

Center of Excellence (CoE)

This Service Description describes the scope of services (including associated Customer requirements and benefits) for the Databricks subscription service identified above (the “**Service**”), and applies to the Service under Customer’s accepted Services Order.

Accelerating past the challenges of Data & AI at scale

Databricks Lakehouse Center of Excellent (CoE) provides flexible access to a concentration of Databricks expertise and resources. Pre-packaged annual or multi-year options help you develop internal expertise to build your own Center of Excellence on the lakehouse, mitigate risk, and accelerate time to value.

	SILVER	GOLD	PLATINUM
Benefits	Access to expertise, guidance, best practices plus ongoing design & build support across all four CoE Focus Areas.	All the features of Silver, plus reference implementations of CoE capabilities across all four CoE Focus Areas.	All the features of Gold plus end-to-end white glove CoE design & build support across all four CoE Focus Areas.
Engagement (by Service type)			
Professional Tech/Strategy	<ul style="list-style-type: none"> Up to 2 Days/Week/Year Up to 1 Days/Week/Year 	<ul style="list-style-type: none"> Up to 3 Days/Week/Year Up to 2 Days/Week/Year 	<ul style="list-style-type: none"> Up to 4 Days/Week/Year Up to 4 Days/Week/Year
Enablement			
Skills@Scale	<ul style="list-style-type: none"> For up to 100 Students 	<ul style="list-style-type: none"> For up to 200 Students 	<ul style="list-style-type: none"> For up to 400 Students

Service Types and Benefits

- “**Professional**” – *Professional Services*, delivered by Resident Solution Architect (**RSA**) or equivalent.
- “**Tech/Strategy**” – *Technical Strategy & Adoption*, delivered by Delivery Solution Architect (**DSA**) or equivalent.
- “**Skills@Scale**” – *Enablement*, delivered by combination of cohort-based Blended Learning, role-based onboarding classes, public half-day classes, private user group and access to certification vouchers – all, orchestrated with the assistance of a Customer Enablement Architect (**CEA**).

View the [Focus Areas Matrix](#) (below) for a set of representative activities and resourcing available within the Professional, Tech/Strategy, and Skills@Scale Service types (tagged as **RSA**, **DSA**, and **CEA** respectively).

For more information on **Skills@Scale**, including benefit details and other terms, visit: <https://www.databricks.com/professional-services/service-descriptions>

CoE focus areas provide the assistance you need

Leading-edge projects come with risks, technical challenges and unforeseen obstacles. Avoid time-wasting trial and error, costly rework of the implementation by working with our experienced professionals to accelerate your outcome. The following provides representative activities Databricks CoE experts can engage with the Customer to help jointly set up the Customer's Lakehouse CoE practice.

* Each representative activity or a group of activities are tagged by role either as **RSA**, **DSA** or **CEA** in a parenthesis at the end of the activity description denoting the primary resource who'll be working on it during delivery.

Focus Areas Matrix	
4 Focus Areas	Representative Activities
Focus Area 1: Data & AI blueprint	
Data and Model Governance	<p>Together, we'll assess and document your lakehouse data estate and governance roadmap (RSA):</p> <ul style="list-style-type: none">• Current state of data architecture, ecosystem, & processes• Review current workspaces, data, storage and catalogs• Review Customer's requirements to provide best practices for data governance, management, lineage, sharing, and discoverability• Documentation including target architecture design• Inventory existing metastores & objects• Metastore registration and design• Migration strategy of existing metastore (internal/external) to Unity Catalog• Review access control patterns to manage users and groups• Review Unity Catalog enabled clusters and SQL endpoints• Review table and view creation, file access
Well-Architected Lakehouse	<p>Together, we'll assess and document your lakehouse architecture against Databricks' Well-Architected Lakehouse framework (RSA):</p> <ul style="list-style-type: none">• Lakehouse Operations - Operations processes that keep a system running in production• Security, Privacy and Compliance - Protecting applications and data from threats• Platform Reliability - The ability of a system to recover from failures and continue to function• Performance and Efficiency - The ability of a system to adapt to changes in load• Cost Optimization - Managing costs to maximize the value delivered• Interoperability and Usability - Processes in place are easy to use for users and compatible with other tools and systems

Focus Area 2: Lakehouse Operations

Platform Readiness

Working together, to assess and document a 12 month Lakehouse roadmap and how your platform owner will onboard new features where required to meet use case requirements. Based on that assessment of platform features and intended use cases support needed, develop standard operating procedure (SOP) on how to monitor various aspects of the Data estate and platform costs (**RSA**).

Operations Design

Collaborate together to assess and document the following standard operating procedures for (**RSA**):

- Running use cases in production on the Lakehouse
- Deploying new use cases in production on the Lakehouse
- Upgrading existing use cases in production on the Lakehouse
- Monitoring existing use cases in production on the Lakehouse

In addition, we'll conduct an assessment and provide recommendation on areas that improves your future scalability in operationalizing new and existing use cases (**RSA**).

