Telefónica Tech

Microsoft Fabric *Adoption* Considerations

telefonicatech.uk

Contents

Introduction....3

Microsoft Fabric Outstanding Features....4

Key Considerations For A Successful Microsoft Fabric Adoption....7

Who Are Telefónica Tech?....9

How Can Telefónica Tech Help You?....9

Conclusion....10

Introduction

Back in May 2023, during Microsoft Build, Satya Nadella, CEO of Microsoft, announced the "biggest launch of a data product from Microsoft since the launch of SQL Server, Microsoft Fabric".

Microsoft Fabric is a new cloud-based SaaS solution for data and analytics. Referred to as the OneDrive for data, it brings together, under the same platform, multiple data and analytical experiences/workloads, such as Synapse Data Engineering, Synapse Data Science, Synapse Data Warehouse, Data Factory, Data Activator, Industry Solutions, Real-Time Intelligence and Power BI.

Six months after Satya's announcement at the Ignite conference, Microsoft released Fabric for general availability (GA). This launch generated great excitement among service providers, independent software vendors (ISV) and customers looking for a new data platform. But, beyond the hype surrounding the new data product, many questions were raised, with one standing out: "is Fabric enterprise-ready and should I look to adopt it?"

Telefónica Tech was invited to join the private preview of Fabric, giving us the chance to experience the rapid improvements and new feature releases firsthand. Having worked with several customers and delivered multiple use cases within Fabric, we are in a position where we have a great understanding of the solution, including pros and cons. This enables us to confidently advise our customers on what their best options are when they decide to migrate or build a greenfield data platform.

This article will hightlight the key features of Fabric and what you should consider when looking to adopt Fabric as your next cloud data platform.

Microsoft Fabric Outstanding Features

In no particular order, here are some of the distinctive features of Microsoft Fabric. Please note that this only considers features that are generally available (GA) or in public preview at the article publish date. A link to the official documentation is provided.

1

Serverless

Fabric is built on a foundation of Software as a Service (SaaS), and therefore, there's no need to spend time setting up and managing the infrastructure.

Fabric aims to remove the complexity of deploying and configuring a large and very fragmented collection of products from the customers and provide them with a "5x5" experience, i.e., 5 seconds to sign in, 5 minutes to get data insights.

Read more

2

OneLake

OneLake is a unified data lake solution that serves as a central repository for storing and managing all types of data across an organisation. It is designed to integrate seamlessly with all Fabric workloads, providing a single, logical data lake experience across various cloud platforms and storage services.

OneLake simplifies data management by enabling a consistent, scalable, and secure environment for storing both structured and unstructured data, making it accessible to various analytics and AI tools. This approach aims to reduce data silos, improve collaboration, and enhance data-driven decision-making by providing a single source of truth for all data assets.

Read more

3

Shortcuts

If you work with data and understand what it takes to move it from source to target, then you can't be indifferent to this feature. Shortcuts allow users to create virtual links to data stored in various locations across different cloud platforms and data lakes. By using shortcuts, organisations can seamlessly access, integrate, and analyse data from multiple sources without needing to move or copy the data into a central repository.

This functionality enhances data accessibility and collaboration by providing a unified view of data, simplifying data management, and reducing duplication and data movement costs. It also supports better governance and security by maintaining data in its original location while still allowing centralised access through Fabric's unified interface.

Read more

4

Mirroring

Mirroring is a feature that enables the synchronisation of data between OneLake and other external data storage systems, such as Azure SQLDB, Azure Databricks or Snowflake. It allows users to create an up-to-date copy of data from various sources directly within OneLake, ensuring that the mirrored data reflects any changes made in the original datasets. This feature provides a consistent and reliable data experience by maintaining synchronised datasets across different platforms, improving data accessibility, and supporting real-time analytics.

Fabric Mirroring is particularly useful for scenarios where users need high availability and consistency of data across diverse environments, enabling better data management, analysis, and decision-making without needing to move or duplicate data manually.

Read more

5

Lakehouse and Warehouse

Lakehouse and Warehouse are two enterprise-scale, open standard format workloads for data storage available in Fabric. A Lakehouse combines the scalability and flexibility of data lakes with the performance and structure of data warehouses. It is suitable for data engineers that want to use Spark and notebooks to work with structured and unstructured data. A Warehouse is a fully managed data warehouse service designed for scalable, high-performance data storage, processing, and querying, enabling users to run complex analytics and business intelligence over structured data.

