



Microsoft Fabric Guidebooks



Real-Time Analytics

“Microsoft Fabric is a new paradigm in how we work with data – it goes beyond BI as we know it.”

“It is probably the biggest innovation in data analytics since Power BI”

Mathias Halkjær
Principal Architect





Microsoft Fabric



Data Factory



Synapse Data
Engineering



Synapse Data
Warehouse



Synapse Data
Science



Synapse Real-Time
Analytics



Power BI



Data Activator
(coming soon)



OneLake

Collaboration

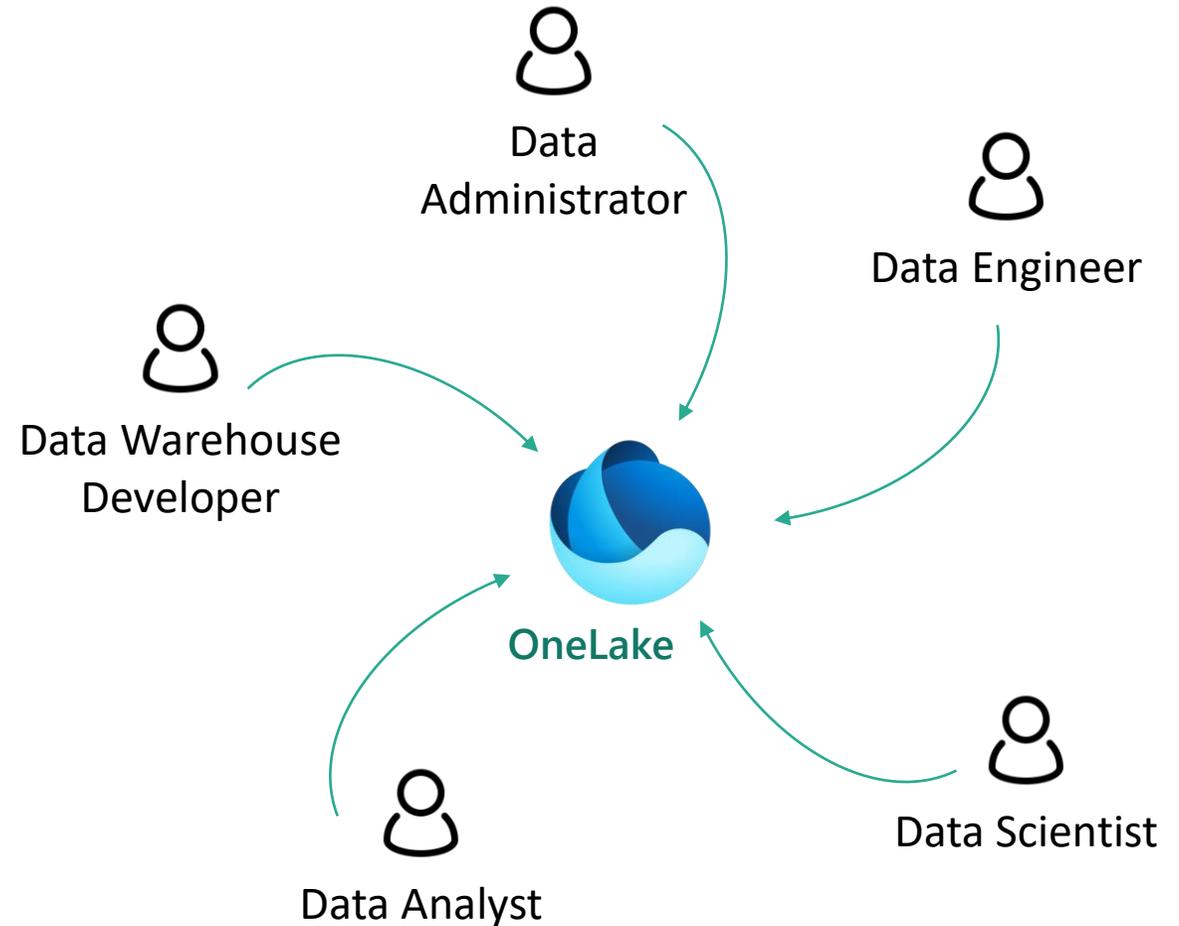


ONELAKE

With OneLake at its core, Microsoft Fabric unifies data disciplines and enhance collaboration across all data professionals.

