www.sqlbi.com

How to choose between Direct Lake and Import mode

from May 2025 onwards

Marco Russo - SQLBI

Why so many choices?

o "It depends": every architecture has pros and cons

o Review existing architectures and their constraints

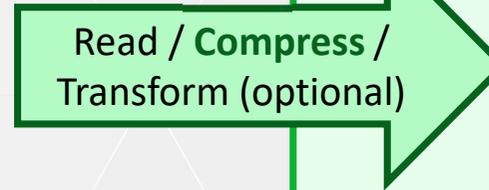o Introduce a new composite model type: **Direct Lake + Import**

Agenda

o Import mode

o Direct Lake

o Composite models (before May 2025)
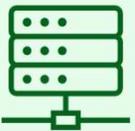
o Direct Lake + Import (from May 2025)

sqlbi.

Import mode

Import mode

Refresh
- Read data
- Compression
- Calc tables
- Calc columns

Copy / Transform

Any format
Any storage

Read / **Compress** /
Transform (optional)

Private
storage

Load in memory

VertiPaq

High compression
Best performance
Local PBIX
Zero-low maintenance

Semantic model

Raw data (bronze layer?)
- OLTP
- Web services
- Logs
- …

Refined data (gold layer?)
- Data mart
- Dimensional model
- …

sqlbi.

Key points for Import Mode

- Simple
  - Default mode for "traditional" data sources (e.g. SQL, Excel, dataflows, …)
    - Not the default for Lakehouse or Warehouse
    - Still supported with Power Query transformations
  - Support any transformation
  - All features available
- Optimized
  - Best query performance
  - Lower cost per query (CU consumption)
  - Better compression
  - Upfront planning for incremental refresh
  - Memory needed only for queried columns (only large format)

- Processing overhead
  - Time to load and compress data
  - Memory required for compression
- Internal binary storage format
  - Analysis Services format is private
  - Storage engine level non-accessible from outside
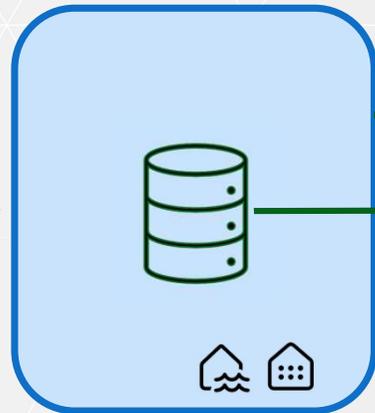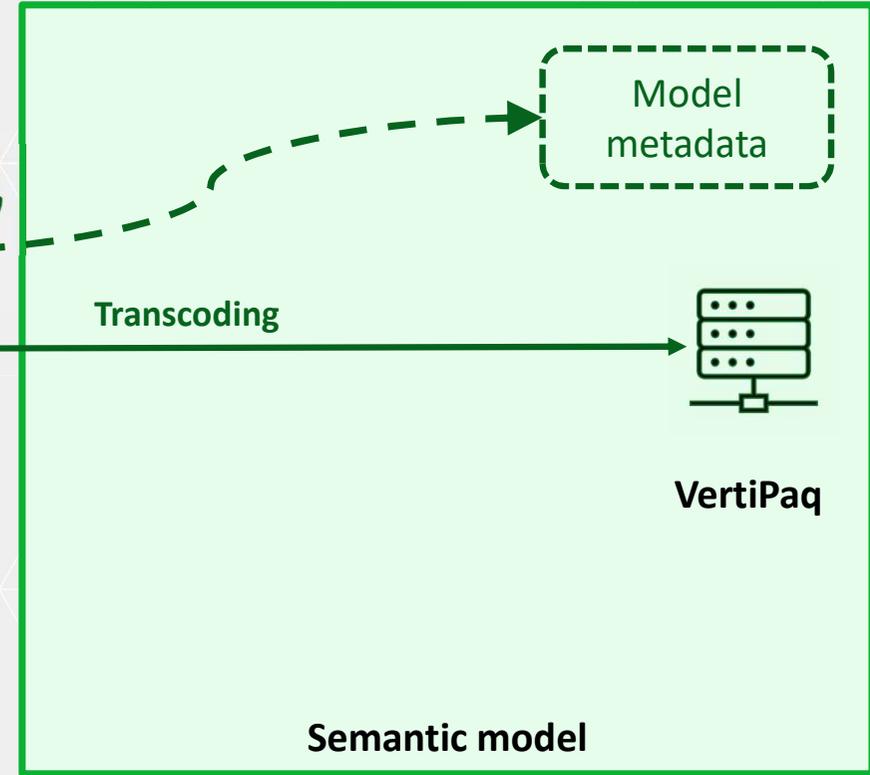 - OneLake integration exports imported data in Delta table format

sqlbi.

Direct Lake

Key points for Direct Lake

- Instant processing

  - Data storage contains compressed data
  - No data ingestion during processing
  - Public format, accessible from outside

- Transformations in semantic model

  - No Power Query → Dataflows Gen 2
  - Calculated tables → coming soon
  - Calculated columns → coming later

- Limitations for Analyze in Excel

  - User-defined hierarchies not available
  - Drillthrough is currently not supported

- Transcoding needed

  - One-time cost for columns used by queries
  - Additional cost for high cardinality columns

- Compression quality "depends"

  - Developers have complete control
    - They are "in charge"
    - Great power, great responsibility
  - Fabric tools help data ingestion
    - Default settings are good for compression
    - Pay attention to incremental updates
  - External parquet files requires more attention
    - Good or bad Parquet compression
    - Bad compression → bad query performance

sqlbi.