Before Fabric, you would likely select one over the other or use multiple services to achieve something similar, however, full interoperability was not possible. With Fabric, thanks to the standardised data format, you can combine both options under a single solution. Developers can use Spark notebooks to transform semi-structured data such as JSON files, store them in a Lakehouse in the delta format and use T-SQL in the Warehouse to explore the data, directly from the Lakehouse with zero data copies, and create new entities that can be used in the semantic model.

Read more

6

Direct Lake

Direct Lake is a semantic model feature suitable for analysing very large data volumes in Power BI.

By combining the advantages of Import Mode and Direct Query, it can query the data directly from OneLake and access it in real-time while still providing great performance.

Read more

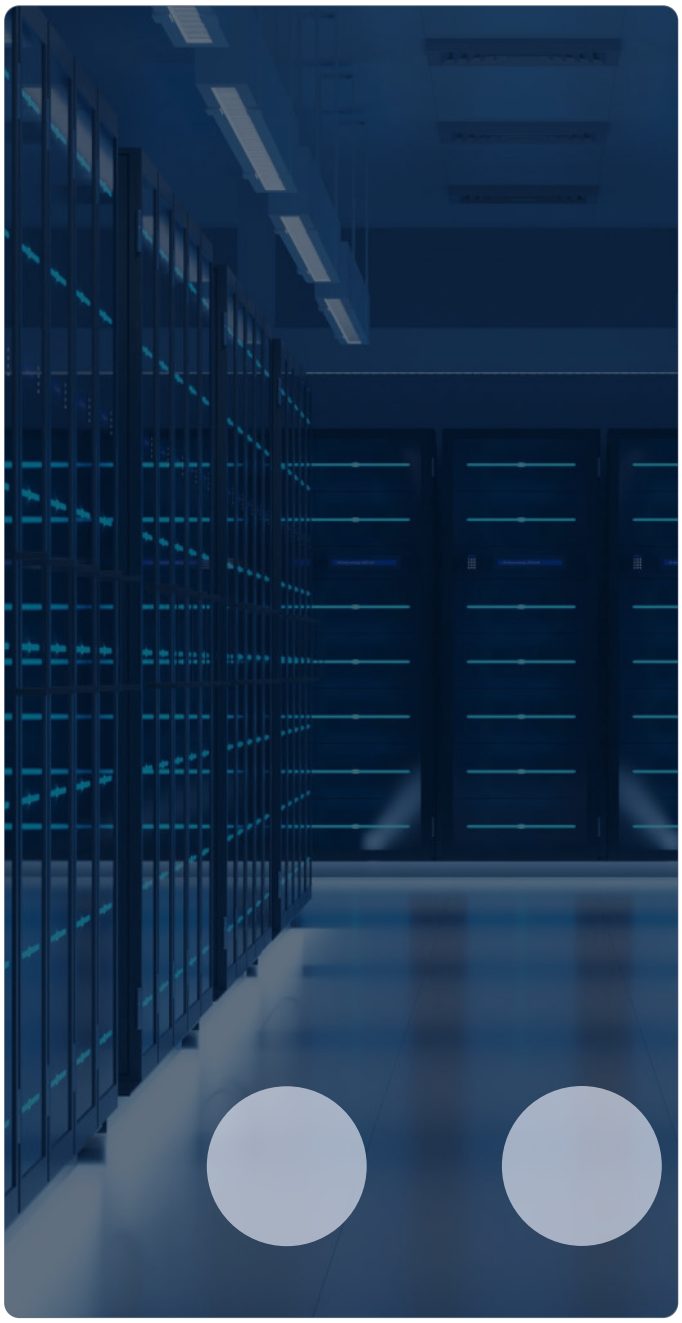

7

Data Activator

Data Activator enables organisations to automate actions based on data insights in real time. It allows users to define triggers and workflows that automatically respond to specific data changes or conditions, helping streamline operations and drive timely decision-making.

By integrating seamlessly with various data sources and analytics tools within the Microsoft ecosystem, Fabric Data Activator helps organisations quickly react to business events, improve operational efficiency, and reduce the need for manual monitoring and intervention. This feature is particularly valuable for use cases that require immediate action based on data, such as alerting, notifications, and automated process management.

Read more

8

Copilot

Copilot is designed to assist users in interacting with their data more efficiently and intuitively. It leverages natural language processing to enable users to query data, generate insights, build visualisations, and automate tasks simply by typing or speaking their requests.

