


www.michaelpage.co.jp

English-Speaking Data Engineer - Enterprise AI Platform

English-Speaking Data Engineer

Job Information

Recruiter
[Michael Page](#)
Job ID

1600688

Industry

Internet, Web Services

Job Type

Temporary

Location

Tokyo - 23 Wards

Salary

8 million yen ~ 10 million yen

Refreshed

July 3rd, 2026 14:39

General Requirements

Career Level

Mid Career

Minimum English Level

Fluent

Minimum Japanese Level

Daily Conversation

Minimum Education Level

Bachelor's Degree

Visa Status

Permission to work in Japan required

Job Description

This role focuses on building and operating cost and usage data foundations for an enterprise AI platform, enabling financial transparency at scale. You will partner with engineering and business stakeholders to design robust pipelines, models, and APIs that support cost allocation, chargeback, and optimization.

Client Details

Our client is a globally recognized leader in advanced mobility and AI technology, operating at significant scale and investing heavily in next-generation enterprise data platforms.

Description

- Design, build, and maintain scalable cost and usage data pipelines using SQL and Python in cloud data environments (e.g., Snowflake, Databricks)
- Own ingestion and transformation logic for platform-wide usage and billing data
- Develop and enhance cost allocation models (compute, storage, and shared overhead), ensuring alignment with cloud billing
- Deliver cost visibility through well-designed data models, warehouse views, and API endpoints for internal stakeholders

- Establish data quality, observability, and reliability standards (SLIs/SLOs) for pipeline freshness and accuracy
- Collaborate cross-functionally with Cloud Engineering, FinOps, and Business Operations teams to translate requirements into scalable data solutions
- Support broader data engineering initiatives, including additional ingestion and warehousing use cases

Job Offer

- Hybrid work style in Tokyo (3 days/week onsite)
- Opportunity to shape FinOps strategy for a cutting-edge AI platform
- Collaborative, cross-functional environment with strong business impact
- Exposure to large-scale AI and cloud data ecosystems

To apply online please click the 'Apply' button below. For a confidential discussion about this role please contact Serena Wu on +81 3 6627 6137.

Required Skills

- 5+ years of experience in data engineering or backend development using SQL and Python
- Strong experience with cloud billing datasets (e.g., AWS, GCP) and understanding of pricing models
- Proven ability to design and operate data models for financial reporting or cost allocation in modern data warehouses
- Hands-on experience with Snowflake, Databricks, or similar platforms
- Understanding of FinOps principles such as chargeback, showback, tagging strategies, and cost optimization
- Experience working cross-functionally with engineering, product, and business stakeholders
- Business-level English communication skills

Nice to Have:

- Experience with BI tools (e.g., Power BI) and financial dashboards
 - Familiarity with FinOps tooling (e.g., Kubecost)
 - Exposure to AI/ML platforms, MLOps, or training/inference infrastructure
 - Experience in multi-cloud or multi-tenant environments
 - Business-level Japanese
-

Company Description

Our client is a globally recognized leader in advanced mobility and AI technology, operating at significant scale with strong investment in enterprise data platforms.