

AGENTIC MLOPS ACCELERATOR | POWERED BY PHDATA FORGE™

Deploy ML Models Fast. Reduce Operational Costs. With **Snowflake ML** and **Cortex Code**.

Full-lifecycle ML model deployment. Five weeks.
Built for production — with observability, reliability, and maintainability from the start.

DELIVERABLE

Full lifecycle ML model deployment

TIMELINE

5 weeks

FIXED PRICE

\$60K

Getting models into production is one of the biggest bottlenecks for data science teams. After go-live, managing those models and automating monitoring for drift detection with performance-based triggers for retraining adds further costs and complexity that are often prohibitive. Teams invest months building and validating models, only to see them stall before they ever reach end users, or quickly abandoned as performance drifts. Leveraging Cortex Code, phData Forge™ Skills, and deep Snowflake ML expertise, phData delivers a production-ready model deployment with an AI-native, skills-based workflow.

The result is speed, self-sufficiency, and zero operational debt.

01 DELIVERABLES

WHAT WE DELIVER

We replace MLOps friction with an AI-native toolkit for deploying and managing agents. Powered by Cortex Code and phData Forge™.

➤ **Full-lifecycle Model Deployment**

Including inference, training, and retraining.

➤ **Observability Dashboards**

Monitoring for model performance and data drift.

➤ **Reusable Cortex Code Skills**

Agentic model deployment and management skills that reduce operational costs for your team.

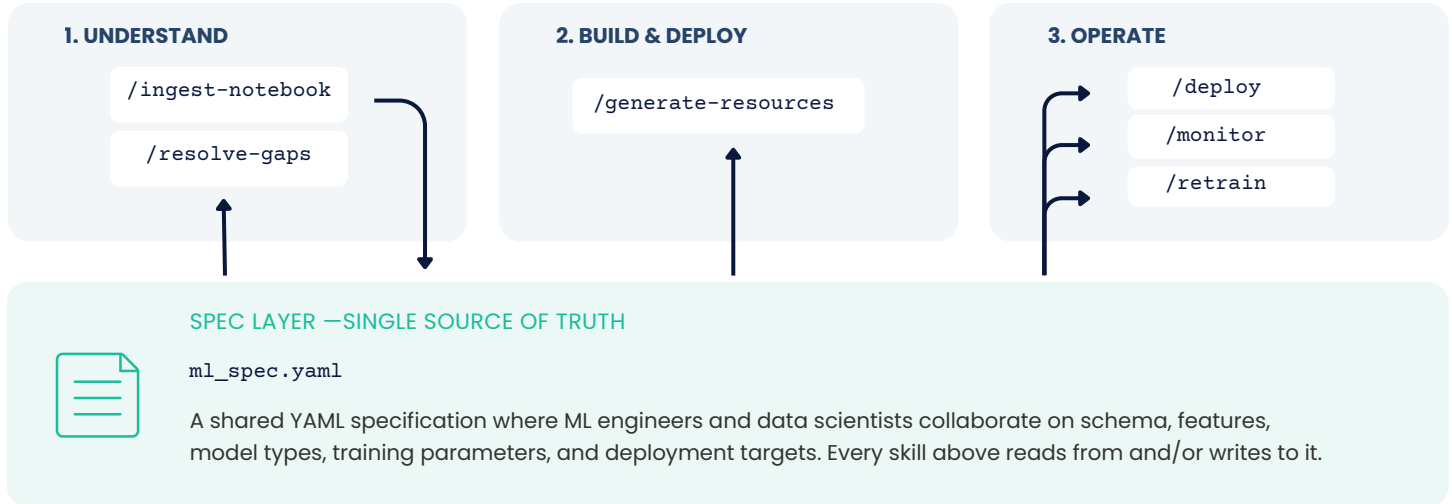
➤ **Enablement and Best Practices**

Workshops and office hours to train your team to manage with the agentic skills we deploy.

02 METHODOLOGY

HOW WE DEPLOY AND OPERATE

A reusable, AI-native deployment pattern. Three workflow phases, one spec-driven source of truth.



03 BENEFITS

PRODUCTION MLOPS, WITH OPERATIONAL CONFIDENCE

- ✓ Cut time to production from months to weeks with AI-accelerated deployment on Snowflake ML
- ✓ Free your data team to focus on building the next model, not maintaining deployment pipelines
- ✓ Reduce dependency on scarce ML engineering resources with a fully managed, expert-led deployment
- ✓ Get a repeatable deployment blueprint your team can use long after the engagement ends
- ✓ Prove AI value faster with a working production model in hand within a month

RESULTS

Faster time-to-value

Higher confidence in model outputs

Lower rework & regression

Repeatable path to scale

Self-sufficiency & institutional knowledge

Ready to accelerate ML deployments on Snowflake?

FULL-LIFECYCLE MLOPS DEPLOYMENT AND ENABLEMENT. FIVE WEEKS. BUILT FOR PRODUCTION.

[Talk to a Snowflake Expert](#)