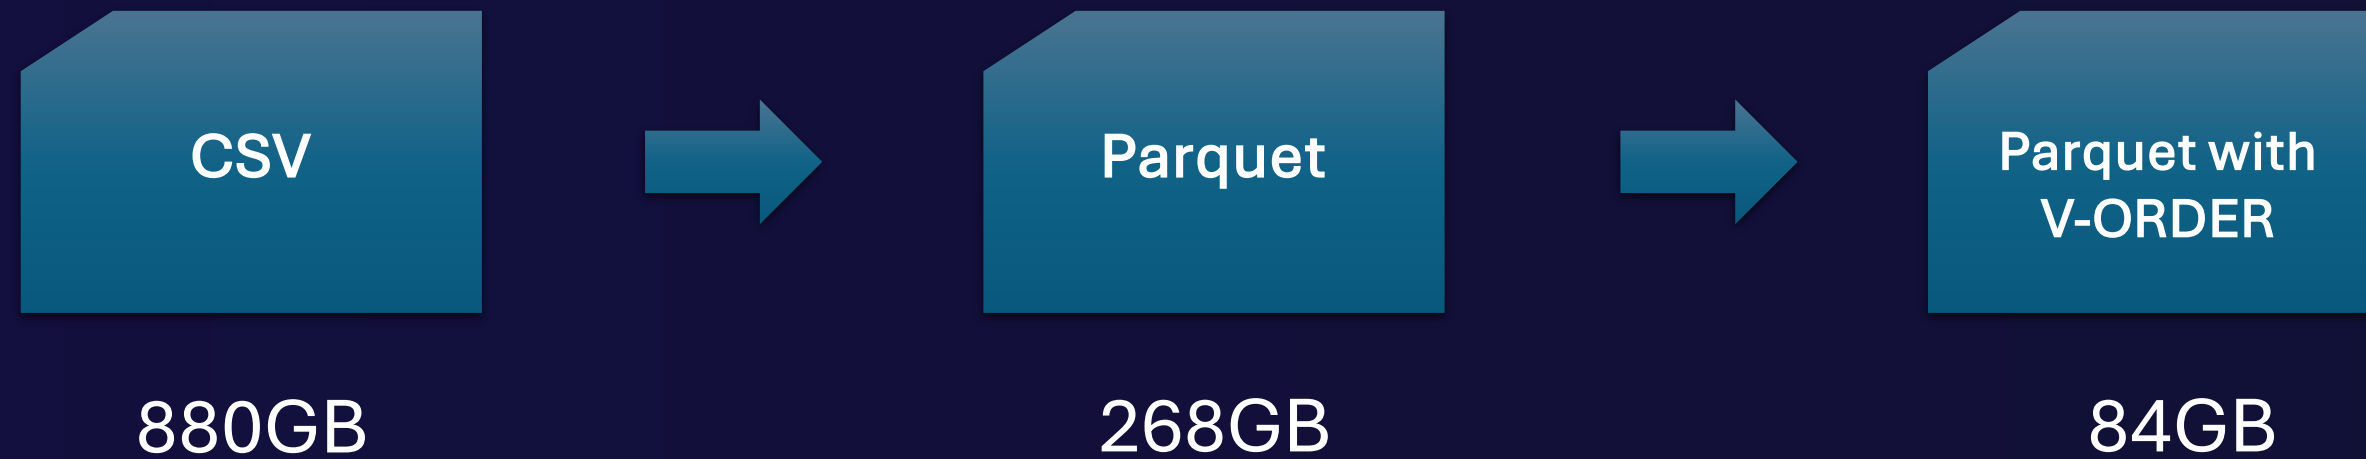
- Less compute for DAX queries

Only Fabric processing can apply V-Order

This is why using Fabric tooling matters

# Priorities : V-Order

Sales Data (162 tables)



**x3.2**

Less I/O for all workloads

# V-Order — The Numbers

Microsoft Sales Data (162 tables)

CSV: 880 GB

Parquet: 268 GB

Parquet with V-Order: 84 GB

3.2x less I/O for all workloads

..... But should we apply V-Order everywhere?