Fabric Copilot helps make data analytics more accessible to a broader range of users, including those without technical expertise, by providing guided, context-aware suggestions and recommendations. Copilot is available across a range of workloads, such as Data Science, Data Engineering, Data Factory, Data Warehouse, Real-Time Intelligence and Power BI. If that is not enough, you can even build your own copilot by using the client advisor AI accelerator tool.

Read more

9

Simplicity

Fabric's goal is to make your life easier. To achieve that, Microsoft provides a unified SaaS experience, suitable for Data Engineers, Data Scientists, Data Analysts and Business Users. It provides a set of no-code, low-code and full code tools with support for various languages, such as T-SQL, PySpark or KQL.

From a billing perspective, you no longer need to worry about forecasting the cost of multiple products, since you have a simplified charging model. Rather than provisioning and managing separate compute for each workload, you simply pay for the amount of storage you use and the amount of compute you provision via the Fabric capacity. Finally, you have the concept of a centralised administration, i.e., you only need to govern one tenant, which contains all the data artefacts. This simplifies security and compliance, since everything is handled in a single place.

Key Considerations for a successful Microsoft Fabric Adoption

It has now been over a year since Fabric was launched and since then we've seen how the product has matured. After the initial hype, we observed how the market reacted, with organisations looking to run proof-of-concepts (PoC) before they commit to a migration or greenfield Fabric data platform.

At this point you may have a good understanding of what Fabric is and its potential. However, you may still be thinking: is it the right solution for your business? What can you do to ensure you are taking the right approach? To get the answer to these questions, you should look to address the following points:

1

Business Value

Having a strategy, a tangible list of objectives and a set of business use cases to determine how a new platform can impact your business is crucial. It's easy to be blindsided by the promise of new and exciting technology, but have you considered how it aligns with your plans? On many occasions, the most difficult step is the first one, so here are a couple of questions to help you understand if you are in the right path towards adopting a new platform.

- What is the driving factor for adopting a new data platform? What will be the Return of Investment (ROI)?
- Do you have a current data platform that does not scale and struggles to cope with the increased data volumes and demand from business users to get quicker access to the data?
- Is your solution on-premises and you need to migrate to the cloud to reduce infrastructure costs, bring additional processing capabilities and explore new areas, such as machine learning and data science?
- Is your solution in the cloud but is convoluted and too expensive to maintain?
- Do you have a good understanding of which insights you want to get from your data?
- What is your current investment in analytics products and how much of it will become redundant?

2

Skillset

Evaluate the technical expertise of your team. Can you upskill your team to learn how to work with new tools and languages or afford hiring new personnel to bridge the skill gap? There is no value in using the latest and greatest technologies and architectures if your business can't support it.

It's important the new platform is designed around your skillset and limitations. For example, you may benefit from having a platform based on Fabric Warehouse rather than using a Lakehouse architecture and the Spark engine.

3

Risk

What is your risk aptitude? Can your organisation afford to adopt a solution that is maturing and may lack on features you consider being key for the success of your business? Can you take a tactical approach where certain limitations may stop the business from adopting a new platform? More importantly, can you continually improve and adapt your solution as new features become available?

It's important to have a list of requirements and understand how many are fully, partially or not met by the targeted solution.

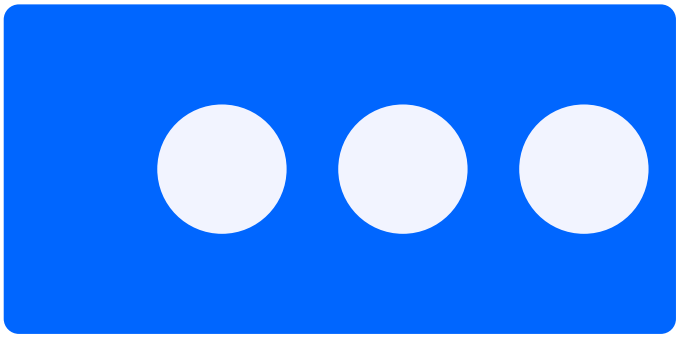

4

Maturity

How mature is your business from a data and analytics perspective? Are you a small business looking for a new solution with low complexity that can provide additional insights beyond what is currently achieved by a small set of disparate products and excel spreadsheets? Or are you a larger enterprise business, with well-established processes and data governance but lacking on a platform that is on par with your growth and business demands?

