Deep dive on ‘Shortcuts’ feature in Microsoft Fabric

Vinodh Kumar
- Cloud data architect @Presidio

19th June 2024

About Me...

Vinodh Kumar

- Cloud Data Architect
- Microsoft Certified Trainer x4
- Microsoft Community Champion
- Global Speaker
- Data Blogger → vinsdata.wordpress.com

Introduction to Microsoft Fabric

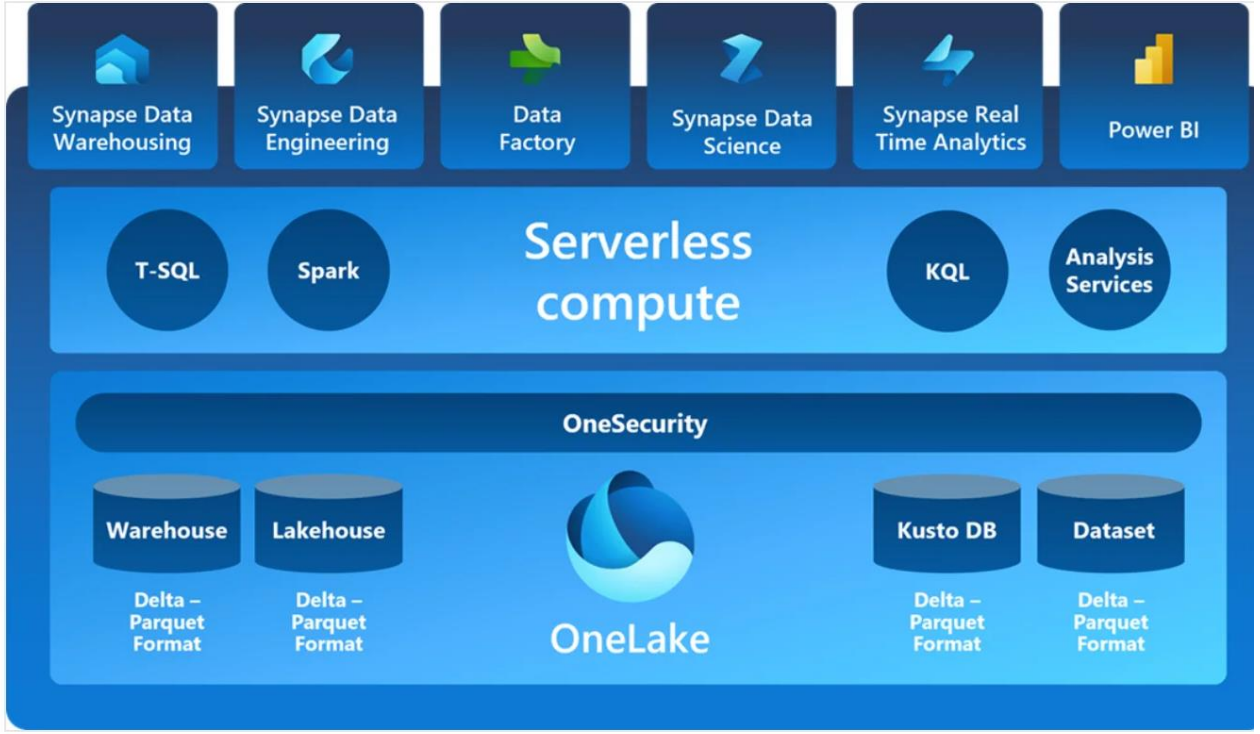

What is Microsoft Fabric?

- A powerful platform for designing, building, and deploying applications.
- Supports various development workflows and integrates with other Microsoft tools.

Key Features:

- Rich UI design capabilities
- Integration with Azure services
- Advanced collaboration tools

Overview of Shortcuts Feature

What are Shortcuts in Microsoft Fabric?

- Shortcuts provide quick access to frequently used actions and components.
- It's designed to be fast, secure, and deliver no-code alternative to building complex pipelines for sourcing data from external storage and to offset development and operational costs of data pipelines.
- Once created, you can quickly access data and share it with other users/workspaces. Since it has properties of **symbolic link**, it will always reflect the current state of data (as it is on external storage).

Key Functions:

- Provides direct access to data sources from the shortcut, bypassing intermediate steps.
- Perform actions without leaving the current context.
- Single source of truth: Ensures that all users are accessing the same version of the data, promoting consistency.

