



databricks

G-Cloud 15

Service Definition

Databricks Professional Services & Training

Lot 3 Cloud Support Services

1. Introduction	3
Company Overview	3
Value Proposition	4
Overview of the G-Cloud Service	7
Databricks Support Services	16
2. Data Protection	29
Information Assurance	29
Privacy by Design	36
3. Using the service	37
Ordering and Invoicing	37
Pricing Overview	37
Implementation Plan	37
Project Management and Delivery	38
Service Management	41
Service Constraints	41
Service Levels	43
Outage and Maintenance Management	45
Financial Recompense Model for not Meeting Service Levels	45
4. Provision of the service	46
Customer Responsibilities	46
Technical Requirements and Client-Side Requirements	46
Outcomes/Deliverables	46
Development life cycle of the solution	46
After-sales Account Management	46
Termination Process	47
5. Our experience	48
Clients	48
Contact Details	49

1. Introduction

Company Overview

With origins in academia and the open source community, Databricks was founded in 2013 by the original creators of Apache Spark™, Delta Lake and MLflow. All seven co-founders, at the time of conceiving Databricks, were UC Berkeley researchers, who were able to successfully bank on the proposition: when combined with A.I., data holds the promise of curing diseases, saving lives, reversing climate change, and even changing the way we live.

As the world's first and only lakehouse platform in the cloud, Databricks combines the best of data warehouses and data lakes to offer an open and unified platform for data and AI.

Databricks is the founder of three of the most popular open source projects Apache Spark, Delta Lake and MLFlow and has pioneered the Lakehouse paradigm which provides a simple, open platform capable of handling all data workloads, from ETL through to advanced Data Science and Machine Learning. Available on all 3 main public cloud platforms (AWS, Microsoft Azure and Google Cloud Platform) and based on open source technologies, processes and data within Databricks are inherently portable.

More than 9,000 organisations worldwide — including Ordnance Survey, Department for Education, DEFRA, US Air Force, European Maritime Safety Authority and NASA — rely on Databricks to enable massive-scale data engineering, collaborative data science, full-lifecycle machine learning and business analytics.

Headquartered in San Francisco, with offices around the world and hundreds of global partners, including Microsoft, Amazon, Tableau, Informatica, Capgemini, Accenture, Deloitte and EY, Databricks is on a mission to simplify and democratise data and AI, helping data teams solve the world's toughest problems.

Value Proposition

Data and data-related capabilities play a crucial role in informing policy decisions within the public sector, offering insights that can enhance governance, improve service delivery, and address societal challenges more effectively.

Evidence-based Decision Making: Data provides empirical evidence that policymakers can use to understand societal trends, assess the impact of existing policies, and identify areas requiring intervention. By analysing data on demographics, economics, health, education, crime, and more, policymakers can make informed decisions grounded in evidence rather than intuition. Databricks platform enables users to process country level data volumes that are traditionally too big to handle with ease. The platform provides scalable compute that does not require code changes to scale up with increased data volumes.

Predictive Analytics: Advanced data analytics techniques, such as predictive modelling and machine learning, enable policymakers to forecast future trends and anticipate potential outcomes of policy interventions. By analysing historical data and identifying patterns, governments can predict the potential effects of various policy options before implementation, thereby minimising risks and maximising positive impacts. Databricks provides end to end capabilities for managing ML/AI/GenAI workloads leveraging open source technologies such as MLFlow for easy exporting/importing of models in/out of the platform fully embracing FAIR principles.

Performance Monitoring and Evaluation: Data allows policymakers to monitor the performance of government programs and initiatives in real-time. By tracking key performance indicators (KPIs) and conducting comprehensive evaluations, policymakers can assess the effectiveness and efficiency of policies, identify areas for improvement, and make necessary adjustments to achieve desired outcomes. Databricks platform provides seamless integration of streaming data into data pipelines, the code used to process batch and streaming data has minimal differences and only where explicitly required. This in fact allows for build once pipelines and use in both batch and streaming use cases. Data coming from streaming sources can easily be integrated with historical data to build both retrospective, predictive and real-time reporting.

Targeted Interventions: Data-driven insights enable policymakers to target interventions more precisely, ensuring that resources are allocated efficiently and effectively. By identifying vulnerable populations, geographic areas with the highest need, or specific risk factors, policymakers can tailor policies and programs to address underlying issues and achieve greater impact. Databricks platform provides state of the art ETL tooling and allows for easy slice and dice over various data assets that can come in any number of open formats. The data can be processed and served using python, R, SQL or scala as programming languages.

Transparency and Accountability: Making data publicly available promotes transparency and accountability in government decision-making processes. By publishing data on government spending, performance metrics, and policy outcomes, policymakers can engage citizens, civil society organisations, and other stakeholders in the policymaking



process, fostering greater trust and accountability. Databricks provides out of the box sharing service – delta sharing, through delta sharing organisations can easily provide any data product both internally and externally without the need to manually build Rest API services. This in effect is increasing transparency and shortening the sharing cycle.

Cross-Sector Collaboration: Data can facilitate collaboration across different sectors and levels of government. By sharing data and insights across agencies, departments, and jurisdictions, policymakers can develop holistic solutions to complex challenges that require coordinated efforts across multiple domains. Databricks platform both through Delta Sharing and through Databricks Marketplace can expose data assets and data derivatives such as notebooks, code, dashboards, models, etc. to organisations both in Public and Private sectors. These assets can be shared free of charge or could be monetized if there is a need for monetization.

Risk Management and Disaster Response: Data enables governments to assess and mitigate risks, particularly in areas such as public health, natural disasters, and national security. By analysing data on disease outbreaks, weather patterns, or social vulnerabilities, policymakers can develop proactive strategies to mitigate risks, enhance resilience, and improve emergency response.

On the infrastructure side, a common challenge across public sector organisations is the time and effort required to produce the data and insights required for decision and policy makers. This is driven in part by a complicated IT and data infrastructure that has evolved over many years across many different generations of technologies, and a wider market challenge of attracting the required skills. The Databricks Lakehouse Platform simplifies your data processing architecture, reduces the data infrastructure management overhead and provides a centralised security and permissions and model for controlling and monitoring access to data assets, such as PII or machine learning models. This creates a significant administrative and maintenance burden, running technologies that are outdated, hard to find talent for, and no longer represent the best solution to the problem they address. Furthermore, these disparate systems create distinct silos of data within the organisation that make it difficult and time consuming to consolidate information for the analysis and insights necessary to inform decision and policy makers. In these highly federated environments, Databricks can act as a single point of entry for data users by peering into different cloud tenancies or on-prem servers. This can reduce the administrative burden on granting data access requests, by consolidating all requests into a single place. Combined with security and governance policies and monitoring this can democratise data across the organisation whilst also ensuring the controls who can access what data are maintained and monitored. This allows an organisation to reap the benefits of increased data access, reduced time to insight and a reduced administrative burden without requiring an initial long term migration.

Another common challenge faced across the public sector is the attraction and retention of engineering and analytical talent with modern skillsets such as python, combined with an existing talent pool of SQL and Excel users. For the former, product

features such as Databricks SQL provide a SQL editor native within Databricks which abstracts away the underlying technologies such as Spark and data lakes

To simplify the learning curve and ease the journey to the cloud, product features such as Databricks SQL offers a route to cloud for users with traditional Business Intelligence or Data Warehousing skills such as SQL. Because the platform is accessible to users of all skill levels, this leads to reduced friction when migrating users off their traditional tools and patterns, such as working in Excel on their local laptop in an ungoverned and unaudited fashion to the cloud where data access and use is easily controlled and monitored.