Understanding how far you are down the analytics journey and how mature your processes are, will help to identify the best solution for your business.

5

Compatibility

If you are not looking to migrate your existing platform, you can skip this point. How compatible is your current data estate with the targeted solution? How many features will need to be re-developed or may not be supported? Will you require connectivity to/from third-party tools remain as is or can you provide, if possible or supported, different alternatives?

The complexity of your integration needs will dictate whether it's safer to adopt a new platform now or wait for further enhancements.

6

Licensing and costs

Cloud is a magnificent and powerful world, where, in many cases, additional power comes with a press of a button, however, it also comes with an associated cost. Luckily, Microsoft simplified the pricing model for Fabric by only using two variables to calculate your bill; provisioned compute and used storage. If you already have a presence in the cloud and use reserved instances to reduce the price, how much of it will be transferrable to the new platform? How much would you lose if you stopped using those services before the end of the reservation? Whether you are new or already have a presence in the cloud, how many capacities will you require to fulfil your requirements? Can you afford to use a lower capacity, or do you need a bigger one (F64+) because you require some of the premium features? How much would you save or pay if you migrated from your existing platform?

Evaluating the operational cost of a new platform and understand the impact it will have in your allocated budget is critical.

Adopting Fabric is not just a technical decision, but one that requires careful consideration. By addressing these key areas, you can set your organisation up for success, and make an informed decision about whether Fabric is the right solution for them and how to approach its adoption.

Who are Telefónica Tech?

Telefónica Tech is a next generation solutions provider that unlocks the power of integrated technology, bringing together a unique combination of the best people, with the best tech and the best platforms, across Data & AI, Cyber Security, Cloud, Business Applications and Modern Workplace, and supported by a dynamic partner ecosystem to make a real difference to every business, every day.

We provide a wide range of services, including data strategy, data platform development, business intelligence, and advanced analytics. We are recognised by Microsoft as an Advanced Solutions Partner for Data & AI and Azure, with Specialisation in Analytics, and AI & Machine Learning.

We'll help your organisation harness its data – whether that means designing a smart, compliant data management strategy, building and evolving your data and AI platform, or managing that platform for you.

A data analytics partner you can rely on, we'll keep you effortlessly at the cutting edge, no matter how fast the tech evolves.

How can Telefónica Tech help you?

With all new technology innovations, understanding how its capabilities can be harnessed effectively for your organisation's requirements is key. The set-up and usability of Microsoft Fabric is designed to be quick and easy. In this 5 Day PoC, our Solution Architects will enable you to quickly understand the features and capabilities of Microsoft Fabric, as well as relevant use-cases that can be delivered by the new solution.

Rapid 5-Day PoC

We can deliver a 5-day PoC for either a data engineering or data science use case, depending on your organisation's requirements.

Day 1
Overview of Microsoft Fabric and organisation stretegy: Exploring the architecture options available, defining the PoC objectives and potential use-cases, the non-functionals and a review of current data strategy and governance.

Day 2
For a data engineering use-case we will review data modelling & mapping, data analysis/profiling, data storytelling and wireframing. For a data science use-case we will assess AI and ML model options.

Day 3
Complete data extraction, ingestion/cleansing patterns and transformation.

Day 4
Deliver semantic layer or ML model.

Day 5
Build reports/dashboards, EDA (Data Insights/Value), or ML model, and review use-case output.

Conclusion

We are excited about the potential of Microsoft Fabric and to see how we can leverage and explore its capabilities and functionalities to address the customers' most complex business challenges.

Having delivered a multitude of data analytics projects across a diverse range of customers and business verticals, we understand that the most successful engagements occur when customers take a bird's eye view of their business and understand how the technology can help them achieve their goals. If you would like to learn more, please speak to a member of our team.

Leading the Way in *Digital Transformation* for our Customers

Telefónica Tech unlocks the power of integrated technology, bringing together a unique combination of the best people, with the best tech and the best platforms, supported by a dynamic partner ecosystem to make a real difference to every business, every day.

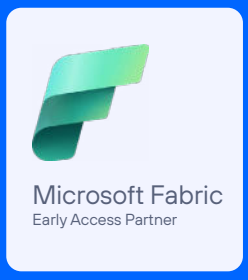

T: 0333 188 7868 | info@telefonicatech.co.uk | telefonicatech.uk