

## Power your organisation with trusted intelligence and real-time data-driven decision making.

Bedrock is a cloud-based enterprise AI platform. It is a platform-as-a-service that provides data-driven enterprises a path to productionising responsible machine learning (ML) products. Bedrock helps you achieve a faster time to market for real-time, massive scale AI engines. In-built governance enables transparency and accountability of AI in production; the foundation for responsible AI deployments within enterprises.

Bedrock provides three layers of functionality around your ML workflows:

1. Metrics and artefact logging: Track the AI model performance and complete ML lifecycle.
2. Governance: Fine-grained permissions, digital audit trail, model monitoring, alerts and security-by-design.
3. Infrastructure Orchestration: Automate your cloud infrastructure for ML workloads.

### Who is Bedrock for?

- Data scientists
- Machine learning engineers
- Data science leaders
- Anyone who is responsible for making sure that ML models are achieving the business objectives in the intended manner

## PRODUCT BENEFITS

### Faster time to market for real-time, massive scale AI engines

---

#### Manage end-to-end ML lifecycle

- Standardized workflow to support end-to-end ML lifecycle from experimentation to deployment.
- Create ML Projects to allow different user personas to collaborate on solving specific business problems.
- Easily set up different environments to separate runtime for different use cases e.g. staging and production environment.
- Managed notebooks for exploratory data analysis and experimentation.
- Set up production grade ML pipelines to minimize ML tech debt.

#### Automate ML workflows

- Compose multi-step training pipelines to modularise pipelines using Argo open-source framework.
- Track resource utilisation and view persistent training logs for better debugging.
- Automatically scale up/down the resources based on the demand.

- Compare hyperparameters and metrics for different model versions.
- In-built scheduler to automate ML pipelines to keep models fresh and performant.
- Flexibility to support all the major open-source ML frameworks including Tensorflow, PyTorch, scikit-learn etc.
- Harness GPUs for training, and managed Spark for feature engineering and distributed training of ML models.

#### Production deployments of new models in minutes

- Supports batch scoring and API deployments.
- Quickly deploy containerised gRPC and REST endpoints with auto-scaling enabled.
- Weight-based traffic splitting for new model version deployments to reduce risk in production deployments.
- Out-of-the-box monitoring capability for model endpoint metrics such as latency and throughput.

### AI governance by design

---

#### Digital audit trail

- Leverage the platform to get complete visibility across ML lifecycle and track the model provenance.
- Track and log distinct model versions.
- Keep track of model history deployed to endpoints.
- Visualise the complete ML lifecycle from training code to model artefact logging to endpoint deployment in a single-pane-of-glass view.

#### Ongoing monitoring of models and closing the algorithmic feedback loop

- Automatically detect feature and inference drift in production.
- Log every inference made by the models in production.
- Supports logging of ground truth data to continuously monitor the model performance.

## Securely collaborate and retain control over your data

- Manage users in organisations and teams.
- Retain control over your data and code.
- Review, compare, validate and discuss models before they are approved to be deployed to production.
- Role based permissions control (Admin, Train, Review, Deploy) so that Data Science, business stakeholders and DevOps can collaborate seamlessly.
- All data stored by Bedrock is encrypted in transit and at rest, complying with the Centre for Internet Security (CIS) benchmarks.
- Secrets management facility to store sensitive information, backed by Hashicorp Vault.

## Responsible AI

- Out-of-the-box visualisations of explainability and fairness metrics.
- Use Bedrock to easily compute and visualise the explainability and fairness metrics for your ML models using the latest interpretability methods such as SHAP.
- Visualise explainability at the individual instance and model level.
- Select fairness metrics and thresholds, and be alerted when unintended bias is discovered.
- Apply mitigations to models to de-bias them.

## Integrate easily, and provision flexibly

### Managed infrastructure on multi-cloud deployments

- Leverage Infrastructure as code (IaC) and Terraform to automate infrastructure provisioning.
- We provision underlying infrastructure in your cluster harnessing modern open-source technology: Consul (service discovery, distributed key-value store), Prometheus (monitoring), Fluentd and Traefik load balancing.
- Self-service infrastructure in the Bedrock managed Kubernetes cluster for all the data science and ML engineering needs including managed notebooks, training and model deployment.
- Cloud agnostic, works on all major cloud providers (AWS, Azure, GCP).
- SLAs of 99.5% availability of AI application endpoints.

- Compliant with CIS benchmarks for AWS and GCP, assessed by third-party cybersecurity experts.

### Integrations with popular and modern ML and DevOps tools

- Unify data science processes across your organisation by leveraging flexible product design to integrate with popular data science tools and frameworks.
- Allow your data scientists to use all popular open source ML frameworks, including Tensorflow, Pytorch, Scikit-learn, Keras, MXnet, XGBoost and any other ML framework of your choice.
- Easily integrate with popular CI/CD systems such as Jenkins, Teamcity and TravisCI using Bedrock APIs.

## ARCHITECTURE

### “Glass-box” deployment

We orchestrate and securely manage ML infrastructure in an isolated virtual private cloud (VPC) on your cloud environment. You have full visibility but we are responsible for management.