# Seamarks V-Order Maxim

V-Order Benefit = # Reads / # Writes

# Seamarks V-Order Maxim

Bronze / Silver layers:

Written once, read once or twice

Ratio is low, V-Order cost outweighs benefit

Gold layer / semantic model source:

Written occasionally, read by every DAX query

Ratio is very high, V-Order benefit is massive

The higher the number, the more benefit you get from V-Order

DAX queries count as reads — they page V-Ordered data into memory

# Configure resource profiles

- Configure per environment

Create a new environment or edit an existing one.

set `spark.fabric.resourceProfile` to one of the following values:

`writeHeavy`

`readHeavyForPBI`

`readHeavyForSpark`

A custom profile name

## Configure at runtime

```
spark.conf.set("spark.fabric.resourceProfile", "readHeavyForSpark")
```

# 3. Data Layout

Some writes can produce lots of tiny Parquet files

- Depends on how often you update the upstream source

- Frequent intraday updates can produce many small files

Why it matters for Direct Lake:

- More files = more overhead when paging

- Fragmented rowgroups impact scan performance

Direct Lake works on any delta table

- But performance depends on how the data is laid out

# Data Layout — The Gold Layer Pattern

Create a layer optimised for Direct Lake:

- Consolidated files (fewer, larger Parquet files)

- Run OPTIMIZE() regularly

- V-Order applied by Fabric processing

Point your semantic model at the final layer

Important: only do this where performance is critical

Don't over-engineer models that don't need it

# How Do I Check All of This?

- Cardinality per column?
- V-Order status?
- File and rowgroup layout?
- Estimated model size?

Delta Analyzer

Available in Semantic Link Labs (Fabric Notebooks)

# Delta Analyzer

Point it at a delta table in your Lakehouse

It reports:

- Column cardinality (find the worst offenders)

- V-Order status (is it applied?)

- File and rowgroup layout (fragmentation?)

- Estimated semantic model size

- DQ fallback risk (will you hit guardrails?)