Flavors of Direct Lake

Direct Lake on OneLake

o   No DirectQuery fallback

o   M expression: AzureStorage.DataLake

o   Enforce OneLake security

Direct Lake on SQL endpoint

o   DirectQuery fallback

o   M expression:
    Sql.Database

o   Enforce SQL-based security

sqlbi.

Creating a Direct Lake model in Power BI Desktop

Direct Lake on OneLake

o   Open OneLake catalog

o   Use "Connect to OneLake"

Direct Lake on SQL endpoint

o   Not possible

Creating a Direct Lake model in Power BI Service

Direct Lake on OneLake

o Not possible

Direct Lake on SQL endpoint
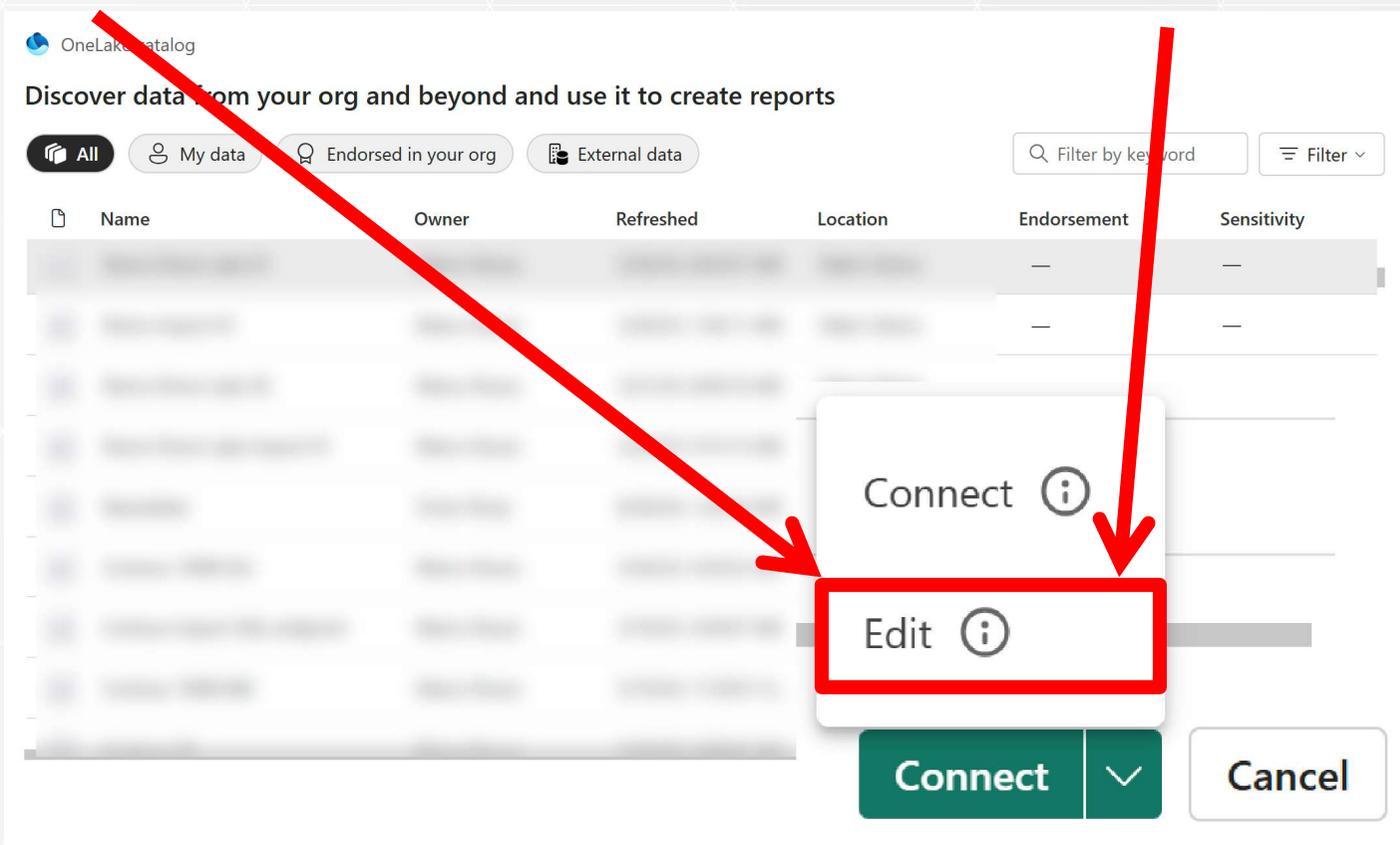
o From Lakehouse or Warehouse

Editing a Direct Lake model in Power BI Desktop

Direct Lake on OneLake

o   Connect / Edit to Semantic Model

Direct Lake on SQL endpoint

o   Connect / Edit to Semantic Model

Editing a Direct Lake model in Power BI Service

Direct Lake on OneLake

o Open data model (*)