Overview of the G-Cloud Service

Databricks Professional Services (refer to our rate card/pricing for detailed pricing)

Delivering rapid success on projects with world-class data engineering, data science and project management expertise. Databricks Professional Services can help you at any point in your data and AI journey:

Key benefits of Databricks Professional Services:

- Accelerate your data and AI journey: Our offerings and expert services help you in your data and AI journey to provide optimal and tailored acceleration from the initial workspace onboarding to building out DataOps, AI capabilities and Center of Excellence practices at enterprise scale.
- De-risk your project: Whether you are migrating from legacy workloads to Databricks, building out new data products or data and AI pipeline and machine learning projects, let us be your partner and trusted advisor to minimise risk and maximise value at each step of your journey.
- Operationalise at scale: Building a data pipeline POC or single node model development is relatively manageable. The challenge lies in how you successfully adopt and scale data and AI practices across the company. Let us help you meet this challenging goal with our prescriptive offerings and expertise.

Databricks will generally provide the Databricks Professional Services remotely, unless we agree to provide Databricks Services at a mutually agreed location at your request. If our personnel travel at your request, you agree to reimburse our reasonable, actually-incurred travel and lodging expenses. Pre-purchased Professional Services will expire after the period indicated on the Order. Professional Services are quoted and provided on the basis of 8-hour service days.

Databricks Professional Services Packaged Offerings:

Delivering rapid success on projects with world-class data engineering, data science and project management expertise in pre-packaged offerings. These includes:

Architecture and Best Practices Assessment

Well-Architected Lakehouse with focused areas:

- Security and Data Governance reviews, unify data management, security and best practices
- Operational Excellence: Optimise build and release processes, including CI/CD, pipeline design, testing and orchestration, etc.
- Delta Lake design and performance efficiency, including data layout, partitioning strategy, table design and Z-ordering
- Dev/Test/Prod workspace design, including cluster policies, tagging, monitoring and logging

Architectural Design Review

Well-Architected Lakehouse Design:



- Security and Data Governance reviews, unify data management, security and best practices
- Operational Excellence: Optimise build and release processes, including CI/CD, pipeline design, testing and orchestration, etc.
- Delta Lake design and performance efficiency, including data layout, partitioning strategy, table design and Z-ordering
- Reliability: high level design for auto scaling and DR best practices, monitoring and logging, etc.
- Usability: standards, open interfaces, data consumption patterns and best practices

Enterprise-Readiness Review

Validating Databricks environment is ready for running mission-critical data and analytics workloads at enterprise scale. Milestones composed of Building, Maintaining, Extending the platform and Implementation Office Hours.

- Ensure enterprise-ready platform that scales
 - Provide advice, guidance and enablement of customers to implement secure Databricks workspaces.
 - Make recommendations on cluster types, cluster management and cost control.
- Findings and best practices recommendation
 - Secure workspace provisioning and administration incl. monitoring and logging
 - Cluster configurations and sizing
 - Cost management recommendations and examples of cluster policies

DevOps best practices

Validating Databricks environment is automated to rapidly develop, test and productionalize new features.

Key milestones:

- Current State Review & Enablement
- Target State Design
- CI/CD Pipeline Implementation Office Hours

Ensure operations best practices are followed

- Provide advice, guidance and enablement of customers to implement automated CI/CD pipelines.
- Make recommendations on productionizing and jobs orchestration.
- Example CI/CD pipeline that tests and deploys code automatically.
- Recommendations on code organisation and testing, jobs orchestration, libraries management

MLOps best practices

Integrate your existing ML workflows with MLflow along with MLOps best practices.
Advising on various ML modelling and deployment approaches

The package includes:

- Advisory sessions focused on ML topics such as architecture, implementation, and operationalization
- Review sessions to discuss existing models and deployment process
- Databricks feature integration working sessions for model development and deployment
- CI/CD architecture design
- Scaling ML pipeline review
- Advisory sessions consolidated into a best practices document

Data Engineering Support

As an assurance offering, this package provides guided expertise, helping your delivery team implement data engineering solutions and derive accelerated value on Databricks. This involves understanding your business requirements, assessing your existing architecture, providing recommendations around best practices for design and implementing your desired use case through your journey to production.

- Support project team delivery of the use case into production
- Define the project scope, success criteria and production timelines
- Provide expert support in performing key data engineering activities per the use case implementation plan – best-practice recommendations, regular office hours, plan out the production implementation, code review and develop exemplar code.

Delta Lake – Use Case Acceleration

The main objective of the Delta Lake use case acceleration is to understand the business context, the details of the business requirements and the existing data architecture. The second objective includes providing recommendations and best practices to ensure the architecture and design meet the requirements, then we support and enable you during the implementation journey and make sure the final solution is production-ready. Note that this type of engagement focuses on one specific Delta Lake-enabled end-to-end use case.

- Requirements gathering
- High-level solutions architecture recommendations — Documenting the high-level architecture and decisions made
- Delta Lake workshop
- Office hours
- Production readiness — Defining the path to production and scale after the MVP

Streaming Use Case Acceleration

The main objective of the streaming use-case acceleration is to understand the business context, the details of the business requirements and the assessment of any streaming architecture that exists. The second objective includes providing recommendations and best practices to ensure the architecture design meets the defined requirements, then we support and enable you during the implementation journey and make sure the final solution is production-ready. Note that this type of engagement focuses on one specific streaming use case.

- Requirements gathering
- High-level streaming solutions architecture recommendations — Documenting the high-level architecture and decisions made
- Streaming workshop
- Office hours
- Production readiness — Defining the path to production and scale after the MVP

Data Estate Optimisation Assessment

Assess, analyse and provide recommendations to optimise the data estate for scale and TCO reduction

- Review deployment & operational architecture
- Review data, pipelines, data products & applications
- Review Databricks workspaces
- Review cloud provider / services configurations
- Explore migration & modernization opportunity for legacy / cloud DW and big data workloads
- Deploy data capture via system table and Lakeview
- Compile observed inefficiencies and root causes

Recommendation and Action Plan

- Categorise & prioritise inefficiencies by impact and remediation difficulty / LoE
- Synthesise optimization recommendations for short, medium and long term

Migration Assurance

The Databricks Migration Assurance package gives you a concrete migration outcome through a prescriptive approach that maximises the value of your existing data and pipeline investments — ensuring seamless migration to the Databricks Lakehouse Platform and best practices conformance while reducing risk and cost. Databricks subject matter experts work as part of your team providing architecture, planning, design and ongoing expert support and technical guidance in the form of an advisory capacity throughout the migration journey to ensure success. Specific activities performed will vary, depending on Customer-specific objectives.

There are three levels of Migration Assurance packages

- Migration Assurance Silver
- Migration Assurance Gold
- Migration Assurance Platinum

The core activities performed are the same across different tiers; however, the duration of engagement is longer for higher tiers, as they tackle increasingly larger scope and higher complexity.

Unity Catalog Reference Implementation

The Databricks Unity Catalog package configures Unity Catalog as the unified governance solution for your data and AI assets. By applying structured and best practices-led approaches to help ensure the right architecture and delivery, the package includes design and delivery of one Unity Catalog reference implementation on a single workspace.

- Unity Catalog readiness
- UC implementation on a single workspace
- UC with Delta Sharing
- Documentation and closeout

SQL MVP

The Databricks SQL MVP package is designed to help you realise lakehouse architecture by implementing SQL warehouses (compute/endpoint) so you can directly query data from your Databricks Lakehouse. It is a sister offering to the Lakehouse Buildout and Unity Catalog packages, in which you'll leverage your existing Silver or Gold tables as the query source and Unity Catalog as a unified governance layer for all your data.

