

Manage thousands of models for drift in real-time.

See ROI in no time.

When model behavior changes, you need to quickly figure out what is going on to fix the problem.

But identifying root cause is hard - and it gets harder as more models are deployed to production.

Now you can break-through alert fatigue, and get to root cause fast. Even at scale.

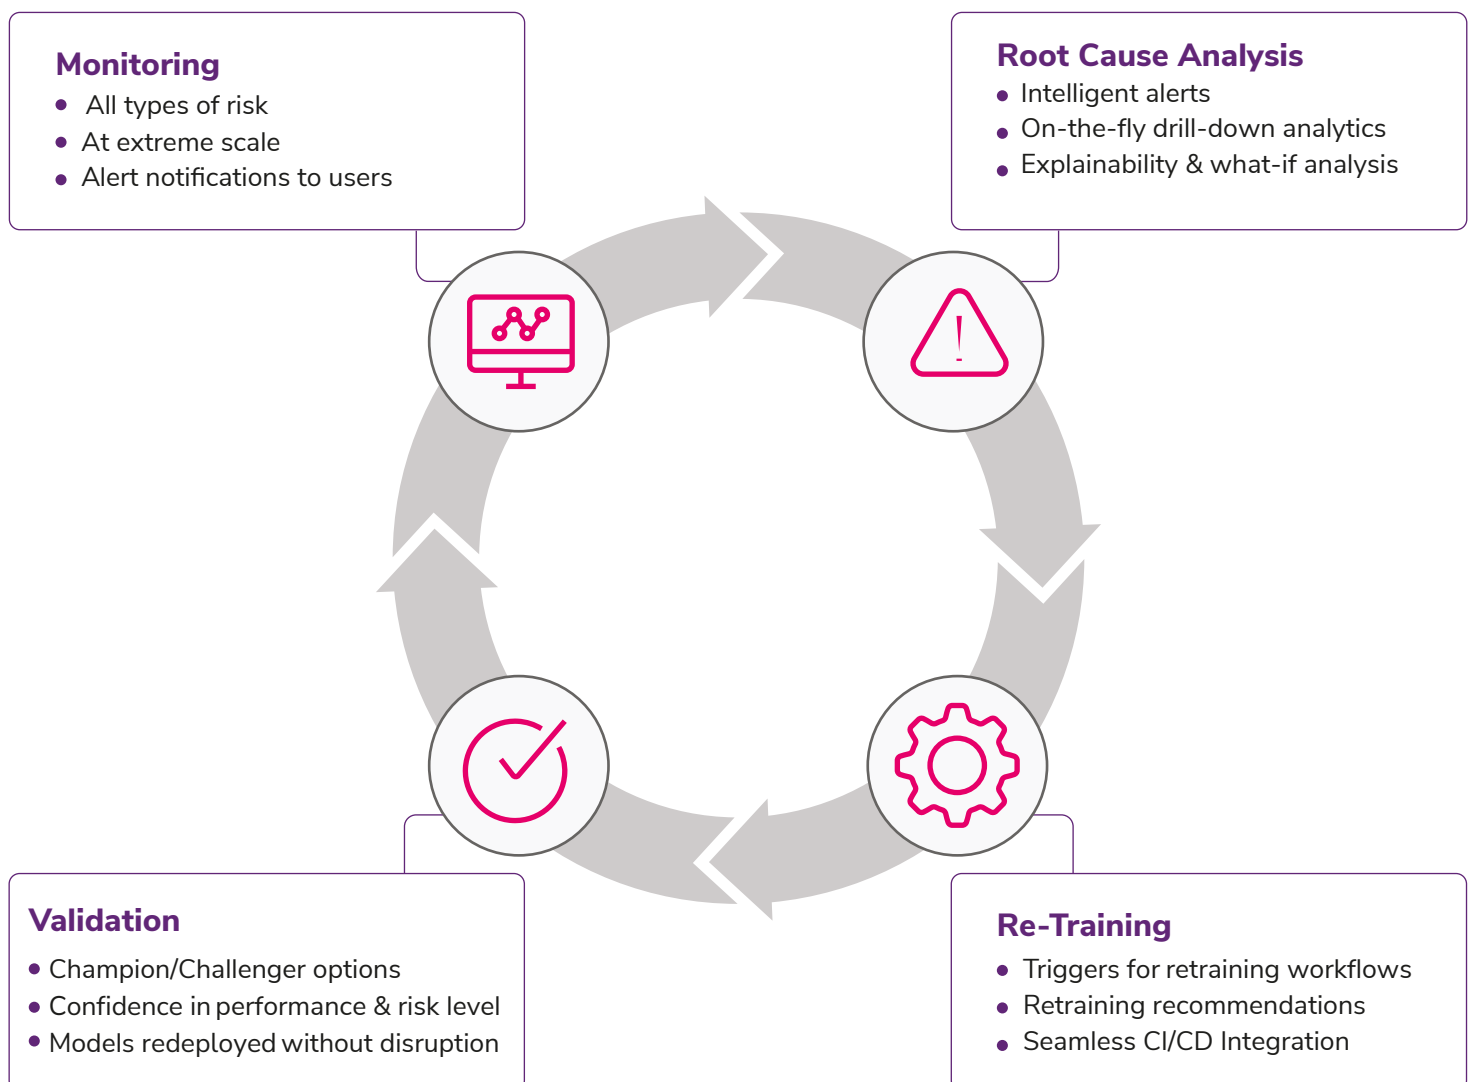
We help you keep models trustworthy and running in production to more quickly realize business value.