Check your data BEFORE you build your model -Works over shortcuts

# Demo

## Delta Analyzer

# A Pattern Worth Considering

Direct Lake over OneLake enables composite models

Dimension tables in Import mode

- Small, lightweight, always hydrated

- No paging delays

- Easy to extend with calculated columns

Fact tables in Direct Lake mode

- These are the large tables

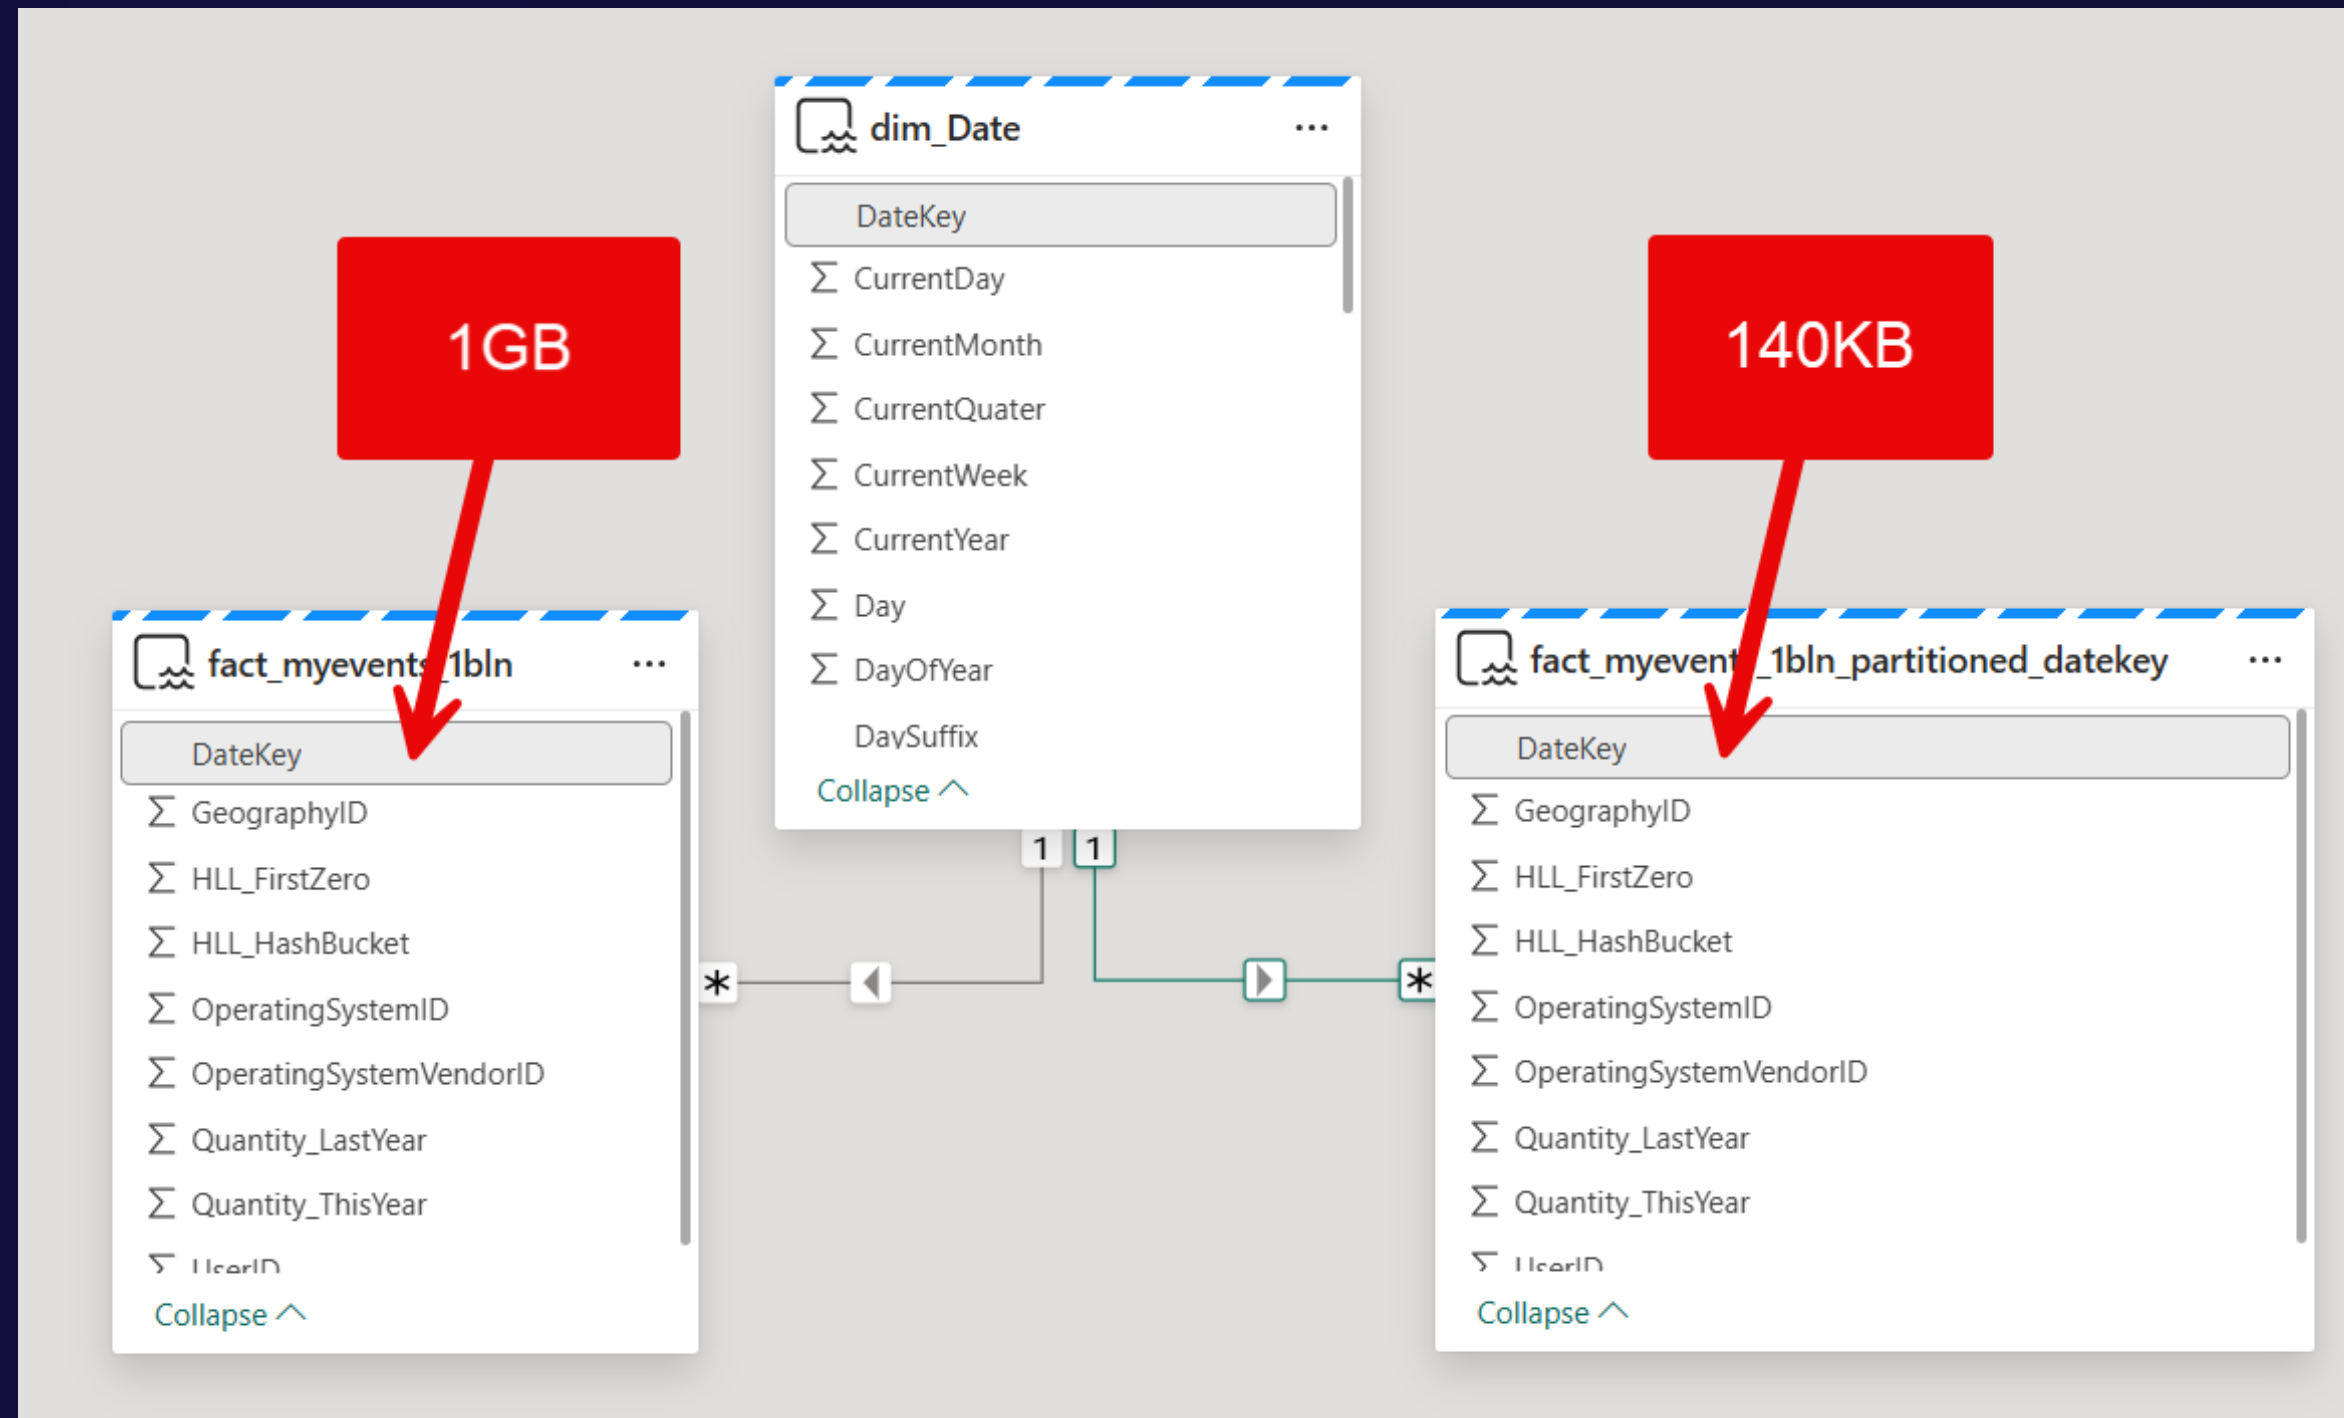
- This is where Direct Lake's paging model shines