- SQL and BI Requirements and Planning
- Access Control for SQL Objects Like Queries, Dashboards, Alerts and Warehouses
- SQL Warehouse (Compute) Design and Validation
- Databricks SQL Dashboard Workshop

Databricks Jumpstart

The Databricks Jumpstart packaged offerings help you accelerate your project timelines by familiarising you with the Databricks platform and key capabilities that follow best practices. You work with our industry-leading experts from the Databricks Professional Services team. The offerings reflect the most common questions our customers face during their early Databricks journey and provide you with a path for successful adoption.



Hadoop Migration

The Databricks Hadoop Migration Package gives you a concrete migration outcome through a prescriptive approach that maximises the value of your existing data and pipeline investments — and delivers a trouble-free migration. We tackle your key migration scenarios, ensuring the journey has a sound foundation and robust plan. And reference implementation serves as a design pattern for similar workloads. The packages include the following:

- Migrate data, code, ETL pipeline and job orchestration
- Execute pilot for ingestion and integration
- Performance testing
- Productionisation hardening best practice
- Code and performance optimization

Databricks offers Hadoop migration service in three tiers. Milestones and the outcome for each tier are produced by our prescriptive methodology, and each tier can be chained for accelerating realisation of your multicloud strategy and lakehouse vision.

- Hadoop migration foundation
 - Planning
 - Reference architecture
 - Reference implementation
- Hadoop migration extended
 - Implement second migration scenario
 - Orchestration integration or BI tool
- Hadoop migration optimised
 - Scaled performance optimization
 - CI/CD and automation
 - Monitoring/alerting/security

EMR Migration

The Databricks EMR Migration Package gives you a concrete migration outcome through a prescriptive approach that maximises the value of your existing data and pipeline investments — and delivers a trouble-free migration. We tackle your key migration scenarios, ensuring the journey has a sound foundation and robust plan. And reference implementation serves as a design pattern for similar workloads. The packages include the following:

- Migrate data, code, ETL pipeline and job orchestration
- Execute pilot for ingestion and integration
- Performance testing
- Productionisation hardening best practice
- Code and performance optimization

Databricks offers EMR migration service in three tiers. Milestones and the outcome for each tier are produced by our prescriptive methodology, and each tier can be chained for accelerating realisation of your multicloud strategy and lakehouse vision.

- EMR migration foundation
 - Planning
 - Reference architecture
 - Reference implementation
- EMR migration extended
 - Implement second migration scenario
 - Orchestration integration or BI tool
- EMR migration optimised
 - Scaled performance optimization
 - CI/CD and automation
 - Monitoring/alerting/security

Multicloud migration

The Databricks multicloud Migration package gives you a concrete migration outcome through a prescriptive approach that maximises the value of your existing Databricks investments (e.g., data engineering, ETL pipelines, ML workflows and supporting artefacts) – and delivers a trouble-free migration of Databricks workloads across different cloud providers.

The package combines our migration best practices, proven methodology and ready-made tooling to automate the migration of Databricks workloads, objects and configurations to a target cloud of your choice. It offers three tiers: Foundation, Extended and Optimised. Milestones and the outcome for each tier are produced by our prescriptive methodology, and each tier can be chained for accelerating realisation of your multicloud strategy and lakehouse vision.

- Multicloud migration foundation
 - Planning
 - Data, workspace and job migration
 - Validation
- Multicloud migration extended
 - For your second workspace:
 - Data, workspace and job migration
 - Validation
- Multicloud migration optimised
 - Cut-over planning and support
 - Hardening best practices
 - Target cloud CICD review
 - Target reference architecture review

Databricks Lakehouse Buildout

Build or optimise a Delta Lake as the structured transaction layer on top of low-cost storage following the most up-to-date best practices, laying the foundation of your lakehouse vision. The package addresses many of the common problems found with data lakes by leveraging Delta Lake and its associated architecture patterns. The package offers three tiers: Foundation, Extended and Optimised. Milestones and the outcome for each tier are produced by our prescriptive methodology, and each tier can be chained for greater impact on your journey of realising your lakehouse vision. Databricks will provide Services from the Technical Focus Areas in assisting customers on Lakehouse Buildout. Specific activities performed will vary, depending on Customer-specific objectives.

- Lakehouse buildout foundation
 - Target lakehouse architecture
 - Build a curated data pipeline
 - Develop Bronze, Silver and Gold
- Lakehouse buildout extended
 - Build a second curated data pipeline
 - Develop Bronze, Silver and Gold
- Lakehouse buildout optimised
 - Production readiness
 - CI/CD and automation
 - Orchestration and governance

Streaming Analytics

Build and optimise a Streaming Analytics workflow following the most up-to-date best practices, guided by experts. By leveraging structured streaming, Delta Lake and the associated architecture patterns, these packaged offerings address many of the common challenges faced when dealing with the complexities of delivering low-latency insights. The packages offer three tiers: Foundation, Extended and Optimised. Milestones and the outcome for each tier are produced by our prescriptive methodology. Each tier can be chained for greater impact on your streaming analytics exploration and adoption.

- Streaming analytics foundation
 - Reference architecture
 - Reference implementation of one streaming pipeline
- Streaming analytics extended
 - Implement additional pipeline
 - Operational visibility
- Streaming analytics optimised
 - Performance optimization
 - Workflow automation
 - Logging, monitoring, security

GenAI MVP

Develop a data strategy and custom Large Language Model (LLM) application leveraging Retrieval-Augmented Generation (RAG) to answer questions based on your organisation's knowledge base, in your own Databricks environment. Databricks will provide Services from the Technical Focus Areas in advising customers on its LLM strategy and implementation approach. Specific activities performed will vary, depending on Customer-specific objectives.

Machine Learning – Model MVP

Translate an existing business problem into an ML solution. This includes defining success criteria, performing exploratory data analysis, developing and benchmarking models, tracking and versioning models and data, as well as ensuring appropriate knowledge transfer to Customer's teams. Databricks will provide Services from the Technical Focus Areas to assist customers to develop an ML Model MVP. Specific activities performed will vary, depending on Customer-specific objectives.

Machine Learning – MLOps optimised

Accelerate time to production for your machine learning pipelines by developing a repeatable, production-grade MLOps framework that enables MLOps best practices at scale. Develop a production-ready repository set up with infrastructure as code, CI/CD workflows, and referenceable ML pipeline, customised for your use case. This includes developing an end-to-end ML pipeline production architecture, setting up batch, streaming and/or real-time inference, and establishing data and model governance with Databricks Unity Catalog and MLflow. Databricks will provide Services from the Technical Focus Areas in assisting customers on MLOps. Specific activities performed will vary, depending on Customer-specific objectives.

Customisable Services

Databricks Account and Engagement teams are open to customise the packages above to more closely meet your requirements via a defined Statement of Works to meet your requirements. This would be based on the day rates outlined in the SFIA Rate card.

Databricks Delivery Solutions Architect:

Databricks Delivery Solution Architect (DSA) can provide expert guidance in delivering use cases on the Lakehouse.

Delivery Expertise:

- Years of experience delivering large scale, technically complex use cases or projects (in the Data & AI domain)
- DSAs know how to deal with complexity and scale and how to get your Lakehouse into production

Technical Expertise

- Experience in building and optimising large scale data transformation pipelines using popular frameworks such as Apache Spark, Delta Lake and MLFlow.
- Knowledge in designing end-to-end architectures on the cloud that serve all your use cases.