Intake & Platform Management

Working together, assess and document how the existing use cases are deployed and maintained in the Lakehouse. The assessment focuses on how you can ensure use cases are still performant with known consumption patterns. In addition, review how you are automating user onboarding onto the Lakehouse based on personas: Data Engineering, Data Science and Business Analysts. We'll provide a summary of findings and areas of recommendation guided by Databricks best practices (**RSA**).

Integrate DevOps Practices

Working together, assess and document recommendation and an action plan that improves overall Lakehouse CI/CD practices (**RSA**).

Centralized Reporting & Monitoring

In many cases, the Customer might only have project-based scope and visibility. In CoE, having a broader, centralized view and reporting across all projects and use cases are helpful to (**RSA**):

- Have the full visibility and collect right KPIs / measurements
- Drive right changes from usage pattern, cost, user access etc.

With that in mind, work together to review/assess the following:

- Usage monitoring by use case
- User access
- Aggregated TCO

Self-Service Onboarding for Users and Teams

Work together to review existing Lakehouse operations with the objective of producing a roadmap of improvements/best practices in collaboration with the Customer. The following areas will be covered (**RSA**):

- User self service
- Data governance
- Monitored and measurable team platform usage

Focus Area 3: Use Case Acceleration

Identification & Prioritization

Collaborate to assess and document how the business can raise and capture new use cases for the CoE team. Recommendations will focus on (**DSA**):

- Applicable use case categories for the Lakehouse
- How the CoE will manage businesses new use case request
- Details to be captured for implementation based on use case category
- Process to review use case backlog and how use cases are prioritized for delivery

Delivery Execution Process

Work together to review how the CoE manages use case delivery support into the business (**DSA**):

- Document the resource profile within the CoE and how business teams can make requests for resourcing assistance from the CoE.
- Document how the CoE resource pool will engage the businesses in their use case delivery.

Define Outcomes & Metrics

Work together to assess how use case business outcomes are mapped against the agreed metrics documented during the use case requirements process (**DSA**).

Reusable Accelerators & Assets

Work together to review your current library of reusable assets such as pipeline development framework, code snippets, deployment scripts and configurations. Databricks will present additional asset type recommendations (if any). In addition, review how your internal business units collaborate with this library and overall asset management practice and governance (e.g. keeping these reusable assets up to date). Work together to develop programs and plans as needed to drive awareness and adoption (**RSA**).

Production Readiness at Scale

Work together to review past use case deployments and how this CoE has improved over time in deployment efficiency, as mutually agreed upon between Customer and Databricks. We review the lessons learned and evaluate how the CoE is managing user adoption for new use cases deployed into production (**DSA, RSA**).

Enterprise Standards & Practices

Work together to review Customer's current practices employed in Customer's CoE. The review is focused on how the CoE continues to iterate to adapt to your Enterprise and business changes (**DSA**).

Curated Asset Repository

Work together to review the degree to which Customer's organization is engaged and leveraging the CoE Asset repository produced from the engagement under these Services (**DSA, RSA**).

Focus Area 4: Organization Enablement

In-House CoE Champions

Work together to develop a mutual enablement plan for the CoE and identify target audience in the organization. This plan documents the personas and the planned certification pathways and courses for the initial CoE members to complete. This plan also identifies 2 CoE champions at Customer who evangelize the lakehouse vision and strategy ([CEA](#), [DSA](#)).

Enablement Journey Mapping

Together, collaborate on building a custom enablement plan via a Learning Needs Assessment. The enablement plan consists of role-based enablement paths that bring Student skills to where they need to be to effectively leverage Databricks for targeted business use cases ([CEA](#)).

Universal Lakehouse Fundamentals

Work together to identify the initial target use case and associated audience. This target audience will be encouraged to complete Lakehouse foundational courses in Databricks Academy as the first stage of their enablement plan execution. CEA will work with the Customer on training completion targets and success metrics. Feedback will be gathered and incorporated into the enablement pathways for an improved learning experience for subsequent Students ([CEA](#)).

Repeatable Persona-Based Cohorts

As the second stage of the enablement plan, we will collaborate to drive Students to persona-based blended learning cohorts, with the goal that those Students will earn a Databricks certification. During this process, feedback is collected and areas of improvement identified to accelerate the next cohort and/or to pursue additional training areas to assist in use case development and deployment ([CEA](#), [DSA](#)).

Digitized Enablement Journeys

Databricks and the Customer collaborate to provide a seamless experience for anyone in Customer's organization to register and learn about Data and AI through the Databricks Academy. Promote cross-team collaboration and skill development through Databricks Community private user group. The private user group also promotes sharing of ideas as well as posting announcements and driving awareness of upcoming events and/or classes ([CEA](#)).

Skills Validations via Badges & Certs

Work together to help the Customer drive certification across the organization in their respective personas. Educate and promote new use cases being deployed in the business across different community channels ([CEA](#), [DSA](#)).