Direct Lake on SQL endpoint

o Open data model

(*) Editing table coming soon

Importing tables from Lakehouse / Warehouse in Power BI Desktop

o   OneLake catalog → Select the Lakehouse / Warehouse (+ other Fabric sources)

o   Use "Connect to SQL endpoint"

o   Select table(s) and "Import"

Using DirectQuery to Lakehouse / Warehouse in Power BI Desktop

o OneLake catalog → Select the Lakehouse / Warehouse (+ other Fabric sources)
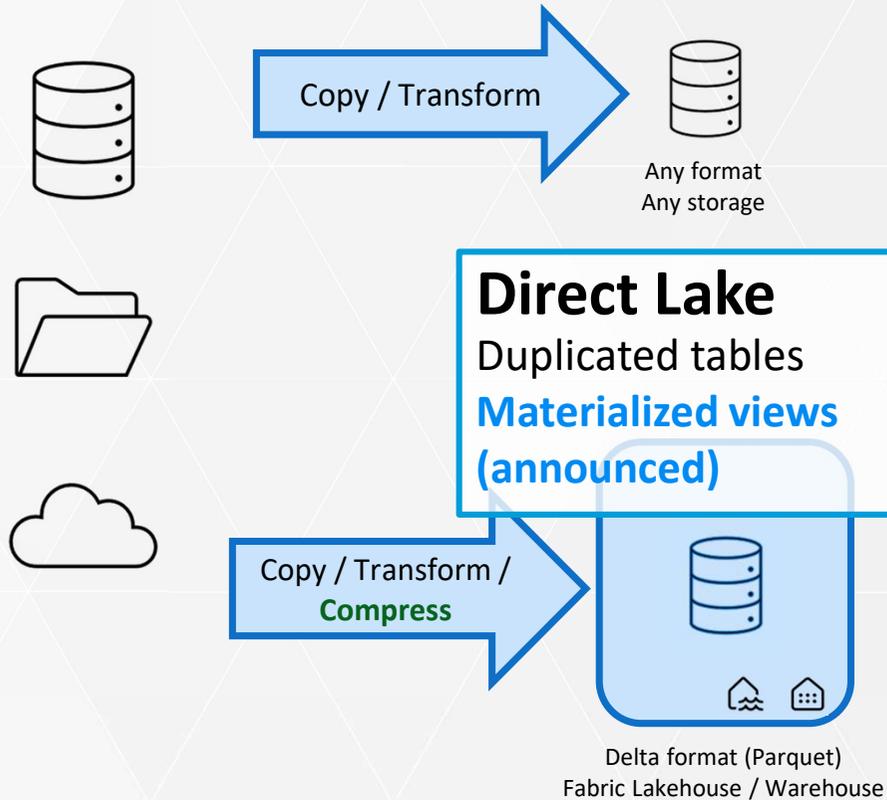
o Use "Connect to SQL endpoint"
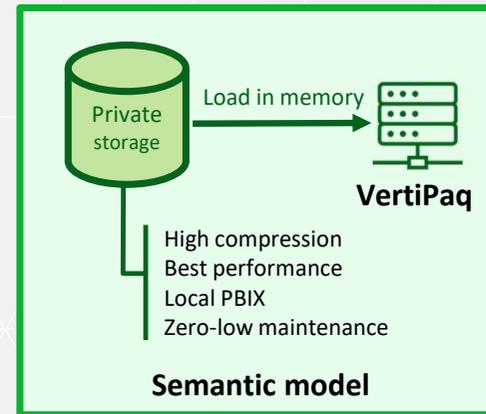
o Select table(s) and "DirectQuery"

Semantic model design

o The physical model must match the semantic model – no abstraction:

- One entity, one table
- One attribute, one column

o Semantic model features may require specific transformations:

- Sort column by (e.g.: month name sorted by month number)
- Any presentation column (e.g.: First name and last name in a single column)
- Any aesthetic request (e.g.: first three letters of month name)
- Any custom grouping (e.g.: price range)

o Impact:

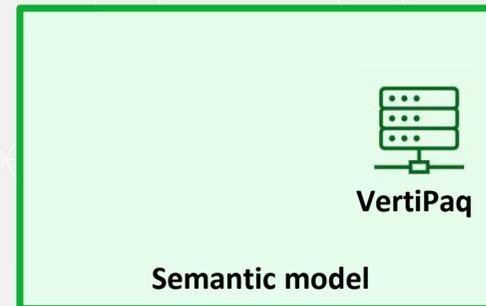- Skills required
- Organization

sqlbi.

