

Databricks Platform Engineer

Location: UK

Role Type: Permanent - Full Time

Job Location: Hybrid Working - minimum of 2 days a week in our London office

A small, solid pink circle.

ABOUT SAGACITY

Clever with data... combining data, insight, and action to drive value

Sagacity are the data intelligence people. Our proposition is simple: we help consumer facing businesses increase customer engagement and make more money. How do we do it? We dig into the data businesses already have, and combine it with our data, insight, and action. Think of us as data co-pilots for sales, marketing, ops, billing, credit and debt - clients typically see 5x returns within the first few months.

Founded in 2005, we are a customer-centric business with a World Class +95 Net Promoter Score. We believe data should be at the heart of every company and while we are sector agnostic, primarily our clients are in the financial services, retail, energy, telecoms & media, water, and not for profit sectors.

We are a purpose led company and we believe that purpose comes before profit. We work as one team both with clients and internally, and are trusted to deliver quality in our standards and in achieving successful client outcomes. We are open, honest and communicate in a jargon-free, collaborative way, with all our teams being agile and curious, and continuously learning to achieve our purpose. We believe that if we achieve our purpose then everything else will follow.

Databricks Platform Engineer

About the Role

You will design, build, and operate the Databricks Lakehouse platform that powers Sagacity's client analytics, predictive modelling, and AI work. Working closely with our Data Scientists, Data Engineers, and Data & Analytics Manager, you will own platform reliability, governance, performance, and cost, enabling teams to deliver clear, insightful analytics to clients across financial services, retail, energy, telecoms & media, water, and the not for profit sector.

Responsibilities

Principal responsibilities will involve:

- Provision, configure and operate Databricks workspaces, clusters, SQL warehouses and pools across our cloud environment
- Implement and maintain Unity Catalog: catalogs, schemas, external locations, row/column-level security, lineage and audit
- Build reusable Terraform modules for Databricks and automate platform changes through CI/CD pipelines
- Tune Spark, Delta and Photon workloads to improve performance and reduce DBU and cloud spend
- Manage identity, SSO/SCIM, service principals, secret scopes and network isolation to meet security and compliance needs
- Operate Databricks Workflows, Delta Live Tables and orchestration integrations used by Data Scientists and Data Engineers
- Support MLflow, model serving and feature store usage for predictive modelling and GenAI work
- Monitor jobs, clusters and warehouses; lead incident response and post-incident reviews
- QA own work and that of others

What Success Looks Like in the Role

- A secure, reliable, well-governed Databricks platform that our analytics teams can rely on day to day
- Reduced DBU and cloud spend through clear cluster policies, tuning and FinOps reporting, without slowing teams down
- Faster, smoother delivery of client analytics and modelling work, with fewer platform-related blockers
- Continued improvement of Sagacity's Product Suite through a robust, scalable data and ML platform
- Ability to take accountability and ownership for client and internal platform deliverables
- Work as part of the Data & Analytics team, providing knowledge transfer, peer reviews and mentoring to drive continuous learning



Databricks Platform Engineer

Competencies and Behaviours

- 3+ years experience as a Data Platform Engineer, Cloud Engineer or Data Engineer, with hands-on Databricks production experience
- Strong knowledge of the Databricks Lakehouse, Delta Lake, Apache Spark and Unity Catalog
- Production experience administering Databricks on at least one major cloud (AWS, Azure or GCP), including networking, IAM and storage
- Proficiency with Terraform (ideally the Databricks provider) and Git-based CI/CD workflows
- Solid programming skills in Python and SQL; comfortable with PySpark and shell scripting
- Experience tuning Spark jobs and Delta tables (partitioning, Z-ORDER, liquid clustering, file sizing, AQE)
- Working knowledge of SSO/SCIM, identity providers and secrets management
- Understanding of data security, encryption, network isolation and frameworks such as SOC 2, ISO 27001 or GDPR
- Highly numerate with strong attention to detail
- Solutions orientated with natural problem-solving skills
- Self-motivated and pro-active, with an ability to multi-task, prioritise, be organised and plan your work
- Excellent communication skills, both written and verbal, with a willingness to engage and influence others
- A team player with a positive attitude; personable, adaptable and willing to learn
- Are flexible and self-motivated
- Are good under pressure
- Able to travel throughout the UK
- Have the right to work in the UK
- Have a commitment to personal development

● ABOUT SAGACITY

People at Sagacity

Success in any business is ultimately about its people; their skills, personality, attitude, qualities, dedication and enthusiasm.

We recognise the value of our people and their commitment to working together to achieve successful outcomes.

If you would like to join a unique working environment, with a sociable culture, where work is done a little bit differently – and we believe 'better' - then we look forward to hearing from you!

At Sagacity, we:

- believe working with our clients collaboratively delivers better results
- coach and mentor our clients' teams so our data and solutions live on after our assignment ends
- believe in delivering benefits as we go along

We achieve this through our core values:

- ✓ One Team
- ✓ Quality Delivery
- ✓ Trusted
- ✓ Agile & Curious
- ✓ Open, Honest, Simple Communication