OneLake both ties together all the tools, experiences and technologies – and by doing so the people working in it.

Never has it been as easy to share ones important and impactful work instantly with the right colleagues.



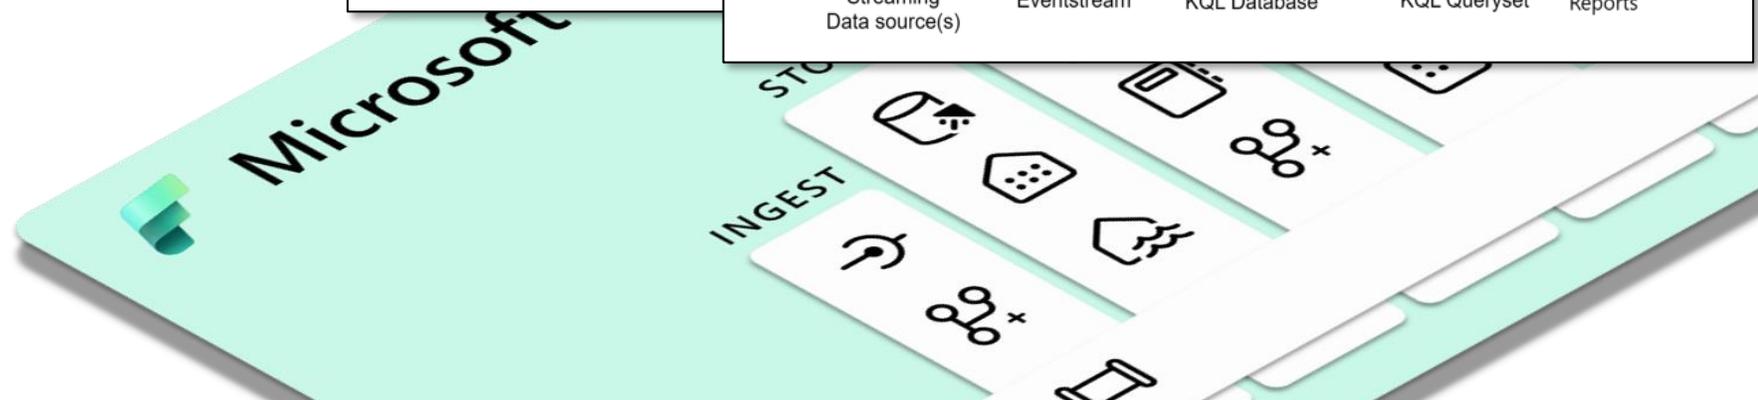
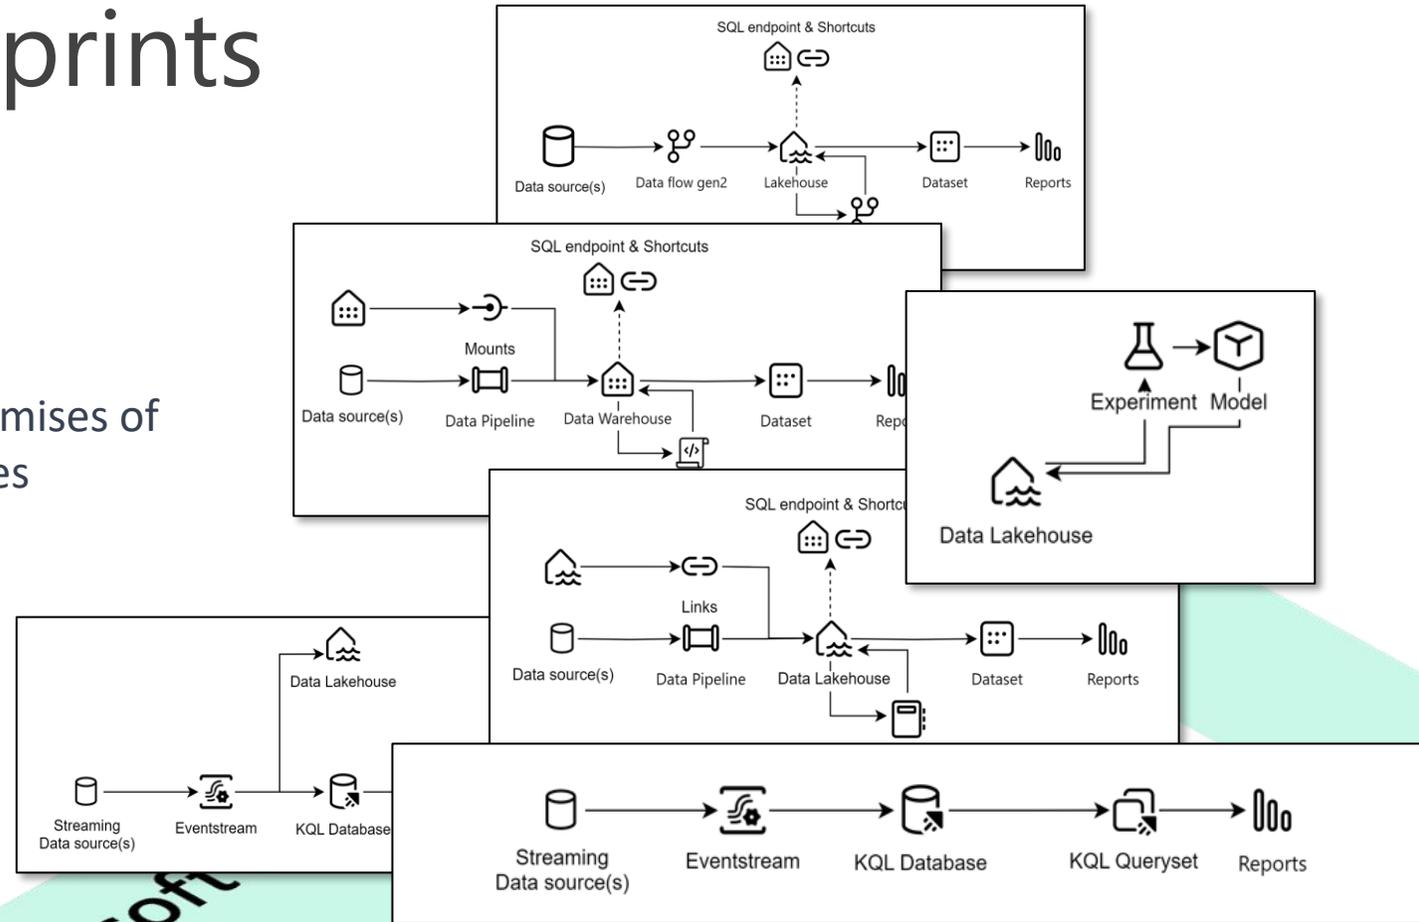