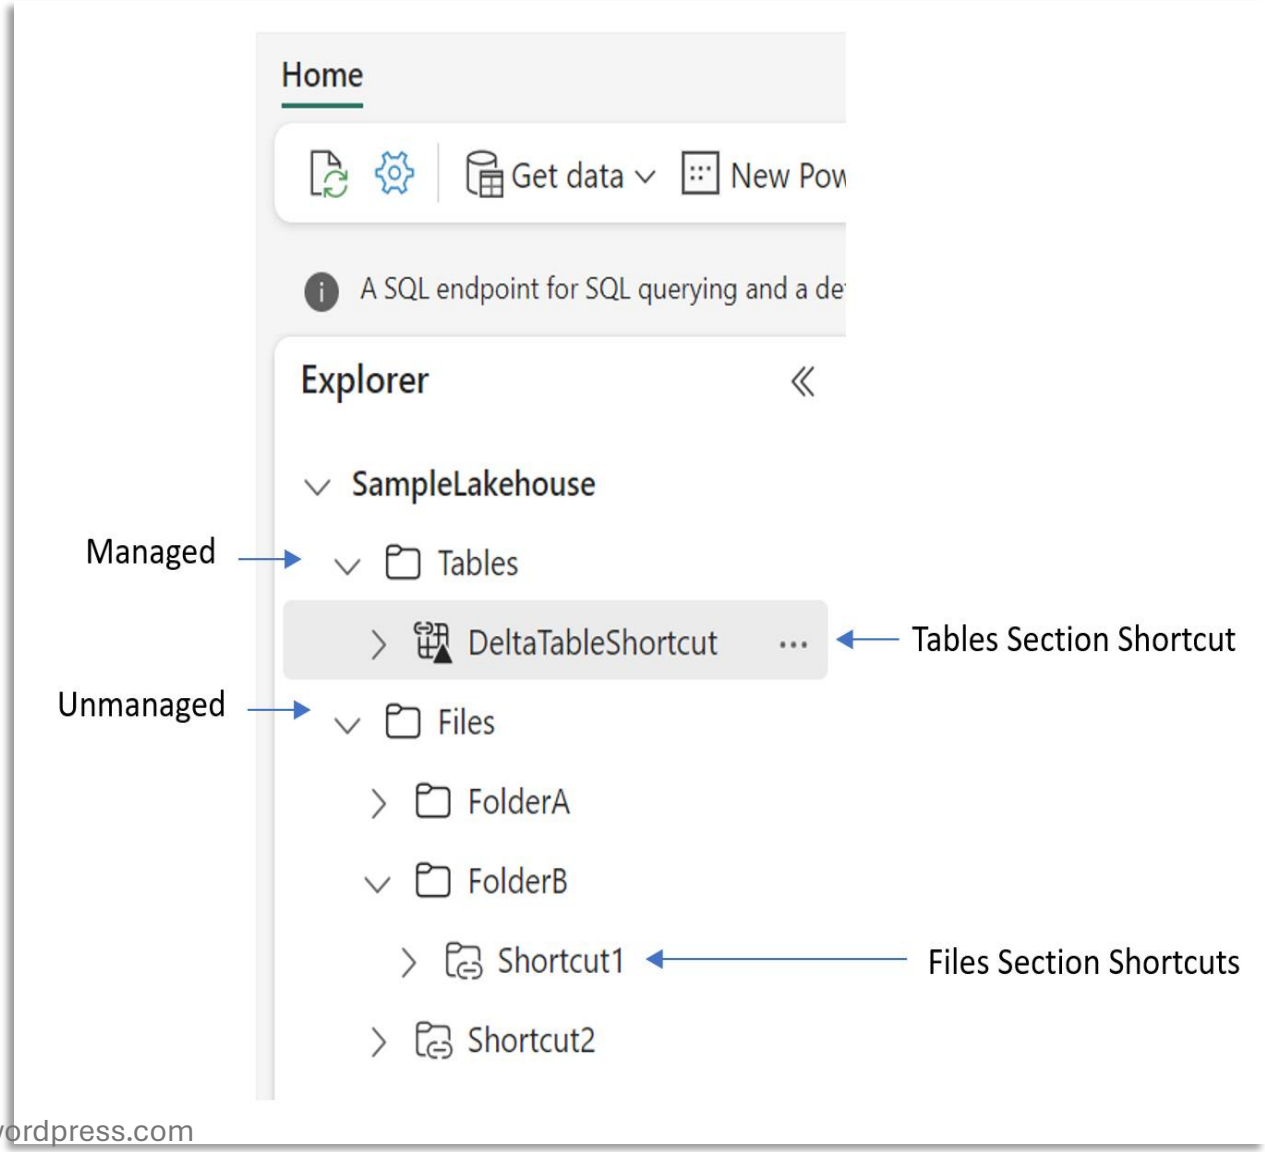

‘Shortcuts’ – Deep dive

- Shortcut does not physically copy or move data to OneLake. A shortcut enables Fabric to perform in-place IO on file(s) wherever they reside.
- Shortcut does not use any compute, as it also does not copy data. It performs in-place reads/writes on remote storage. You can consider them as symbolic links, meaning that they rely upon the IO performance of the remote/external storage where data is being read from.
- If you are reading data from a shortcut, then IO sits with external storage be it ADLS gen2, S3, or GCS. Compute (or CUs) will be dictated by the type of operation you are doing on Fabric. This will sit with the Fabric engine you are using to interact with a shortcut.
- For example, when you read a data file from the ADLS gen2 shortcut using Fabric Spark, the CUs cost is for whatever analysis you are trying to do with data. The IO required will be handled by remote storage.
    - The alternative in this scenario would be to physically store data on OneLake which involves CUs (for compute and also storage transactions) and raw storage cost (to physically store data in OneLake)
- Expect minimal latency because Shortcut is carrying out in-place read/write operations.
- Reading files stored in OneLake using Shortcuts from the same region the Fabric capacity is, can provide better latency than reading files from different regions.
- In some cases, customers are ready to trade slight performance in return for ease of use; avoiding pipeline costs, etc.
- There are various performance-enhancing strategies in place such as Intelligent Cache which would reduce latency. Intelligent cache in Microsoft Fabric - Microsoft Fabric | Microsoft Learn