Cultivate Data + AI Community

Databricks and the Customer will review the organization's culture to Data & AI and provide any further recommendations or learnings from industry best practices ([DSA](#)).

Executive Thought Leadership

Work together to drive Executive knowledge in Data and AI. Leverage industry events that Databricks hosts that showcase the organization's maturity and advanced use cases pushing the boundaries of Data and AI ([DSA](#)).

Additional definitions and terms

Key definitions

- **“Day”** means an 8-hour Service day, based on local business hours per Customer’s Service location
- **“Focus Area”** means each of four Lakehouse building blocks (*Data & AI Blueprint, Lakehouse Operations, Use Case Acceleration, Organization Enablement*)
- **“Focus Area Matrix”** means Databricks’ outline of CoE activities across the Focus Areas, used to jointly define Customer’s Service Implementation Plan
- **“Learning Needs Assessment”** (or **“LNA”**) means an analysis of existing skills (typically self-identified by a sample of team members) within a Customer relative to the skills needed to achieve business objectives that results in the creation of an enablement plan. An outcome of the LNA is an enablement plan.
- **“Month”** means a calendar month
- **“Student”** means a named individual who is registered to access and use the Service (and who may also be referred to as a user or learner)
- **“Week”** means Monday–Friday, excluding local holidays based on Customer’s Service location. (Customer’s default Service location, for purposes of defining primary contact hours/location, is selected at time of Order)
- **“Year”** means the 12 Month period beginning on the Start Date indicated in Customer’s Order Form

General

Resourcing

For flexibility, Customer may arrange to consume Weekly allocations of Service Days via 1 or more assigned FTEs (no fractional Days). Any unused Service Days or Training benefits in a given Week or Month do not carry over to future periods and are forfeited. Resourcing assignments require a minimum 4-weeks advance request (though Databricks will endeavor to reduce lead times where reasonably possible). Accordingly, Databricks recommends Customer coordinate with Databricks CoE planning at least a month before placing its CoE order. Databricks recommends that the periods of Professional Services and Tech/Strategy Services be the same, for maximum Customer benefit.

Delivery

CoE Service benefits are generally provided remotely. If onsite engagement is requested and agreed, Customer is responsible for reasonable travel and living expenses of Databricks personnel. Primary Databricks CoE contact coordinates other Databricks experts as needed; Databricks may update designated Databricks team member contact(s) from time to time.

Annual Benefits

Access and benefits described in this Service Description are provided on an annual basis. If a multi-year subscription is purchased, entitlements from a future annual period cannot be pulled forward into the current annual period, and unused entitlements cannot be rolled over for use in a subsequent period.

- **Professional Services** are available in weekly increments across each Customer’s annual CoE subscription period.
- **Tech/Strategy Services and Enablement** are available in weekly increments across each Customer’s annual CoE subscription period.
- **Skills@Scale** benefits are provided for the numbers of Students included in your subscription level, during each annual CoE subscription period.

Relating to Professional and Tech/Strategy Services

Service Implementation Plan (or "SIP")

The SIP is a jointly-planned selection of CoE activities from the Focus Area Matrix to support Customer's unique requirements. The SIP is developed at the beginning in the first 4-weeks of the annual CoE subscription period, based on desired Customer outcomes confirmed during the planning phase and drawing from the Lakehouse best practices outlined in the Focus Areas Matrix).

Assumptions

During the planning phase, with Databricks assistance, Customer will define its desired use of the Databricks Professional and Tech/Strategy resourcing included in Customer's subscription. Databricks does not make any guarantees that desired outcomes will be delivered within a fixed timeframe or fixed cost. Estimates may vary with new requirements or information. Successful engagement is contingent upon joint and active participation of Customer resources with required domain expertise. We assume our Advisory Service delivery will be based on 1-4 working Days per Week with Customer resources able to work on the implementation along-side the Databricks resources.

Databricks Responsibilities

Assign appropriately-qualified Resources to deliver Professional Services and Tech/Strategy Services.

Customer Responsibilities

Throughout engagement, make available appropriate Customer technical, business and domain experts with an understanding of Customer's business requirements for inclusion in design and implementation of Customer's CoE capabilities and processes. Such individuals need to be reasonably available to answer Databricks questions, and provide necessary context and implementation details, access to necessary environment, data and artifacts to successfully complete the project, as well as information (as relevant) on current data pipelines, data products, data sources, downstream systems, data consumers and their needs. Customer to identify point(s) of contact to run Customer's CoE program and its relevant services/capabilities.

Out of Scope

Items not included in the SIP (including deployment into production unless explicitly specified in the SIP), or work exceeding the allocation of hours and service types included in Customer's subscription. Databricks does not validate Customer business requirements, perform user acceptance testing, Customer change management, provisioning of any Workspaces, or configuration of/integration with non-Databricks solutions. Products that integrate with Databricks as source or sync but not part of the Databricks platform itself, e.g., streaming sources, orchestration tooling. Databricks does not guarantee the percentage of Customer employees who will take and/or pass a course.