Architecture blueprints

ALL

“One Architecture” is one of the inaugural promises of Microsoft Fabric, that in many ways streamlines architectural complexities.

It does, however, present a variety of options and patterns, enabling users to customize their experience and maximize its potential according to their needs.





Capacity & pricing



CAPACITY



DOMAIN



WORKSPACE

Microsoft Fabric offers a variety of purchasable capabilities divided into SKUs, each providing unique computing power quantified by Capacity Units (CU).

Fabric features two SKU types:

- **Azure** – Billed per second with no commitment.
- **Microsoft 365** – Billed monthly or yearly, with a monthly commitment

SKU*	Capacity Units (CU)	Power BI SKU	Power BI v-cores
F2	2	-	0.25
F4	4	-	0.5
F8	8	EM/A1	1
F16	16	EM2/A2	2
F32	32	EM3/A3	4
F64	64	P1/A4	8
F128	128	P2/A5	16
F256	256	P3/A6	32
F512	512	P4/A7	64
F1024	1024	P5/A8	128
F2048	2048	-	256

Organization



CAPACITY



DOMAIN



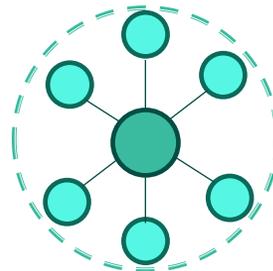
WORKSPACE

Warehouses, Lakehouses, Data Marts, Pipelines and Notebooks. Microsoft Fabric launched with more new gadgets and technologies than we could have ever dreamed of.

