Fabric Accelerator

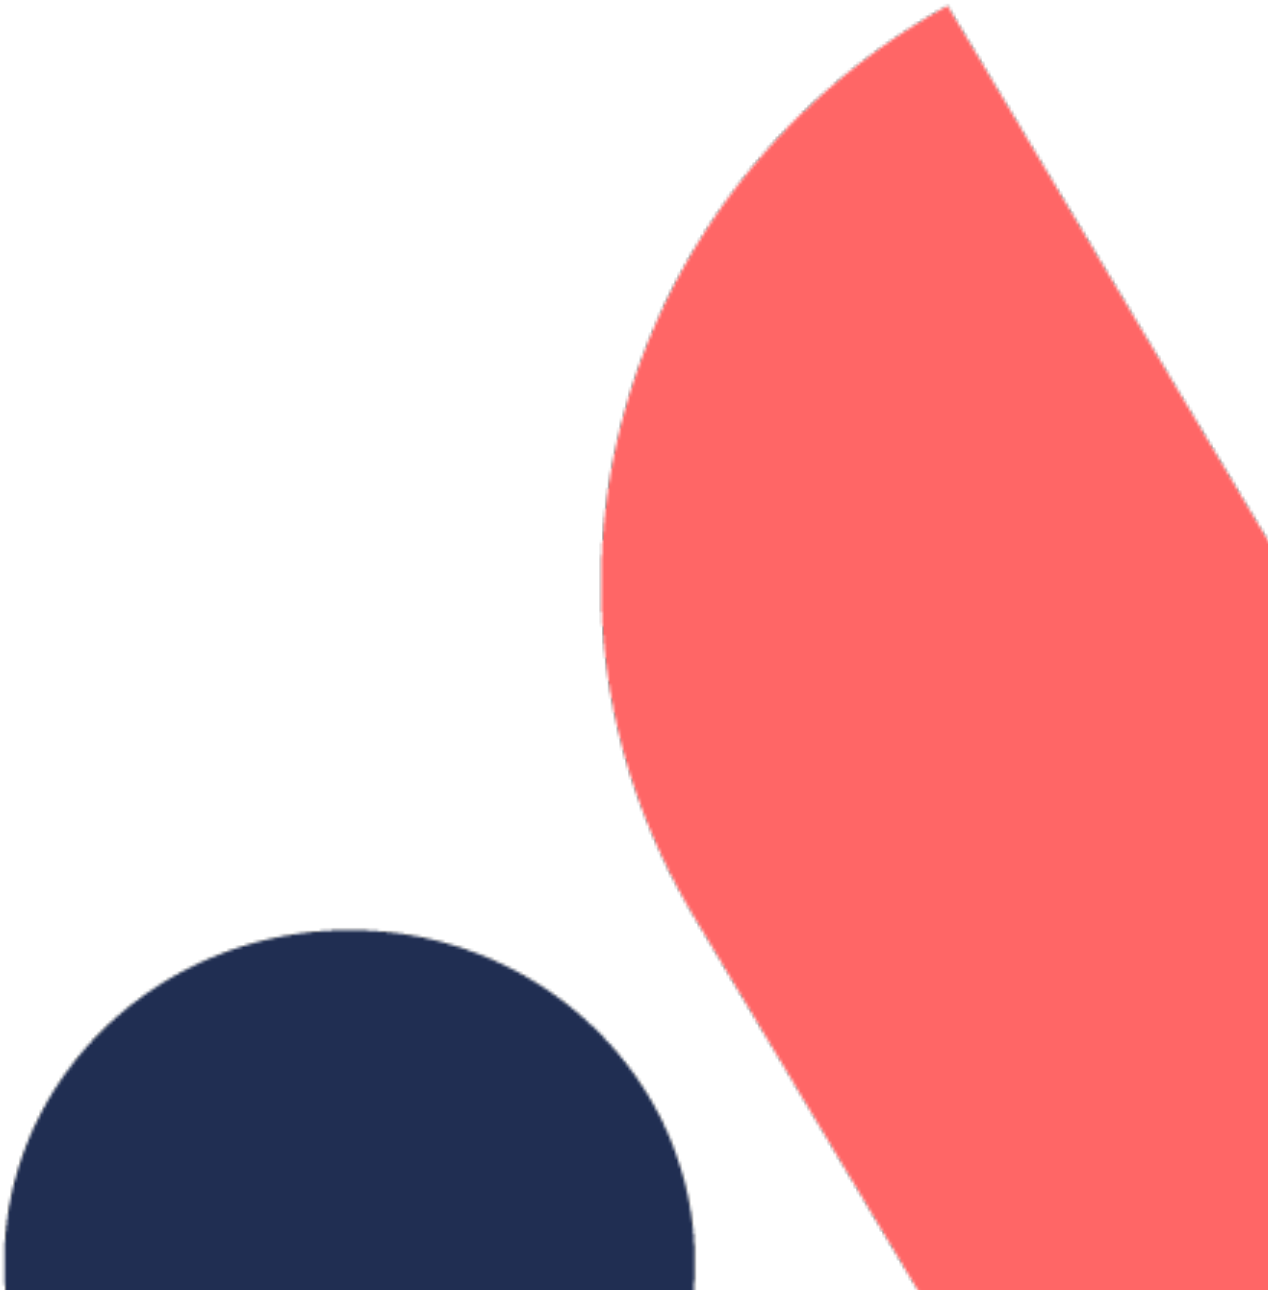

Contents

1. Fabric Accelerator 3

1.1. Scope of Work 3

1.2. Solution Blueprint 4

2. Service Descriptions and Prerequisites 6

Notice

1. Fabric Accelerator

The Fabric Accelerator by ANS is a service designed to deliver a proof-of-concept Fabric capability tailored for organisations, leveraging Microsoft's Azure and Power cloud platforms and ANS' best practice framework. The Accelerator aims to establish a dedicated Fabric environment. Fabric simplifies the consolidation of data at any scale, offering insights through analytical dashboards, operational reports, and advanced analytics. It follows a structured methodology that crafts a customised design while ensuring adherence to best practices and essential requirements. The key outcomes aim to provide you with:

A Accelerator-Ready Data Platform: Designed to handle modern ETL or ELT processing workloads, aligned with the organisation's compliance, security, and architecture standards.

A Defined Operating Model: Ensuring operational readiness by establishing necessary tooling, processes, and skills.

Governance: Offering transparent cost structures and highly governed configurations to meet stringent data security and governance requirements of the organisation.

Stored data within this platform serves Data Scientists and Business Analysts, enabling them to conduct diverse analytics such as data processing, real-time analytics, and machine learning. This facilitates deeper insights and supports improved decision-making processes.

1.1. Scope of Work

The Fabric Accelerator focuses solely on deploying and/or configuring specific logical components for the Fabric Accelerator:

- Data Ingestion & Orchestration
  - Raw data ingestion into OneLake from the data sources detailed below.
  - Configuration of Azure Fabric pipelines
- Data Storage
  - Design and configuration of zones
  - Configuration of access controls
  - Firewall setup for OneLake
  - Disaster Recovery configuration
- Data Transformation
  - Application of cleansing operations to ingested data, such as calculating a current state position based on incremental loads.
- Data Serving
  - Presentation of the cleansed data via a SQL endpoint using the Fabric Lakehouse
- Power BI
  - Creating of a Power BI dashboard to visualise data (limited to a maximum of two days' effort)

1.2. Solution Blueprint

In line with the scope defined above in this document, the following diagram provides a logical representation of the data platform within Fabric. The solution will be deployed in the UK South Region.