# Performance Tip!

# Partition Lakehouse table by DateKey

- Partition your Fact Delta Table by DateKey
  - Do this in the lake house
  - Just for the Delta Table used by the Semantic Model
  - Use the FK column used to connect to Date Dimension
- Keep your eye on guardrails!!

# Partition Lakehouse table by DateKey



# Performance numbers – Warm Cache

	V-Order	No V-Order	Partitioned (V-Order)
Query 1			
Total Time	223	863	47
FE CPU	1,449	12,840	550
Query 2			
Total Time	1,594	2,891	94
FE CPU	11,890	39,125	1,070
Query 3			
Total Time	4,817	4,729	5,851
FE CPU	29,937	43,933	42,523

# Why This Matters More Than Ever

Semantic models *ARE* your AI-ready data layer

Clearly defined metrics and KPIs

Business-friendly object names

Inferred relationships

AI can look across semantic models

A well-structured model isn't just good for reports

It's the foundation for trustworthy AI answers

# What You Can Try on Monday

Three things to check:

1. Cardinality — find and address high-cardinality columns
2. V-Order — make sure it's applied (use Fabric tooling)
3. Data layout — consolidate fragmented files (OPTIMIZE)

One tool to check them all:

Delta Analyzer (Semantic Link Labs)

One pattern to consider:

Import dims + Direct Lake facts

# Questions?

Sound off.  
The mic is all yours.  
Influence the product roadmap.