Databricks Expertise

- A DSA is up-to-date with the latest Databricks product features, and will ensure your use cases are powered in the most cost optimised and simplified way possible by leveraging the latest and greatest Databricks technologies
- DSAs also are able to pull in the right expert Databricks resources for the right job

Databricks Delivery Solutions Architect (DSA) service can be requested in following tiered subscriptions:

- Delivery Solutions Architect subscription – Silver: 1 day / week
- Delivery Solutions Architect subscription – Gold: 2 days / week
- Delivery Solutions Architect subscription – Platinum: 4 days/ week

Databricks Resident Solutions Architect:

Databricks Resident Solution Architect (RSA) are Industry-leading experts on Spark, ML, Streaming, & Data Engineering. RSA can provide deep technical expertise solving your toughest data + AI problems at scale.

- Architect and implement workloads, test, and optimise pipelines (write production-grade code)
- Conduct expert advisory sessions
- CI/CD frameworks and coach which ones fit you best
- UC governance structure options and design approach
- Spark streaming best practices and optimization

Databricks Resident Solutions Architect service can be requested in following tiered subscriptions:

- Resident Solutions Architect subscription – Silver: 1 day / week
- Resident Solutions Architect subscription – Gold: 2 days / week
- Resident Solutions Architect subscription – Platinum: 4 days/ week

Databricks Support Services

At Databricks, we are committed to providing a comprehensive suite of support services designed to meet the varied needs of our customers. Our support packages are structured into four distinct tiers: Mission Critical, Production, Enhanced, and Business Support, each tailored to support different operational needs and urgency levels. Additionally, we offer the option of a Designated Support Engineer to augment any support package, providing personalised technical assistance.

Databricks will provide you with the level or type of Support Services specified in an Order in accordance with the Support Policy located at databricks.com/support. If Support Services are not specified in an Order, your support shall be limited to public Documentation and fora.

1. Mission Critical Support

Mission Critical is our premier support tier designed for customers who require immediate attention to business-critical issues. Features include:

- 24x7 Priority Team: Our dedicated team is available around the clock to manage communications with the customer and strategic stakeholders, driving urgency and ensuring timely resolution.
- Rapid Response: Customers receive a 15-minute response time for issues classified as Mission Critical.
- Escalation Management: Direct access to an Escalation Manager within 15 minutes of filing a mission critical ticket, ensuring high level oversight and swift action.
- Proactive Monitoring: Continuous monitoring of Production workspaces with direct alerts via Slack or Email for incidents undetectable by regional health monitoring systems.
- Chat Support Provided – direct access to Databricks Spark experts.
- Comprehensive Issue Resolution: Handling from intermediate technical issues to advanced troubleshooting, ensuring smooth operation of your production environments.

This tier is uniquely designed to provide robust support in critical situations, ensuring minimal disruption to your business operations.

2. Production Support

Production Support is geared towards customers operating in high stakes environments where uptime and reliability are paramount. This tier includes:



- **24x7 Team:** Our dedicated team is available around the clock to manage communications with the customer and strategic stakeholders, driving urgency and ensuring timely resolution for Urgent and High Severity incidents.
- **Fast Response:** Customers receive a 1 hour response time for issues classified as Urgent.
- **Chat Support Provided** – direct access to Databricks Spark experts.
- **Comprehensive Issue Resolution:** Handling from intermediate technical issues to advanced troubleshooting, ensuring smooth operation of your production environments.

3. Enhanced Support

Enhanced Support offers an elevated support experience ideal for clients who need more than standard support but less than the intensive care of Production or Mission Critical tiers. Further details can be found in the table below.

4. Business Support

Business Support provides foundational support suited for environments with less critical workloads such as Development and Staging environments.

Additional Services: Designated Support Engineer

For any of the above tiers, a Designated Support Engineer can be purchased to provide:

- **Tailored Technical Assistance:** Direct access to a dedicated engineer who understands your unique environment and needs.
- **Specialised Expertise:** Expert support in Spark/ML/AI/Delta/Streaming and edge technologies to maximise the value and strategic impact of our products.
- A more detailed review of the offering is available here ->
<https://cms.databricks.com/sites/default/files/legal/dse-overview-datasheet.pdf>

Our team is equipped to manage everything from basic queries to complex architectural challenges across the full spectrum of the Big Data ecosystem on Databricks. Whether you are dealing with day-to-day operations or enterprise-wide transformations, our support structure is designed to meet your needs at every level.

Databricks Support Policy

Databricks provides a number of plans that provide you with dedicated support and timely service for the Databricks platform and Apache Spark™

Feature	Business	Production	Mission Critical
Multicloud Support Support for Databricks on permitted Cloud Service Providers and Databricks-powered services; complimentary Success Credits available based on commitment size	—	✓	✓
Single-Cloud Support Support for platform services on a single-chosen Cloud Service Provider	✓	✓	✓

Support Portal Access Online repository of documentation, guides, best practices and more	✓	✓	✓
Updates and Patches Receive updates, bug fixes and patches without impact to your business	✓	✓	✓
Service Level Agreement (SLA) Receive support responses according to issue severity	Initial Response Time		
Mission Critical Severity	—	—	Within 15 minutes

A mission-critical production system is down or severely impacted such that routine operation is impossible			
Standard Support Severity 1 Production system is down or severely impacted such that routine operation is impossible	Within 1 business day	Within 1 hour	Within 1 hour
Standard Support Severity 2 Production issue where the system is functioning but in degraded or restricted capacity		Within 4 hours	Within 2 hours
Standard Support Severity 3 Issue where minor functionality is impacted or a development issue occurs		Within 1 business day	Within 4 hours

Standard Support Severity 4 Request for information or feature request with no impact on business operations			Within 1 business day
Support Service Hours Live support during customer's choice of time zone	9 AM–6 PM, business days	Severity 1 and 2: 24x7x365 Severity 3 and 4: 9 AM–6 PM, business days	Severity 1 and 2: 24x7x365 Severity 3 and 4: 9 AM–6 PM, business days
Databricks Standard Support Limited to break-fix support for the Databricks platform	✓	✓	✓

Databricks Chat Support Channel Per customer dedicated real-time messaging (e.g., Slack, Microsoft Teams) channel available during business hours for informal communication between the two teams, such as basic questions and information exchange	—	✓	✓
Max Number of Technical Contacts The number of technical contacts with access to the Databricks Help Center or Chat Support channel (if available)	2	8	16
Databricks Support for Spark Prioritized access to the world's leading Spark technical experts for troubleshooting problems using the product and services	Up to 2 hours/month	Up to 8 hours/month	Up to 8 hours/month

Additional Benefits — Mission Critical			
Proactive Monitoring			
Proactive monitoring of workspace-level health metrics by the Databricks Support Operations team	—	—	✓
Escalation Management			
Direct access to Databricks Escalation Managers for cases with mission critical severity	—	—	✓

Support Services for Multi-Cloud Training (MCT, formerly MosaicML) Platform Services

Supported Customers who subscribe to Business or Enhanced tier Support Services for MCT Platform Services will receive *Production-level SLA benefits during scheduled Hero Reservations.*

Unsupported Customers who schedule Hero Reservations must *purchase standalone Hero Reservation Support* for the period of their Hero Reservation. Customers with standalone Hero Reservations support will receive Production-level SLA benefits during the period of their Hero Reservation with the following differences:

Spark Support	No
Max Number of Technical Contacts	2

Further Details can be found at <https://www.databricks.com/support>

Definitions and Terms

- Support business hours are from Monday through Friday, from 9 AM to 6 PM (North America time zones (EST/CST/PST/MST), Central Europe (CET), Singapore/China (SGT/CST) and Australia Eastern (AET) Timezone). Support business hours exclude local holidays in each timezone.
- Chat Support channel is not covered under the Support SLA response times.
- "Contact" means a unique named user at customer (whether by email address, chat ID); accounts may not be shared.

