This Service Description describes the scope of services (including associated Customer requirements) for the Advisory Service identified above (the "**Service**"), and applies to the Service under Customer's accepted Service Order.

Service Overview

Your data is your competitive advantage, and the best way to harness it is to build a foundation model that leverages your proprietary data with full control, transparent privacy, and cost control. The expertise of a ML Researcher helps you optimize your Hero Run results to align with your objectives.

Objective

Hero Runs (also sometimes referred to as Hero Reservations) are for foundation model training or very large fine tuning on a Customer–dedicated cluster (with other minimum sizing requirements as described in Customer's Services Order when reserving the Hero Run).

ML Researchers are optionally involved in three critical phases of a Hero Run:

● First, before a customer embarks on their Hero Run, to review the configurations and make recommendations, with the goal of improving model performance or avoiding issues during the Hero Run itself.
● Second, to monitor the Hero Run to ensure the cluster is operating normally and loss metrics, etc are proceeding directionally on course.
● Lastly, at the end of the Hero Run, to help interpret the final output of the Hero Run.
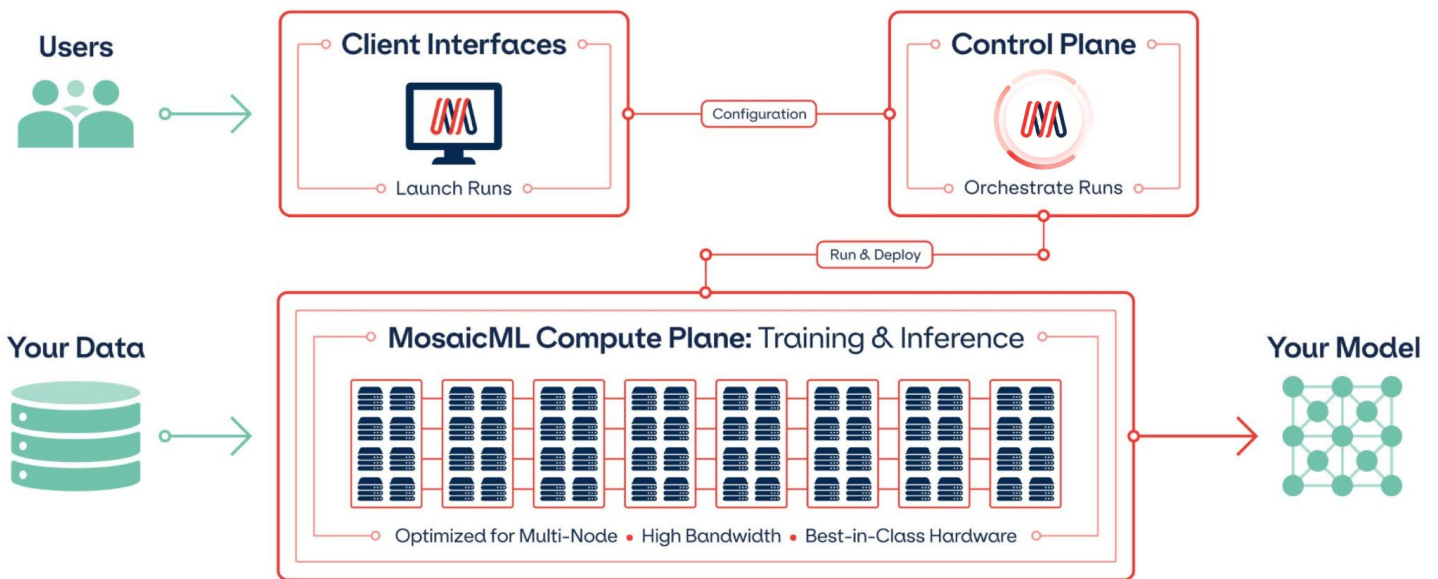
Description of Services

Databricks will provide Services from the Technical Focus Areas and Representative Activities described below in advising Customer on Hero Runs. Specific activities performed will vary, depending on Customer-specific objectives.

Technical Focus Area	Maximum Duration*	Representative Activities
Advising Pre–Hero Run *Review Hero Run configurations*	Up to 8 hours	● Validate model architecture, data mix/selection, training duration, and optimization settings ● Review configurations and make recommendations
Monitoring *Monitor health and progression of Hero Run*	Up to 2 hours per week of Hero Run	● Co-Captain to monitor the Hero Run at least twice a week to check that all reserved GPUs in the Hero Run Cluster are operating normally, monitor that the Hero Run (including loss metrics) is proceeding directionally on course, and reach out to Customer with abnormal findings (if any).
Review *Interpret final output*	Up to 2 hours	● Review and explain the results of the Hero Run ● Provide recommendations on how to use the results of the Hero Run

Customers who schedule a Hero Run extending beyond 30-days, may purchase an additional 2-hour increment of ML Researcher time per each additional 30-day period to fulfill the Representative Activities listed above.

The below visual illustrates a high-level user workflow interacting with MosaicML to train your model.

Prerequisites

Throughout the engagement, Customer will assure that the following requirements are met, to enable the Services:

- Customer must use the full MosaicML stack, including Mosaic Command Line Interface (mcli) as the orchestration tool, Composer, Streaming Dataset, and Mosaic modeling code (if applicable).

- Customer to appoint a Captain for the Hero Run, responsible for actively monitoring and managing the Hero Run, and be the single point of contact for the Databricks team.

- Customer to ensure at least one engineering team member and one ML team is on-call throughout the duration of the Hero Run in case any issues arise.

Out of Scope

- Setting up an inference pipeline
- On-call support beyond any Support Services explicitly included and provided under the Service Order
- Documentation of the solution
- Training of end users
- Work exceeding the allocation of Days and Services included in this engagement (see **Resources and Schedule**)
- This is a time and materials engagement - Databricks does not make any guarantees on the success of the stated goals of the Hero Run

Resources and Schedule

ML Researcher time is sold and consumed in **2 hour increments**, applied against the Representative Activities in the Description of Services table above. (Additional assistance of a Co-Captain is also provided, on an as-available, reasonable basis.)

Resourcing assignments require execution of a Services Order before we can schedule resources. (While Databricks makes reasonable efforts to accommodate scheduling requests, personnel availability is subject to Databricks resourcing and discretion).

Note that the initial phase of Services, the ***Advising Pre-Hero Run*** phase, must take place before the Target Start date, so timing of Services Order placement should take into account this runway requirement before setting the Target Start date.

Additional Definitions and Terms

- "**Agreement**" means your agreement with Databricks, providing general terms for our Services.
- "**Co-Captain**" means a Databricks-assigned individual responsible for monitoring the Hero Run, as described in the Representative Activities table above.
- "**Captain**" means a Customer-assigned individual responsible for the end result of the Hero Run.
- "**2 hour increments**" means two contiguous working hours during local business days, excluding holidays.
- "**Services Order**" may be any of these mutually-accepted formats placed under your Agreement: an Order, Success Credit redemption request, written statement of work, or similar document.
- "**we**", "**us**" or "**our**" or "**Databricks**" means Databricks, Inc. or its Affiliates.
- "**you**" or "**your**" means the Customer organization that placed the Services Order.