Join the Fabric User Panel



Share your feedback directly with our Fabric product group and researchers.

<https://aka.ms/JoinFabricUserPanel>

Join the SQL User Panel



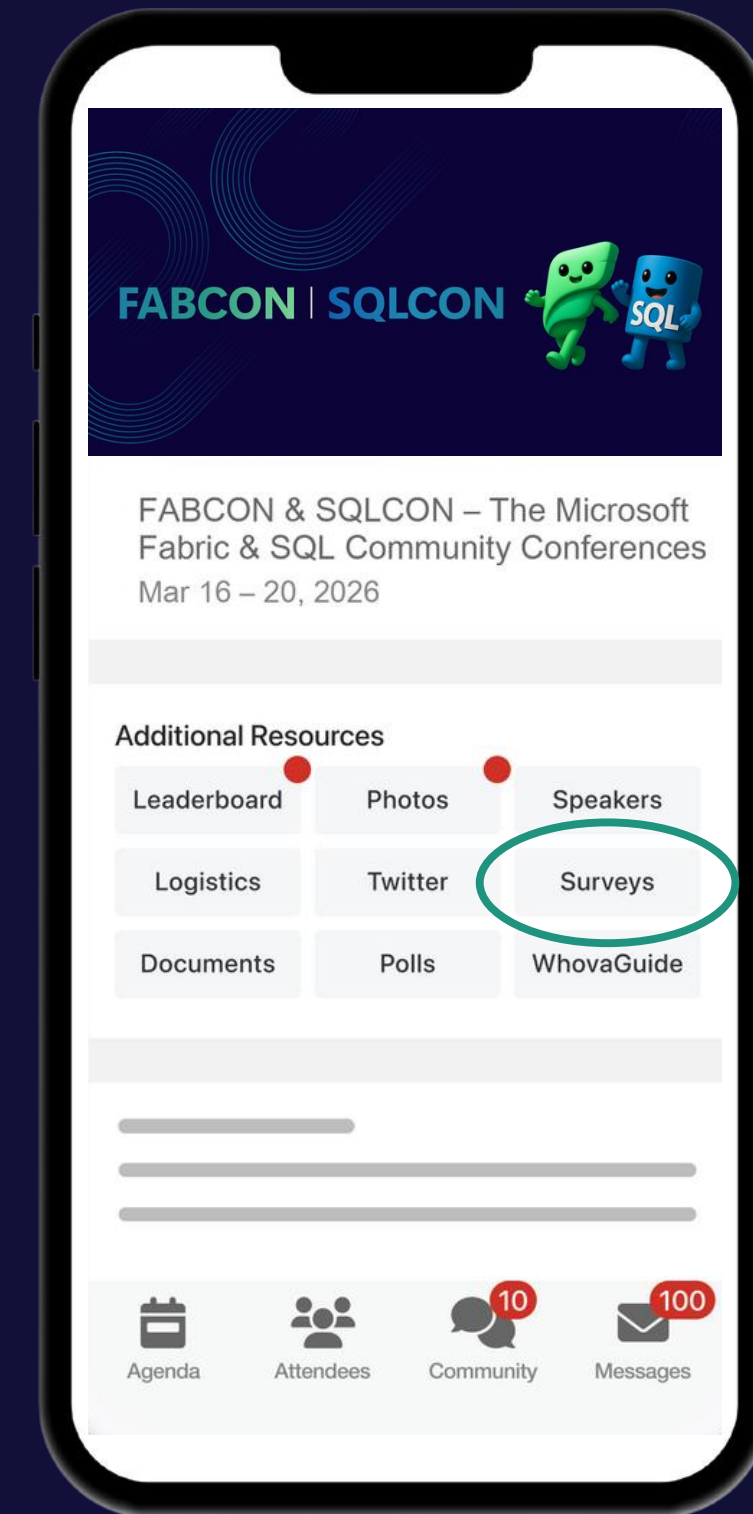
Influence our SQL roadmap and ensure it meets your real-life needs

<https://aka.ms/JoinSQLUserPanel>

# How was the session?



Complete Session Surveys in  
*Whova* for your chance to WIN  
PRIZES!



# Get Two Fabric Certifications for FREE

Attendees of FABCON can take the Fabric Analytics Engineer or Fabric Data Engineer exam for free. Be part of the 2 fastest growing role-based certifications in Microsoft history.

**Request your voucher by March 23, 2026.**

<https://aka.ms/fabcon/cert100>