- For reported mission critical severity issues, Databricks will assign an Escalation Manager to monitor and oversee case resolution and provide updates to the customer on a 15-minute recurring basis.
- Additional assistance beyond what is included in a support plan can be purchased as Advisory Services which will be delivered by our Professional Services team.
- If no support is specified in an order, the customer's support is limited to public documentation and open resources on the [Databricks Help Center](#).
- View the resources available through our [Customer Support Handbook](#) for more information on support definitions, processes and terms.

Training & Enablement

Databricks Training Services:

Data, analytics, and AI roles are growing, however a skill gap prohibits enterprise growth. Databricks has been the leader in data, analytics, and AI, and recognizes that the roles in these domains have grown and will continue to. There is a gap with where skills are and it halts continued organisational growth. The World Economic Forum shares that 44% of workers' core skills are expected to change in the next five years and that 60% will require training before 2027. Additionally, 60% growth of data and big data is expected by 2027. The skill gap is large and Databricks is prepared to provide comprehensive training to meet the needs of practitioners and organisations.

Databricks will generally provide the Training Services remotely, unless we agree to provide Databricks Services at a mutually agreed location at your request. If our personnel travel at your request, you agree to reimburse our reasonable, actually-incurred travel and lodging expenses. Pre-purchased Training Services will expire after the period indicated on the Order.

Depending on the Courses you Order, Training Services may be delivered (i) in scheduled real-time session(s) by instructors ("**Instructor-led Training Services**"), either virtually (online) or in person (as a private Course); or (ii) as a self-paced online training course ("**Self-Paced Training Services**"). For Training Services, we provide qualified training personnel and/or suitable training materials. You will not provide Databricks any Customer Content for use with the Training Services, including in any shared temporary Workspaces we may make available to Students to facilitate demonstration and learning during a Course.

If we have agreed to provide you with Instructor-led Training Services in person (e.g. a private Course), then in support of our delivering such Training Services, you will (i) be responsible for providing a reasonable training venue, and all training venue expenses and arrangements, including providing us access to appropriate contacts to assist in logistics coordination, audio-visual and the like; (ii) provide Databricks access to your designated training venue as reasonably required by Databricks to setup and perform the Instructor-led Training Services.

Key Benefits:

- Accelerate business outcomes: Achieve the full potential of your data, analytics, and AI initiatives on the lakehouse
- Deploy industry-leading technology: Utilise cutting-edge and proven technology to drive results for your organisation
- Foster career development: Enable your organisation with role-based trainings for career acceleration and employee retention

To help customers increase productivity and accelerate business outcomes, Databricks offers many role-based learning pathways across key data, analytics, and AI roles. Each learning pathway starts with Databricks Fundamentals, where learners can gain exposure to the basics of the data intelligence platform. With each of the roles, there are accreditations (short quizzes, unproctored) or certifications (exams, proctored) to test a learner's knowledge. The badges that come from accreditations and certifications can be shared on resumes and LinkedIn to showcase a learner's skills.

Learning Pathways for All Roles: Tailor made for distinct personas




Figure 1 - Databricks Learning Pathways

Databricks role based learning pathways

Databricks has learning solutions for every data and AI role, but also takes into account where learners are in their learning journey. From self-paced training on Databricks Academy, which includes role-based learning pathways to certification

overviews, to Databricks Academy Labs, hands-on, practical learning in Databricks environments to Blended Learning, cohort-based training taught by expert instructors with certification preparation, there is a learning solution for every practitioner.

For organisations that want to focus and prioritise their team's enablement, Databricks has learning solutions for enterprise organisations, as well. From public and private instructor-led training in ½ day to two-day formats to learning subscriptions that include instructor-led training and Skills @ Scale, which includes cohort-based learning, private onboarding sessions, and certification vouchers, each of these offerings can be provided as a standalone or in a bundle.

Stand-alone training offerings:

Self-paced training	<u>Databricks Academy Labs</u>	<u>Blended Learning</u> (no minimum # of users)
Role-based learning pathways Get started courses Certification overviews	Hands-on, practical learning available on-demand	Cohort based training taught by expert instructors Certification preparation

Bundled training offerings:

<u>Public/private instructor-led training (ILT)</u>	<u>Learning subscription</u>	<u>Skills @ Scale</u> (100+ users)
½ day–2 day courses taught by leading experts	Instructor-led training Certification vouchers	Blended learning Private onboarding sessions Certification vouchers & preparation

2. Data Protection

Information Assurance

Our trusted platform is built by embedding security throughout the software development and delivery lifecycle. We follow rigorous operational security practices such as penetration testing, vulnerability assessments and strong internal access controls. We believe transparency is the key to winning trust — we publicly share how we operate, and work closely with our customers and partners to address their security needs. We have offerings for Cyber Essentials Plus, PCI-DSS, HIPAA, DOD IL5 and FedRAMP compliance, and we are ISO 27001, ISO 27017, ISO 27018 and SOC 1, 2 & 3 Type II compliant.

Beyond the documentation and best practices that you will find in our Security and Trust Center, we also provide a contractual commitment to security written in plain language to all our customers. This commitment is captured in the Security Addendum of our customer agreement, which describes the security measures and practices that we follow to keep your data safe.

The Security Addendum is incorporated into and made a part of the written agreement between Databricks, Inc. (“Databricks”) and Customer that references this Security Addendum (“Agreement”).

Databricks maintains a comprehensive documented security program that is based on industry standard security frameworks including NIST 800-53 and ISO 27001 (the “Security Program”). Pursuant to the Security Program, Databricks implements and maintains administrative, physical, and technical security measures to protect the Platform Services and Support Services and the security and confidentiality of Customer Content (including any Personal Data that may be contained therein) (each as defined in the Agreement) under Databricks’ control that is processed by Databricks in its provisioning of the Platform Services or Support Services (the “Security Measures”). Databricks’ compliance with this Addendum shall be deemed to satisfy any more general measures included within any Agreement, including the Service Specific Terms.

Databricks regularly tests and evaluates its Security Program, and may review and update this Security Addendum at any time without notice, provided that such updates are equivalent (or enhance) security and do not materially diminish the level of protection afforded to Customer Content by these Security Measures.

Architecture. Databricks is a hybrid platform-as-a-service offering. The components responsible for managing and controlling the Platform Services are referred to as the ‘Databricks Control Plane’. The compute resources that perform data processing operations are referred to as the “Data Plane”. For certain Cloud Service Providers, the Data Plane may either be deployed in Customer’s Cloud Service Provider account (known as the ‘Customer Data Plane’) or, for Databricks Serverless Compute, in a Databricks-controlled Cloud Service Provider account (known as the ‘Databricks Data Plane’). Data Plane shall refer to both Customer Data Plane and Databricks Data Plane unless otherwise specified.



Shared Responsibility. Databricks operates in a shared responsibility model, where both Databricks and Customer maintain security responsibilities.

Most Customer Content is stored within Customer's own Cloud Service Provider account at rest (e.g., within Customer's AWS S3 bucket) or within other Systems under Customer's control. Please see the Documentation for more details.

Databricks Control. Small amounts of Customer Content may be stored within the Databricks Control Plane, including metadata about Customer Content. Databricks offers Customers options regarding the storage of certain Customer Content within the Platform Services such as query results. Please see the Documentation for more details.