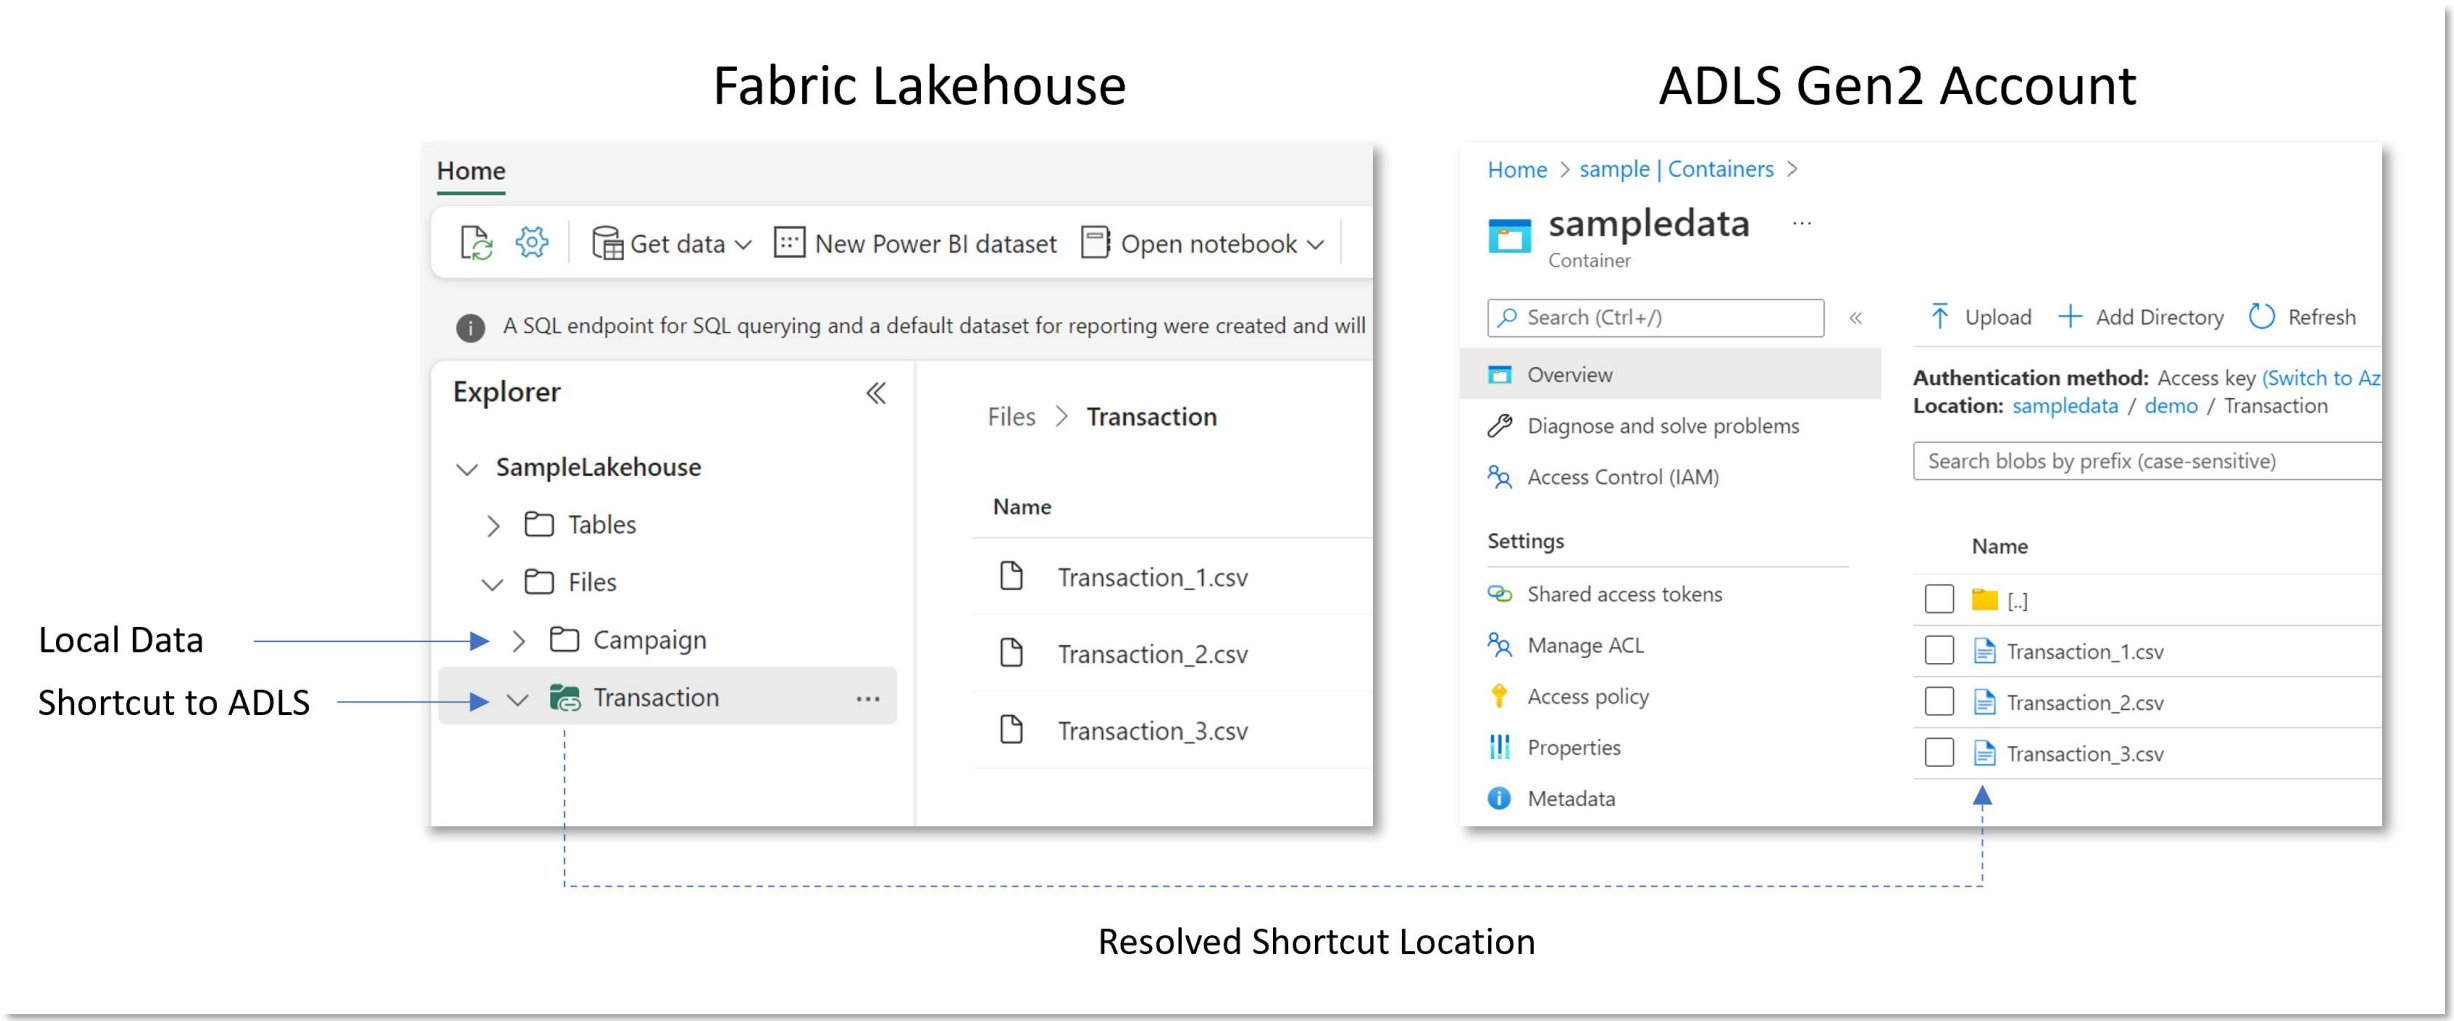

- If the target of the shortcut contains data in the Delta\Parquet format, the lakehouse automatically synchronizes the metadata and recognizes the folder as a table.
- In the Files folder, there are no restrictions on where you can create shortcuts. You can create them at any level of the folder hierarchy.
- Table discovery doesn't happen in the Files folder

Benefits of Using Shortcuts

Improved Data Access:

Quick Access to Frequently Used Data: Fabric Shortcuts allow users to create shortcuts to frequently accessed data sources, reducing the time spent navigating through complex data structures.

Simplified Data Navigation: Users can easily locate and access important data without needing to remember the exact location within the data hierarchy.

Enhanced Collaboration:

Shared Shortcuts: Team members can share shortcuts, ensuring that everyone has easy access to the same important data sets, which enhances collaboration and consistency.

Standardized Data Access: By using shared shortcuts, organizations can standardize the way data is accessed, reducing the risk of errors and inconsistencies.