Organizational issue

Copy / Transform

Any format
Any storage

Direct Lake
Duplicated tables
**Materialized views
(announced)**

**Copy / Transform /
Compress**

Delta format (Parquet)
Fabric Lakehouse / Warehouse

Data Engineer

Private
storage

Load in memory

VertiPaq

High compression
Best performance
Local PBIX
Zero-low maintenance

Semantic model

**SEMANTIC MODEL
TRANSFORMATIONS**

Import
SQL query
Power Query / M
Calculated columns
Calculated tables

VertiPaq

Semantic model
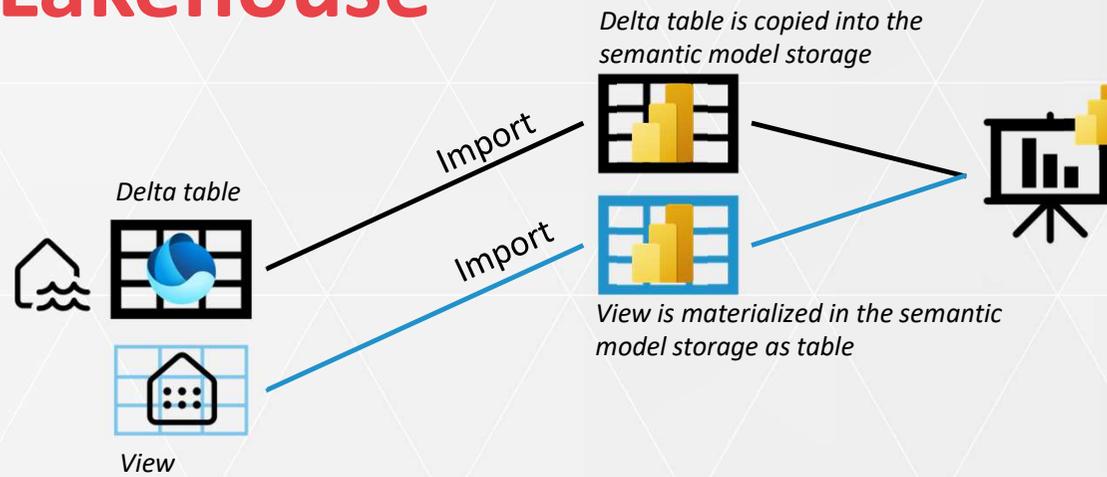
Direct Lake
**Calculated tables
(coming soon)**

**Calculated columns
(coming later)**

Semantic Model Developer

sqlbi

Views on Lakehouse

Delta table

View

Views on Lakehouse

Delta table is copied into the semantic model storage

Delta table

Import

Import

View is materialized in the semantic model storage as table

View

sqlbi.

Views on Lakehouse

Delta table is copied into the semantic model storage

Delta table

Import

View

Import

View is materialized in the semantic model storage as table

Direct Lake

Delta table stays in OneLake and data read as needed

DirectQuery

View stays **in SQL endpoint** and data read as needed

Views on Lakehouse

Views on Lakehouse

Delta table is copied into the semantic model storage

Delta table

Import

Import

View is materialized in the semantic model storage as table

View

View is materialized in the OneLake storage as delta table

Direct Lake

Direct Lake

Delta table stays in OneLake and data read as needed

Delta table stays in OneLake and data read as needed

sqlbi.

Composite models

Composite models (DQ / AS)