Customers may specify the region(s) where their Platform Services Workspaces are deployed. Customers can choose to deploy the Data Plane into any supported Databricks region. The Databricks Control Plane is deployed into the same region. Databricks will not, without Customer permission, move a Customer Workspace into a different region.

Databricks' Audits & Certifications. Databricks uses independent third-party auditors to assess the Databricks Security Program at least annually, as described in the following audits, regulatory standards, and certifications:

Cyber Essentials Plus

HITRUST

SOC 1 Type II

SOC 2 Type II

SOC 3 Type II

ISO 27001

ISO 27017

ISO 27018

ISO 27036

HIPAA

PCI DSS

FedRAMP High

DoD IL5

Databricks' Chief Security Officer leads the Databricks' Information Security Program and develops, reviews, and approves (together with other relevant internal stakeholders) Databricks' Security Policies and Procedures.

Databricks maintains a documented change management policy, reviewed at least annually.

G-Cloud 15



Databricks has implemented a formal Information Security Management System (“ISMS”) in order to protect the confidentiality, integrity, authenticity, and availability of Databricks' data and information systems, and to ensure the effectiveness of security controls over data and information systems that support operations.

Databricks employs monitoring and logging technology to help detect and prevent unauthorised access attempts to its network and equipment.

Active users with privileged access to the Platform Services are reviewed at least quarterly and are promptly removed upon termination of employment. As part of the personnel offboarding process, all accesses are revoked and data assets are securely wiped.

Databricks maintains a comprehensive Third Party Risk Management program that assesses the security compliance of applicable third parties, including vendors and subprocessors, in order to appropriately measure and manage risk.

Personnel receive comprehensive training on the Security Policies upon hire and refresher training annually. Personnel are required to certify and agree to the Security Policies and personnel who violate the Security Policies are subject to disciplinary action, including warnings, suspension and up to (and including) termination.

All personnel undergo background checks prior to onboarding (as permitted by local law), which may include, but are not limited to, criminal record checks, employment history verification, education verification, and global sanctions and enforcement checks. Databricks uses a third-party provider to conduct screenings, which vary by jurisdiction and comply with applicable local law. Personnel are required to sign confidentiality agreements.

Physical and Environmental Controls

Databricks Corporate Offices. Databricks has implemented administrative, physical, and technical safeguards for its corporate offices. These include, but are not limited to, the below:

Visitors are required to sign in, acknowledge and accept an NDA, wear an identification badge, and be escorted by Databricks personnel while on premises

Databricks personnel badge into the offices

Badges are not shared or loaned to others without authorization

Physical entry points to office premises are recorded by CCTV and have an access card verification system at every door, allowing only authorised employees to enter the office premises

Equipment and other Databricks-issued assets are inventoried and tracked

Office Wi-Fi networks are protected with encryption, wireless rogue detection, and Network Access Control



Cloud Service Provider Data Centers. Databricks regularly reviews Cloud Service Provider audits conducted in compliance with ISO 27001, SOC 1, SOC 2, and PCI-DSS. Security controls include, but are not limited to the list below:

Biometric facility access controls

Visitor facility access policies and procedures

24-hour armed physical security

CCTV at ingress and egress

Intrusion detection

Business continuity and disaster recovery plans

Smoke detection sensors and fire suppression equipment

Mechanisms to control temperature, humidity and water leaks

Power redundancy with backup power supply

Databricks leverages multiple layers of network security controls, including network-level isolation, for separation between the Databricks' development and production environments.

Firewalls are implemented as network access control lists or security groups within Databricks' production environment.

Databricks employs industry standards to harden images and operating systems under its control that are deployed within the Platform Services, including deploying baseline images with hardened security configuration such as disabled remote root login, isolation of user code, and images are regularly updated and refreshed.

For Systems under Databricks control supporting the production data processing environment, Databricks tracks security configurations against industry standard baselines such as CIS and STIG.

Encryption of data-in-transit. Customer Content is encrypted using cryptographically secure protocols (TLS v1.2 or higher) in transit between (1) Customer and the Databricks Control Plane and (2) the Databricks Control Plane and the Data Plane. Additionally, depending on functionality provided by the Cloud Service Provider, Customers may optionally encrypt communications between clusters within the Data Plane (e.g., by utilising appropriate AWS Nitro instances).

Customer Content is encrypted using cryptographically secure protocols (AES-256 bit, or the equivalent or better) while at rest within the Databricks Control Plane.

Cryptographic standards are periodically reviewed and selected technologies and ciphers are updated in accordance with assessed risk and market acceptance of new standards.

Customer(s) may choose to leverage additional encryption options for data in transit within the Customer Data Plane or Databricks Data Plane as described in the Documentation (e.g., Customer may utilise AWS Nitro EC2 instances within the Customer Data Plane to provide additional encryption in transit). Customer shall, based on the sensitivity of the Customer Content, configure the Platform Services and Customer Systems to encrypt Customer Content where appropriate (e.g., by enabling encryption at rest for data stored within AWS S3).

Databricks leverages security capabilities provided natively by Cloud Service Providers for security detection.

Databricks generates audit logs from Customer's use of the Platform Services. The logs are designed to store information about material events within the Platform Services.

Customers may, depending on the entitlement tier of the Platform Services, enable delivery of audit logs. It is Customer's responsibility to configure this option.

Databricks stores audit logs in a manner designed to protect the audit logs from tampering.

Databricks stores audit logs for at least one year.

Databricks conducts third-party penetration tests at least annually, in addition to maintaining in-house offensive security personnel and a public bug bounty program.

Databricks regularly runs authenticated scans against representative hosts in the SDLC pipeline to identify vulnerabilities and emerging security threats that may impact the Platform Services. Databricks will use commercially reasonable efforts to address critical vulnerabilities within 14 days, high severity within 30 days, and medium severity within 60 days measured from, with respect to publicly declared third party vulnerabilities, the date of availability of a compatible, vendor-supplied patch, or for internal vulnerabilities, from the date such vulnerability is confirmed. Databricks leverages the National Vulnerability Database's Common Vulnerability Scoring System (CVSS), or where applicable, the U.S.-Cert rating, combined with an internal analysis of contextual risk to determine criticality.

Databricks deploys new code to the Databricks Control Plane on an ongoing basis.

New Data Plane virtual machines use the latest applicable source code and system images upon launch. Customers are encouraged to restart always-on clusters on a periodic basis to take advantage of security patches.

Databricks personnel are authenticated through single sign-on (SSO), 802.1x (or similar) where applicable, and use a unique user ID and password combination and multi-factor authentication. Privileges are consistent with least privilege principles. Security Policies prohibits personnel from sharing or reusing credentials, passwords, IDs, or other authentication information. If your identity provider supports the SAML 2.0 protocol, you can use Databricks' SSO to integrate with your identity provider.



Databricks enforces RBACs (based on security groups and access control lists). Only authorised roles, which are defined based on the principle of least privilege and segregation of duties, are allowed to access production systems.

Databricks enforces certain security controls on its workstations used by personnel, including:

Full-disk encryption

Anti-malware software

Automatic screen lock after 15 minutes of inactivity

Secure VPN

Breach Detection & Response

Databricks' dedicated Detection engineering team deploys and develops intrusion detection monitoring across its computing resources, with alert notifications sent to the Security Incident Response Team (SIRT) for triage and response. The SIRT employs an incident response framework to manage and minimise the effects of unplanned security events.