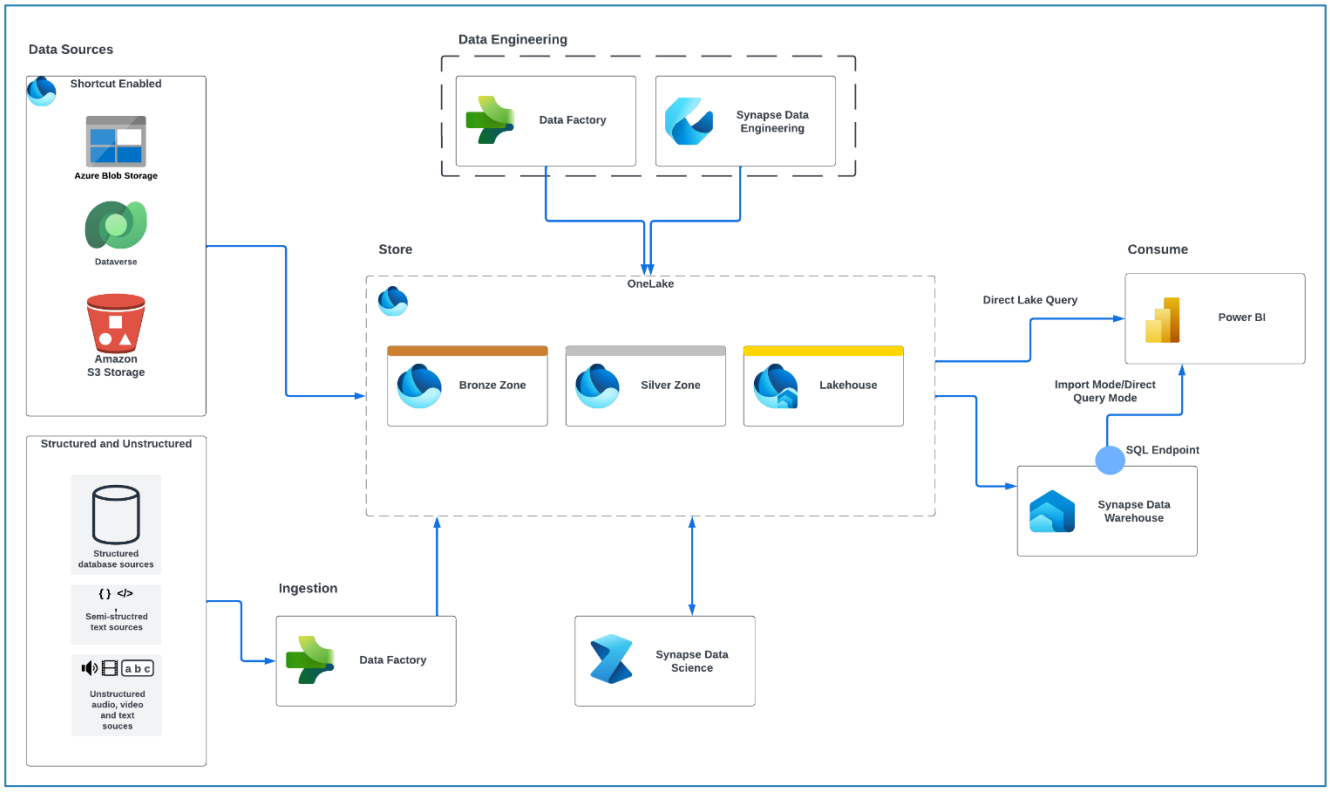

Figure 1: Fabric Logical Architecture

The architecture consists of the following components:

- **Data Sources**
  Any data solution starts with one or more data sources. Examples include:
  - Relational databases, such as application data stores.
  - Static files produced by applications, such as web server log files.
  - APIs which can be called to extract data on demand, such as pricing details.
- **Data Ingestion**
  Acquisition of data, transfer into Fabric and initial storage in the bronze zone of OneLake is performed using Data Factory Pipelines and Dataflows.
- **Data Store**
  Data is stored in efficient file formats in the infinitely-expanding, cost-effective OneLake. OneLake is the single centralised data store underpinning Fabric, used for raw and transformed data, as the storage location for lakehouse and warehouse content, sandboxes, and any other storage purpose.
- **Data Engineering**
  There are multiple options for performing transformations in Fabric. These include the use of Spark-based Data Engineering, low-code Data Factory Dataflows, and SQL in the Data Warehouse.

- **Data Science and Machine Learning**
  Native support for advanced analytics is provided with integrated workspaces for machine learning.
- **Data Serving**
  Fabric provides SQL lakehouse and warehouse choices, either or both of which can be used within the environment. The Data Warehouse provides a familiar full SQL endpoint supporting standard relational database operations for inserting, updating, deleting and reading data, whereas the lakehouse provides SQL querying coupled with Spark-based engineering. Both offer the full SQL enterprise security model and SQL query language support.

2. Service Descriptions and Prerequisites

The following is an outline of what is and is not provided within this engagement:

- ANS Group requires access to the necessary services / applications as required for the duration of the project:
  - Fabric Capacity (Administrator Role)
  - Usernames & passwords required by ANS consultants to allow them access to deploy services into the environment.
  - Connection details to the data sources
- ANS Group will be delivering this project remotely and all meetings will be arranged in advanced and using Microsoft Teams.

What is included within project scope?

- Deployment of the Fabric capacity as detailed above, into your tenant.
- Configuration and development of the data services detailed above, including connecting to the listed data sources, ingestion of data, cleansing of data.
- Configuration of security and access as detailed above.

What is NOT included within project scope?

- Installation or configuration of Data Gateways.
- Configuration of connectivity between your network and Microsoft Fabric.
- Development of any data models (such as star schemas) or data warehouse content
- Anything not specifically listed elsewhere in this document.