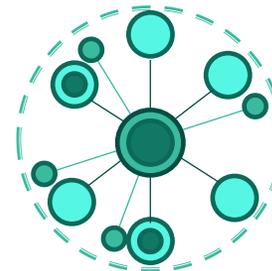
However, it's essential that as organizations, we don't overlook the foundational aspects such as our internal structure, objectives, and strategic planning.

One common organizational decision to consider when deploying a data platform like Fabric, is to choose between a centralized, decentralized, or hybrid implementation approach.

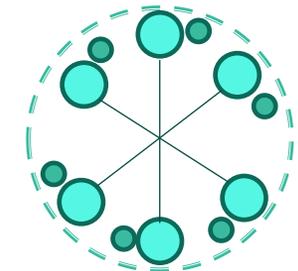
ENTERPRISE



HYBRID



SELF-SERVICE





Microsoft Fabric



Data Factory



Synapse Data Engineering



Synapse Data Warehouse



Synapse Data Science



Synapse Real-Time Analytics



Power BI



Data Activator
(coming soon)



OneLake

Synapse Real-Time Analytics

With the new real-time analytics service in Microsoft Fabric, the ability to read, handle and work with immutable data sets is now available to everyone - without having to pay for a stand-alone service.

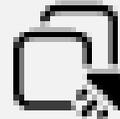
The Kusto engine which resides as the data processing engine under the hood, is second-to-none compared to querying huge amounts of data and return the result set within a matter of milliseconds.

The current implementation of Kusto in Microsoft Fabric can save data in the OneLake directly as parquet files and later the Delta files will be supported.

TOOLS



KQL DATABASE



KQL QUERYSET



EVENTSTREAM

Database solution for real-time data



KQL DATABASE

The KQL Database (Kusto Query Language) can now be enabled and started directly from the portal of Microsoft Fabric.

The ability to read and handle large amounts of data in a second-to-none engine for querying gives the business the ability to gain new insights to their data in a matter never seen before.

Once the KQL Database has been started, it is empty and needs data loaded to it. This data can come from the new Event stream service.

Once data is loaded to the service, it can be used by all the other services using OneLake with a few clicks.

“The Kusto engine is now available to everyone and gives the business near real-time insights to their data entities”

Brian Bøn
Principal Architect



Stored Query Artifacts with KQL Queryset



KQL QUERYSET

The KQL Queryset (Kusto Query Language) is a script which can modify, add, delete etc. the data stored in the KQL Database.

It builds on the KQL language and is easily adoptable from the SQL language.

With additional build in features like outlier detection, vector database functionality, shortest distance between two points etc., the service leverages brand new possibilities to gain new insights on the business data and react to them on near real-time scenarios.

Streaming data with ease



EVENTSTREAM

The brand-new event stream service, leverages the ability to get data from several sources of streaming data and save it to a wide variety of destinations, including OneLake, KQL databases and Azure services.



The service computes the data once and can pipe it out to several destinations at once. All configured and maintained from within the Microsoft Fabric portal and “coded” with your mouse.

Imagine scenarios of IoT devices loading data to both the data warehouse and other 3-rd party destinations – this can now be done using the low-code approach from Event Stream.

New Kusto Magic in Notebooks



NOTEBOOK

The KQL language is now fully supported by the Jupiter Notebooks. Just write your magic with KQL like this:

```
%%KQL
```

The authentication is somewhat new as we are presented with an authentication popup when connecting to the Kusto database.

This enables the data engineer to use Jupiter Notebooks in the entire data pipelines and data transformation directly in a well-known tool.

Get started today

Try Microsoft Fabric

[🔗 Try Fabric \(microsoft.com\)](https://microsoft.com/fabric)

Watch Fellowmind's monthly Power BI Update

[🔗 Power BI Update \(fellowmindcompany.com\)](https://fellowmindcompany.com/power-bi-update)

Connect with our Microsoft Data Platform MVPs

[🔗 https://www.linkedin.com/in/mhalkjaer](https://www.linkedin.com/in/mhalkjaer)

[🔗 https://www.linkedin.com/in/brianbonk](https://www.linkedin.com/in/brianbonk)