"Security Breach" means a breach of security leading to any accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to Customer Content under Databricks control. A Security Breach shall not include an unsuccessful attempt or activity that does not compromise the security of Customer Content, including (without limitation) pings and other broadcast attacks of firewalls or edge servers, port scans, unsuccessful log-on attempts, denial of service attacks, packet sniffing (or other unauthorised access to traffic data that does not result in access beyond headers) or similar incidents. Databricks maintains a record of Security Breaches that includes description, dates and times of relevant activities, and disposition. Suspected and confirmed security incidents are investigated by security, operations, or support personnel; and appropriate resolution steps are identified and documented. For any confirmed Security Breaches, Databricks will take appropriate, reasonable steps to minimise product and Customer damage or unauthorised disclosure.

In accordance with applicable data protection laws, Databricks will notify Customer of a Security Breach for which that Customer is impacted without undue delay after becoming aware of the Security Breach, and take appropriate measures to address the Security Breach, including measures to mitigate any adverse effects resulting from the Security Breach.

Upon written request and at no additional cost to Customer, Databricks shall provide Customer, and/or its appropriately qualified third-party representative, access to reasonably requested documentation evidencing Databricks' compliance with its obligations under Security Addendum in the form of the relevant audits or certifications listed above. Such audits are performed (a) at least once annually; and (b) by independent third-party security professionals selected by Databricks. Such audits result in the generation of a confidential audit report collectively, "Audit Reports").

Only to the extent Customer cannot reasonably satisfy Databricks compliance with Security Addendum through the Audit Reports, Customer may send a written request to conduct an audit of Databricks applicable controls during the term of the Agreement on an annual basis. Following receipt by Databricks of such a request, Databricks and Customer shall mutually agree in advance on the details of the audit, including the reasonable start date, scope and duration of, and security and confidentiality controls applicable to, any such audit. The Audit Report, audit, and any information arising therefrom shall be considered Databricks Confidential Information and may only be shared with a third party (including a third party controller) with Databricks' prior written agreement.

Notwithstanding any other audit provisions in the Agreement, Customer requests for audits are limited to once per year.

Databricks Business Continuity (BC) and Disaster Recovery (DR) plans are reviewed and drills are conducted annually.

Data Resiliency. Databricks performs backups for the Databricks Control Plane, generally managed by the Cloud Service Provider capabilities, for data resiliency purposes in the case of a critical systems failure. While Databricks backs up certain service elements that persist in the Databricks Control Plane as part of its systems resiliency, those backups are maintained only for emergency recovery purposes and are not available for Customer.

No Data Restoration. Due to the hybrid nature of the Databricks Platform, Databricks does not back up Customer Content. To assist Customers in performing their own backups, Databricks provides certain self-service features within the Platform Services (like the ability to synchronise notebooks with a customer's Github or Bitbucket account).

The Platform Services provide Customer with functionality that permit Customer to delete Customer Content under Databricks' control.

Customer Content contained within a Customer Workspace is permanently deleted within thirty (30) days following cancellation of the Workspace.

MosaicML Security Standards. Notwithstanding anything to the contrary in the Agreement including its schedules or annexes, the security standards and measures set forth in the Agreement ("Security Measures") do not apply to products and services related to MosaicML ("MosaicML Services"). Instead, the following section comprises the Security Measures for the purpose of MosaicML Services. For clarity, the Security Measures described below do not apply to any other Databricks Services.

Databricks implements and maintains commercially reasonable technical, physical, administrative, and organisational safeguards, and other applicable aspects of system and content management, designed to protect the security, confidentiality, integrity, and availability of all Customer Content. Provider shall use commercially reasonable efforts to promptly notify Customer in the event of any breach of security impacting Customer

Content (“Security Breach”) within three (3) business days from the date the Security Breach is confirmed.

Privacy by Design

Databricks takes privacy seriously. We understand that the data you analyse using Databricks is important both to your organisation and your customers, and may be subject to a variety of privacy laws and regulations. We incorporate privacy by design into our product development methodology, with input from our legal privacy team throughout the development lifecycle.

Databricks itself is compliant with key privacy regulations such as the GDPR, CCPA, Swiss Data Protection Act, and UK Data Protection Laws, which are included in our default DPA language (<https://databricks.com/dpa>). Additionally, due to the way the Databricks platform is architected, customers retain control of most of their Customer Data within their own cloud service provider storage buckets and generally store little Customer Personal Data within the Databricks services themselves.

3. Using the service

Ordering and Invoicing

When a customer considers deploying Databricks would recommend getting in contact via gcloud@databricks.com to discuss requirements and how best to provide a solution. Databricks can then align account management, technical solutions support and a delivery partner as appropriate to requirements.

Databricks can provide assistance to the customer in completing the Call-Off contract once the buyer makes a request for our services, using terms that are consistent with existing agreements across other government organisations.

Pricing Overview

In the Pricing document with this service detailed pricing information is provided.

Part A – Databricks Support Services refers to the Support tiering related to the cloud platform upon which Databricks is deployed (either AWS, GCP or Multicloud).

Part B – Databricks Training refers to the Training services and Certification on offer for both public and private instructor led training. Furthermore Education Subscriptions are bundled offers for learning packages provided.

Part C – Databricks Professional Services Offerings & Packages includes pricing for prepackaged Professional Services. To meet the Buyer's specific requirements if any adjustment is needed to the packages then we will use the SFIA rate card to tailor the package to meet the Buyer's specific needs.

Implementation Plan

Databricks provides an Account Manager to support our customers and build an ongoing relationship and commercial requirements. In addition, our account teams include Solution Architects who can provide technical advice and guidance so that customers can leverage the Databricks platform to deliver their business outcomes.

The core account team of our customers are supported through their entire journey via various programs, to ensure they are constantly engaged and updated with the latest developments, so they can decide on an optimised method to ensure their success. A sample of the various options available to customers are :




Fig 2 - Databricks Engagement

The ongoing customer engagement is also maintained via a governance process, agreed during the initial kickoff meeting. In addition to the Account Manager, the Databricks team comprises of Solution Architects, Delivery Solution Architects, Engagement Managers, Project Managers, Resident Solution Architects, Solution Engineers along with a plethora of other functional teams who are available to work with the customer throughout the lifecycle to ensure success.




Fig 3 - Databricks Customers

Project Management and Delivery

As part of the customer journey, we follow the below milestones based approach. Each milestone represents the key elements of the experience of becoming a new customer of Databricks, ending in successful realisation of the platform's potential for business value to an organisation. Each step in the customer journey is a necessary component in the natural maturation of a customer from onboarding through to production and achieving value.



Fig 4 - Implementation Plan

The overall customer journey can be summarised as below:




Fig 5 - Customer Journey

During this process, after the customer has been onboarded to Databricks and the appropriate solution has been scoped, the actual delivery is carried out by our team of experienced Consultants and Architects. Our team of Delivery Solution Architects, Engagement Managers and Project Managers provide service reports at agreed timelines to the appropriate stakeholders. Some suggested sample project milestones followed by Databricks are:



Fig 7 - Project flow

Methodologies

We follow an agile methodology as our implementation approach. This enables us to demonstrate value within short timescales. For instance, our recommended timeline for a POC is around 2 weeks.




Fig 8 Agile practise

Reporting

The health and status of the projects are tracked and reported back to the stakeholders in the form of weekly and monthly meetings and reports. We also track the overall success of the engagement and value delivered to the customer through the quarter business reviews and implement any course corrections, as required.