Efficiency and Productivity:

Time Savings: Reduces the time required to locate and access data, leading to increased productivity.

Streamlined Workflow: By minimizing the steps needed to reach data, workflows become more efficient and streamlined.

Organizational Benefits:

Centralized Management: Allows for centralized management of data access points, which can improve data governance and oversight.

Consistency: Ensures consistent access to data across the organization, which can help in maintaining data integrity and quality.

Types of shortcuts

Onelake	ADLS	S3
Internal OneLake shortcuts allow you to reference data within existing Fabric items. These items include lakehouses, KQL databases and data warehouses.	Must point to the DFS endpoint	Target path must contain bucket name
When you create a shortcut across items, the item types don't need to match. For instance, you can create a shortcut in a lakehouse that points to data in a data warehouse.	Hierarchial namespace should be enabled	You can include prefixes in the shortcut path to further narrow the scope of data accessible through the shortcut. When you access data through an S3 shortcut, prefixes are represented as folders
The calling user must have permissions in the target location to read the data.	Supported deleted types are Org account, SAS, Account Key & service principal	Supported deleted type is a Key and Secret for an IAM user with following permissions
		S3:GetObject, S3:GetBucketLocation, S3:ListBucket
	https://accountname.dfs.core.windows.net/ *https://nyc-tlc.dfs.core.windows.net/trip-data/yellow_tripdata_2022-01.csv*	https://bucketname.s3.region.amazonaws.com/ *https://nyc-tlc.s3.amazonaws.com/trip+data/yellow_tripdata_2022-01.csv*

How to Create and Use Shortcuts

Shortcuts can be created in Lakehouses and KQL databases within fabric.

Creating Shortcuts:

- Navigate to the settings or preferences menu.
- Select the "Shortcuts" tab.
- Choose the action or component you want to create a shortcut for.
- Assign a key combination or a specific button.

Deletion:

- Deleting a shortcut simply removes the reference without affecting the original data.
- Logs and audits can be maintained to track the creation, modification, and deletion of shortcuts for governance purposes.

Demo

Questions and Feedback

Let’s Connect!