Continuous Operations

Once models are deployed to production, they begin to drift. But **why** their behavior has changed is not always easy to pinpoint. And as more models are deployed and monitored in production, it becomes difficult to know which alerts are critical and need further investigation.

Often, organizations choose to retrain at regular intervals to address drift – but this can be expensive. In the past, AI operations and the lifecycle of ML models were viewed as something that could be pieced together across teams. With the introduction of new AI tools, it is clear that a strong operational approach to how these systems are monitored and managed is foundational to any enterprise AI strategy.

We help teams maintain a continuous operations loop by identifying root cause for changes in model behavior and providing the insight and data needed to address problems.



Open, Extensible, Flexible

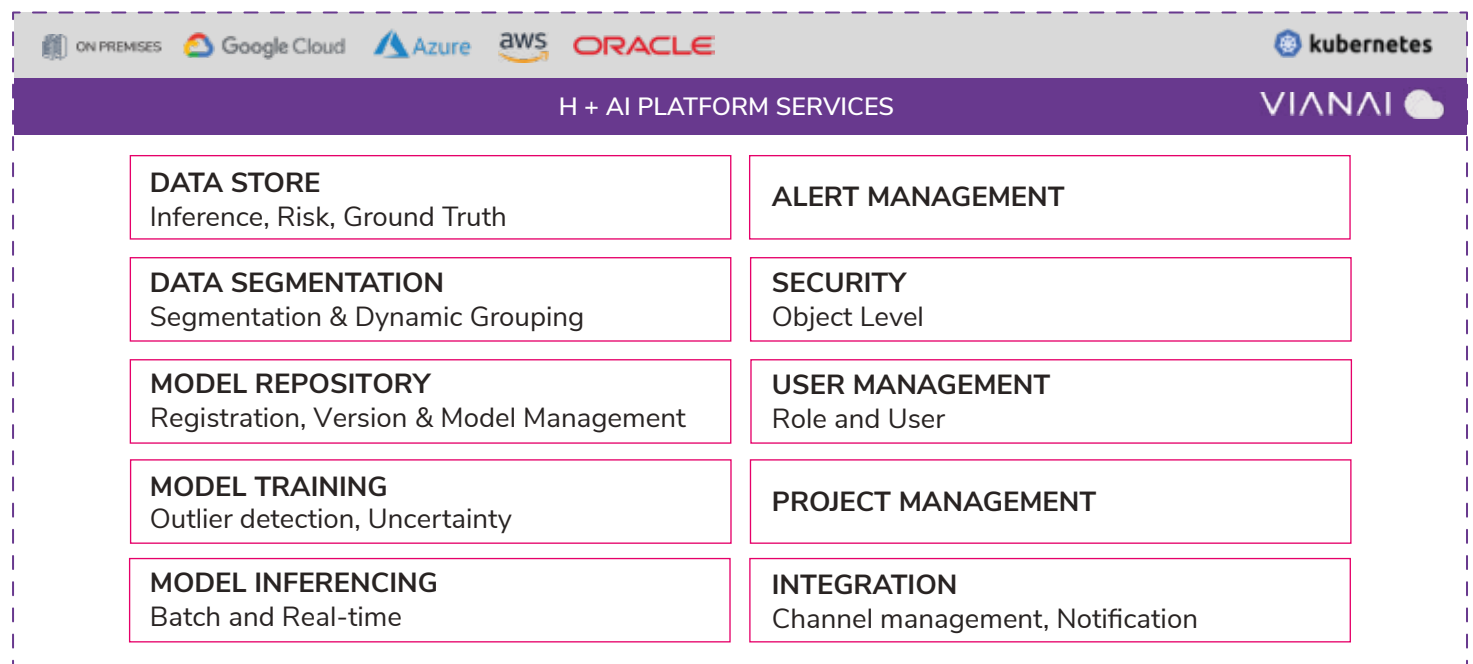
We help data scientists, machine learning engineers, and IT teams work better together.

In each organization, there are different people across the ML workflow with varying levels of data science expertise, and limited access to the data they need to make good decisions.

Our monitoring application easily integrates into your ecosystem. Supported by our human-centric (H+AI) service based platform, users can access our continuous operations capabilities through an intuitive, easy to use UI, or choose to access capabilities directly through APIs to integrate monitoring seamlessly into the existing ML workflow.

We monitor models that were built anywhere and run anywhere.

Monitor at scale Identify root cause in real-time Re-train models automatically Validate models continuously



Data Scientist

- Build in notebooks
- Automate package & deploy
- Gain deep insights to model behavior



ML Engineer

- Validate data quality