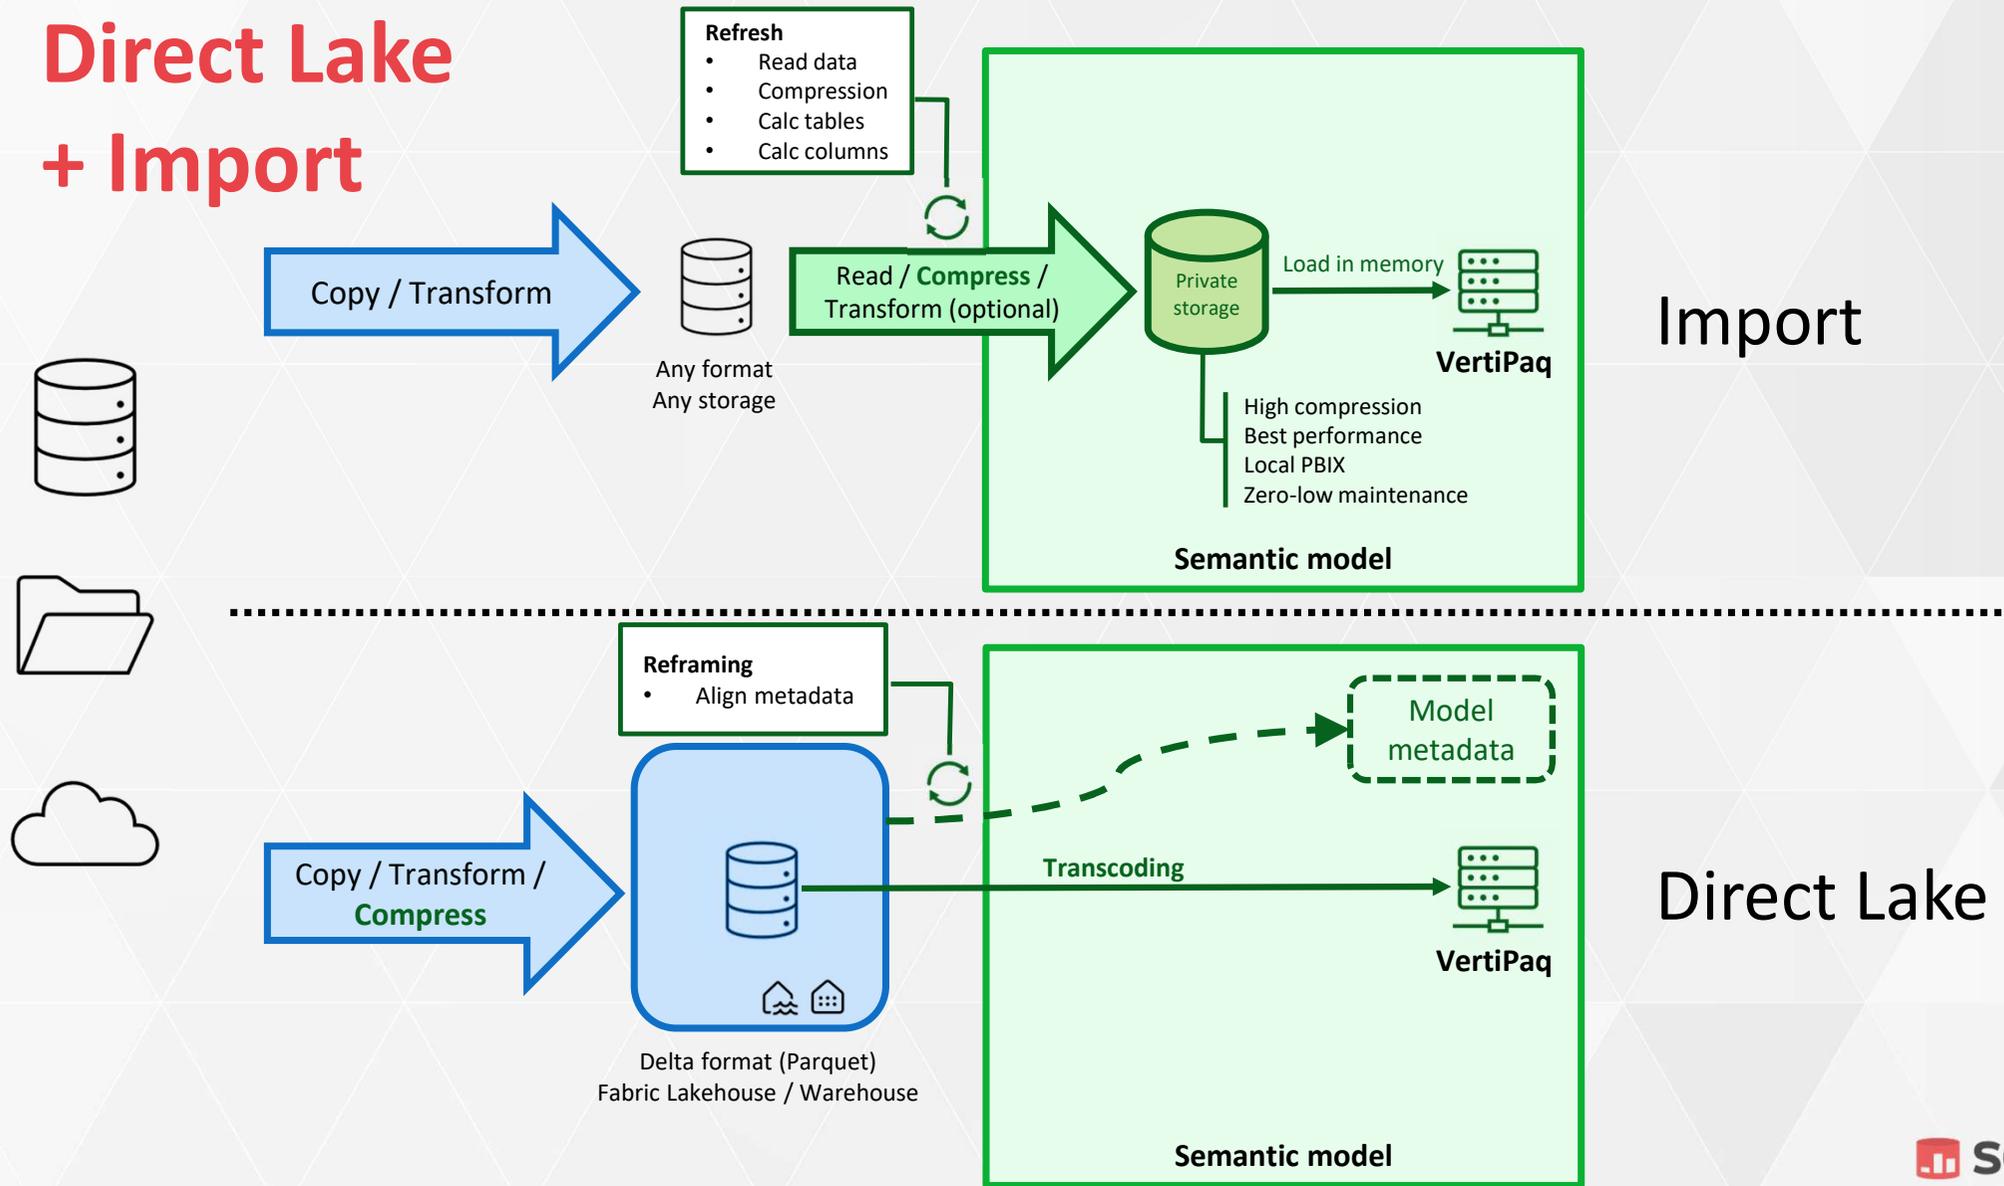o Combine different semantic models

o Relationships between models

- Limited relationships – do not detect referential integrity issues (missing blank row)
- Should be low cardinality
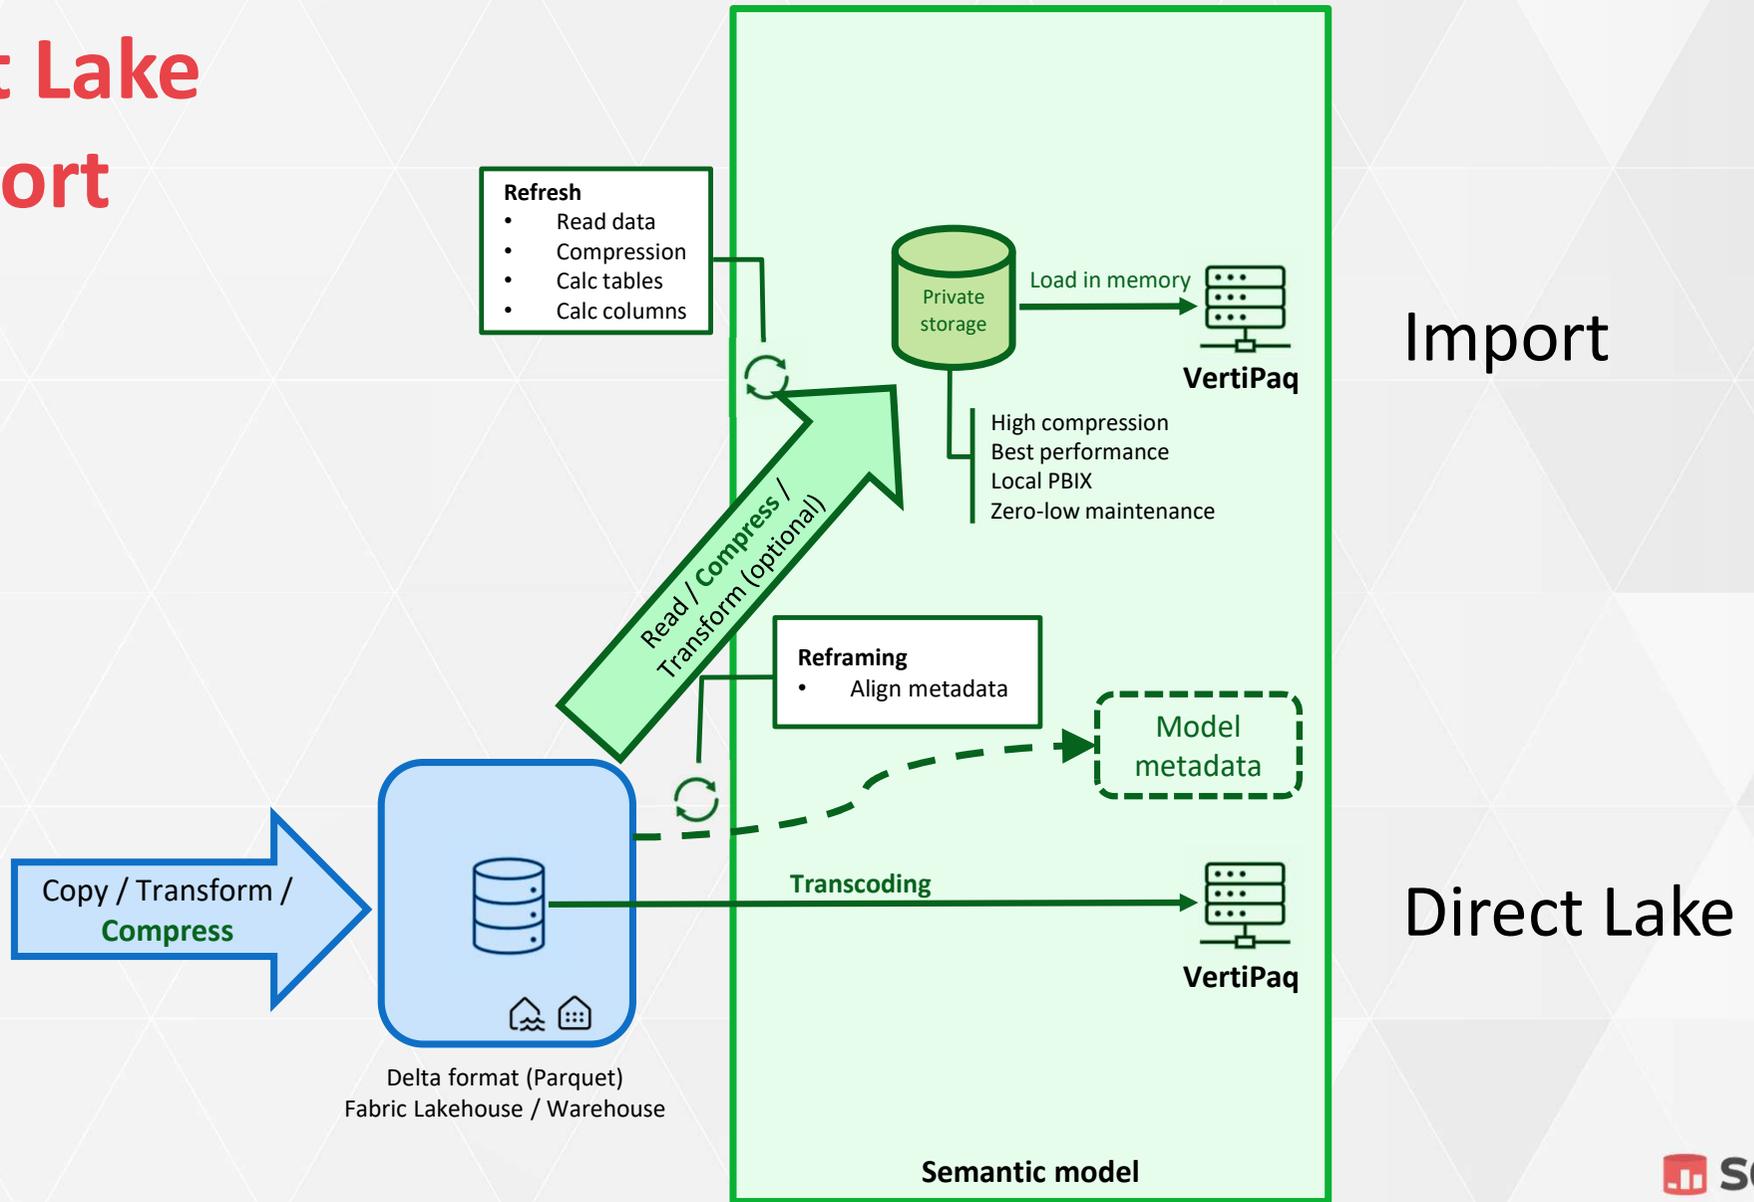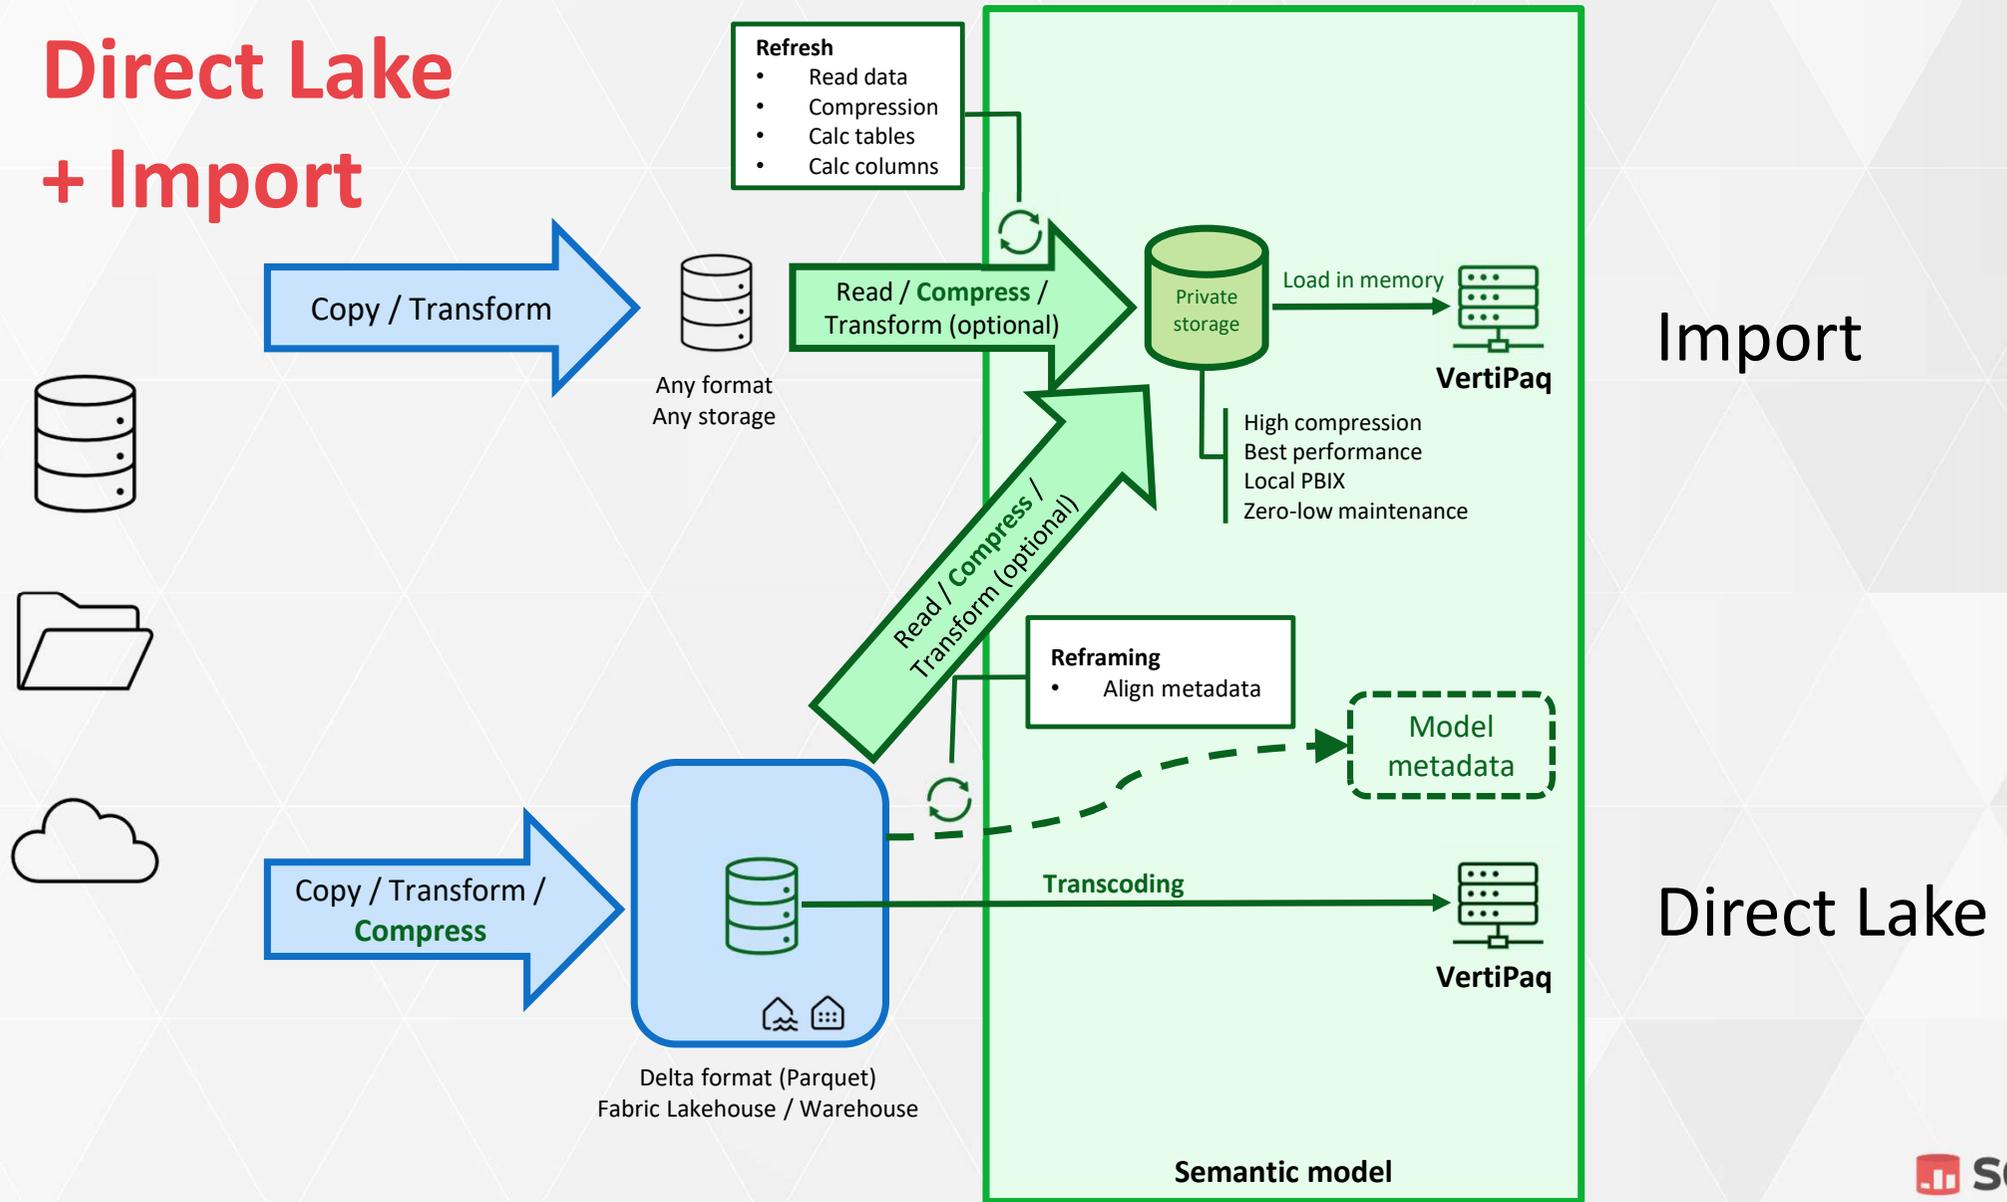- Do not work well to connect dimensions and fact tables

sqlbi.

Direct Lake + Import

Direct Lake + Import

o A new form of composite model – the best of both worlds, finally!

o Direct Lake for fact (large) tables or fast changing tables

- Minimal refresh time
- Usually, no needs to create transformations at the fact table level
- Measures can implement calculations/aggregations

o Import for dimensions / slowly changing tables

- Change tables for reporting needs
- No propagation of aesthetic changes to Direct Lake
- Access to DAX, Power Query, and SQL views to implement transformations
- Access to user-hierarchies also in Analyze in Excel
- Reasonable refresh time (reframing would be needed anyway)

sqlbi.

Conclusion

o Import: best compression, no skills required, any data source, refresh time

o Direct Lake: no data duplication, no refresh time, developers are in charge of compression

o Composite models: combine different semantic models in one

  • High cost for high-cardinality columns in relationships
  • Could produce slow query execution time

o Direct Lake + Import: new flavor of composite model

  • Balance between Direct Lake and Import mode

sqlbi.