Service Management

There are many components to the Databricks product – the user front end, APIs, the Databricks Runtimes (DBRs), and the backend architectural / configuration options. Each of these receives regular updates. As an example, the DBR is the execution engine which powers workloads running on Databricks. The DBR is regularly updated with bug fixes and performance updates. Whilst there is no strict schedule, with releases occurring when ready, historically minor versions are released monthly, with major versions released biannually or less. Databricks releases runtimes as Beta and GA versions. Databricks supports GA versions for six months, unless the runtime version is:

- A Long Term Support (LTS) version. LTS versions are released every six months and supported for two full years.
- A Databricks Light runtime. Light versions are supported until either 12 months after release or two months after the next Databricks Light release, whichever comes first.

Service Constraints

Databricks support service constraints:

- Support business hours are from Monday through Friday, from 9 AM to 6 PM (North America time zones (EST/CST/PST/MST), Central Europe (CET), Singapore/China (SGT/CST) and Australia Eastern (AET) Timezone). Support business hours exclude local holidays in each timezone.
- Chat Support channel is not covered under the Support SLA response times.
- “Contact” means a unique named user at customer (whether by email address, chat ID); accounts may not be shared.
- For reported mission critical severity issues, Databricks will assign an Escalation Manager to monitor and oversee case resolution and provide updates to the customer on a 15-minute recurring basis.
- Additional assistance beyond what is included in a support plan can be purchased as Advisory Services which will be delivered by our Professional Services team.
- Support is only available in English.
- If no support is specified in an order, the customer’s support is limited to public documentation and open resources on the [Databricks Help Center](#).
- View the resources available through our [Customer Support Handbook](#) for more information on support definitions, processes and terms.



Service Levels

Service uptime is agreed and documented in the Order Form between both parties.

Support levels are contingent on the support version purchased. There tiers of support available are here:

Databricks provides a number of support options to GCP Databricks customers with pre-paid consumption plans that provide them with dedicated support and timely service for the Databricks platform.

FEATURE	BUSINESS	PRODUCTION	MISSION CRITICAL
MULTI-CLOUD SUPPORT Support for Databricks on permitted Cloud Service Providers and Databricks-powered services;	-	✓	✓
SINGLE-CLOUD SUPPORT Support for platform services on a single-chosen Cloud Service Provider	✓	✓	✓
SUPPORT PORTAL ACCESS Online repository of documentation, guides, best practices and more	✓	✓	✓
UPDATES and PATCHES Receive updates, bug fixes and patches without impact to your business	✓	✓	✓
SERVICE LEVEL AGREEMENT (SLA) Receive support responses according to issue severity	Initial Response Time		
MISSION CRITICAL SEVERITY A mission critical production system is down or severely impacted such that routine operation is impossible	-	-	Within 15 minutes
STANDARD SUPPORT SEVERITY 1 Production system is down or severely impacted such that routine operation is impossible	Within 1 business day	Within 1 hour	Within 1 hour

STANDARD SUPPORT SEVERITY 2 Production issue where the system is functioning but in degraded or restricted capacity		Within 4 hours	Within 2 hours
STANDARD SUPPORT SEVERITY 3 Issue where minor functionality is impacted or a development issue occurs		Within 1 business Day	Within 4 hours
STANDARD SUPPORT SEVERITY 4 Request for information or feature request with no impact on business operations			Within 1 business Day
SUPPORT SERVICE HOURS Live support during customer's choice of time zone	9 AM–6 PM, business days	Severity 1 and 2: 24x7x365 Severity 3 and 4: 9 AM–6 PM, business days	Severity 1 and 2: 24x7x365 Severity 3 and 4: 9 AM–6 PM, business days
DATABRICKS STANDARD SUPPORT Limited to break-fix support for the Databricks platform	✓	✓	✓
DATABRICKS CHAT SUPPORT CHANNEL Per customer dedicated real-time messaging (e.g., Slack, Microsoft Teams) channel available during business hours for informal communication between the two teams, such as basic questions and information exchange	-	✓	✓
MAX NUMBER OF TECHNICAL CONTACTS The number of technical contacts with access to the Databricks Help Center or Chat Support channel (if available)	2	8	16
DATABRICKS SUPPORT FOR SPARK Prioritised access to the world's leading Spark technical experts for troubleshooting problems using the product and services	Up to 2 hours/month	Up to 8 hours/month	Up to 8 hours/month

ADDITIONAL BENEFITS — MISSION CRITICAL			
PROACTIVE MONITORING Proactive monitoring of workspace-level health metrics by the Databricks Support Operations team	-	-	Yes
ESCALATION MANAGEMENT Direct access to Databricks Escalation Managers for cases with mission critical severity	-	-	Yes

Outage and Maintenance Management

Databricks does not provide SLA details as they are contractual documents managed through the legal process. The current status of all Databricks availability regions can be checked and subscribed to from <https://status.databricks.com>.

Maintenance and updates to the Databricks platform are carried out without interruption to the service due to the hybrid nature of the Databricks Platform as a Service – new features, updates and maintenance releases are integrated without impact.

Downtime SLAs are negotiated on a per contract basis.

Financial Recompense Model for not Meeting Service Levels

Not offered with this service.

4. Provision of the service

Customer Responsibilities

In a professional services engagement with Databricks, the customer has several responsibilities to ensure successful delivery of the offering. These responsibilities are typically outlined in a Checklist associated with each offering. The Checklist stipulates all the items that need to be at a particular state to ensure the Databricks team can complete the deliverables for a given package.¹

For instance, the customer may need to provide Databricks with access to certain resources (people, environments and/or documents), ensure that prerequisites are met, and collaborate with the Databricks team on various aspects of the project. The customer is also responsible for accepting the deliverables based on the acceptance criteria agreed upon at the outset of the project.

Moreover, the customer is expected to work closely with the Databricks team, including participating in regular status meetings and supporting issue resolution and risk mitigation.

It's important to note that the specific responsibilities may vary depending on the nature of the service offering and the complexity of the customer's environment. Therefore, it's crucial to review the specific Checklist and Statement of Work (SOW) associated with each offering for detailed responsibilities.

Technical Requirements and Client-Side Requirements

Customers will need an AWS, GCP or Azure cloud environment, with the ability to access Databricks via the respective cloud marketplace (in the case of AWS/GCP, first party service on Azure).

Outcomes/Deliverables

This will vary from the specific delivery package and has been outlined above in more detail. Typical delivery assets include project plans, architecture design documents, weekly status reports, risk/issue tracking, project close out documentation and training delivery assets.

Development life cycle of the solution

We follow an agile methodology as our implementation approach. This enables us to demonstrate value within short timescales, typically following 2 week sprint cycles (can vary depending on customer requirements). For instance, our recommended timeline for a POC is around 2 weeks.


After-sales Account Management

Databricks is a value-led, software and services business. and values customer obsession as a key business objective. It has an ongoing engagement with customers, with resources are based globally, but crucially for UKG, in the UK and Europe, including security cleared personnel.

The lead role for after-sales engagement is the Databricks Delivery Solution Architect , who engages in order to provide:-

- Technical account manager
- User onboarding & Value Realisation
- Project management & triaging customer issues
- Technical Escalation and Resolution
- Joint Success Planning and Realisation
- Business and stakeholder Liaison

Their engagement is part of a wider supporting ecosystem with Databricks:-



Your Team + The Databricks Team = Success

Fig 8 - Databricks Team

Termination Process

The specifics of the termination process will be outlined in the Call-Off Contract based on the service offered.

5. Our experience

Clients

Below is an example of UK Public Sector customers who benefit from the power of Data and AI through the Databricks platform.



Fig 22 - Public Sector Customers

Global adoption

OVER 12,000 CUSTOMERS, ACROSS EVERY INDUSTRY



Fig 23 – Private Sector Customers

Contact Details

gcloud@databricks.com