- Train & compare multiple model versions
- Identify data distributions that impact performance



IT Ops / ML Ops

- Import and manage model versions, artifacts, and transformations
- Ensure reproducibility
- Automate deployment to any endpoint



CXO, Business Users

- View model performance
- Gain insight to key factors that influence behavior
- Measure & track business value

Keep More Models Running in Production

Operationalize models at scale

We go beyond traditional lightweight monitoring tools and enable root cause analysis, retraining, and model validation in a continuous loop across large, complex, feature-rich models.

Some of the ways we do this:

- Wizard-driven policy creation for each model - distance or window based
- Drill-down insights across hundreds of features to identify root cause
- Highly flexible and granular control for setting warnings & critical alerts
- Automated triggers to kick off retraining workflows
- Challenger/Champion models to validate performance - before redeployment
- Automated model replacement without business disruption
- Intuitive UI so that any stakeholder can perform tasks across the ML workflow

We work some of the largest manufacturing companies, banks, and other enterprise organizations and have tested our platform's performance at billions of events per day.

HUNDREDS

features

TENS OF THOUSANDS

inferences/sec

BILLIONS

rows

VIAN Comprehensive Risk Assessment

We are building the capabilities below to provide the most comprehensive risk monitoring capabilities to augment the ability of humans to better understand and manage risk, meet internal and regulatory compliance requirements, and deliver trustworthy models.

- Drift - Feature, Prediction, and Concept
- Explainability - Global and Local
- Uncertainty & Robustness
- Outliers
- Data Quality
- Bias & Fairness

 Root-Cause Analysis  Alerts & Recommendations for Re-training  Validation