

Principal Data Engineer - Fabric

Position Description

Directorate	Digital Innovation and Technology	Department	Digital and Data
Reports To	Manager Digital and Data	Direct Reports	No
Queensland Local Government Industry Award - State 2017 - Stream	Stream A - Division 2, Section 1 - Administrative, clerical, technical, professional, community service, supervisory and managerial services	City of Moreton Bay EBA6 Certified Agreement 2025 Wage Level	Schedule 1, Level 7

Position Purpose

The Principal Data Engineer (Microsoft Fabric) is the senior technical leader and domain expert responsible for the end-to-end architecture, governance, and operational excellence of the Microsoft Fabric platform. This role is pivotal in establishing the Fabric Lakehouse as the unified, trusted data foundation for the strategic City Brain initiative, enabling Copilot, real-time analytics, application integrations and advanced AI/ML capabilities. The Principal Engineer will own the capacity planning, security, data quality and integrity of data assets within OneLake, ensuring data is managed, governed, and seamlessly integrated from diverse sources (e.g., Azure, AWS, and on-prem).

Key Responsibilities and Outcomes

As a Principal Data Engineer - Fabric and member of the <Department> Department you will:

Fabric Platform Architecture & Governance (The Data Authority)

- Serve as the technical authority for all Microsoft Fabric components, including OneLake, Lakehouse, Data Warehouse, and Real-Time Analytics (Eventstreams).
- Own the Fabric platform's operational model, including capacity planning, workspace topology design, environment lifecycle management, and rigorous cost optimization.
- Establish enterprise-grade data governance standards across Fabric by implementing Microsoft Purview for data cataloging, lineage, sensitivity classification (PII), and policy enforcement.
- Define and implement sophisticated security models (RBAC/ABAC) and secure data sharing mechanisms to control access to data assets and ensure compliance with privacy regulations.

Data Product Ownership and Pipeline Engineering

- Lead the design and development of modular, reusable data products and optimal Lakehouse schemas (e.g., Medallion Architecture) using Delta Lake/Parquet within OneLake.
- Design the data ingestion strategy and build scalable ETL/ELT pipelines using Data Factory Pipelines and Spark Notebooks (PySpark/Spark SQL) to handle high-volume, real-time, and batch workloads.
- Drive performance tuning initiatives, including optimizing Spark job definitions, data partitioning, compression, and query efficiency within the Fabric compute layer.

- Implement and own the operational monitoring and alerting (DataOps) to maintain data quality (DQ) and service level agreements (SLAs) for all production data assets.
- Migrate priority reporting loads from legacy platforms to Fabric; retire redundant pipelines.

MLOps, API and Application Integration

- Collaborate with Data Scientists to design and build the feature store and data processing components for robust MLOps workflows within the Fabric Data Science experience.
- Lead the strategy for consuming and publishing data, including defining clear REST/GraphQL API contracts and governance standards for data consumption (e.g., using Azure API Management, Azure Functions, or similar services).
- Own the DataOps/CI-CD strategy for the Fabric platform, leveraging Git integration and Deployment Pipelines for repeatable, automated promotion of Lakehouses, Notebooks, Pipelines, and Semantic Models across Dev, Test, and Production environments.

Technical Leadership & Stakeholder Management

- Set the gold standard for engineering excellence, coding standards, and architectural patterns for the entire Data & Analytics team.
- Actively mentor and guide other Data Engineers and Data Analysts, fostering technical growth, knowledge sharing, and a culture of continuous improvement.
- Champion innovation within the Microsoft Fabric ecosystem, staying ahead of new features, capabilities, and best practices (e.g., Copilot integration) to maximize business value.
- Partner with Microsoft specialists and approved vendors to deliver project objectives while preserving architectural ownership.
- Work with diverse stakeholders across departments to understand requirements and deliver scalable solutions.

Our Values

Our values:
guiding how we
work and shaping
our future.



SEE THE PERSON.
SEE THE CITY.



OWN THE SOLUTION.
FOR TODAY AND
TOMORROW.



FALL, LEARN, RISE.
TOGETHER
WE GROW.



THINK BIGGER.
LET'S GO.

Decision Making

Budget

N/A

Delegations

Delegations under the Local Government Act 2009 and as directed and published in Council's Delegation Register

Knowledge & Experience

- 8+ years of progressive technical data engineering experience, with a minimum of 3 years in a Lead or Principal capacity.
- Deep expertise in designing and deploying scalable, high-volume data platforms in a cloud environment (Azure/Fabric mandatory).

- Expert-level coding proficiency in Python (for Spark/Notebooks) and SQL (T-SQL/Spark SQL).
- Proven experience with DataOps, Infrastructure-as-Code (IaC), and implementing CI/CD for data assets (e.g., using GitHub, Azure DevOps, or Fabric Deployment Pipelines).
- Mandatory deep hands-on expertise across the Microsoft Fabric ecosystem: OneLake, Lakehouse/Delta, Data Engineering (Spark), Data Factory Pipelines, Eventstreams/Real-Time Analytics, and Power BI.
- Expert-level proficiency in Microsoft Purview for implementing comprehensive data governance, lineage, and compliance.
- Strong experience with cross-platform and hybrid data ingestion patterns (e.g., integrating data from AWS S3, Azure Data Lake, and on-premises sources).
- Demonstrated ability to lead with influence—driving architectural consensus and setting technical strategy across engineering, data science, and product teams.
- Strong fundamentals in distributed systems, data modeling (dimensional, data vault, Lakehouse), and optimization techniques.
- Direct experience collaborating on MLOps pipelines and supporting data science feature engineering requirements.
- Ability to identify, analyze, and resolve complex production issues in high-volume, low-latency data pipelines.
- Experience integrating enterprise geospatial (ArcGIS) data with operational analytics is highly desirable.
- Ability to lead engineers, influence stakeholders, and work across multi-disciplinary teams.

Note: This position description reflects a summary of the key accountabilities of the position, it is not intended to be an all-inclusive list of duties, steps and tasks. Leaders may direct team members to perform other duties at their discretion.