

Running AI Workloads in a Multi-Cloud/ Hybrid Cloud Environment

As enterprises embrace multi-cloud and hybrid cloud strategies, running AI workloads across different cloud environments can pose challenges. However, with Run:ai's powerful platform, organizations can seamlessly manage and optimize AI workloads in a multi-cloud/hybrid cloud environment, ensuring efficient resource utilization and delivering a consistent user experience.



Cross-Cloud Visibility & Control

Empower platform teams with comprehensive cross-cloud visibility, providing insights into resource allocation across different cloud platforms. This centralized view enables organizations to make informed decisions and set policies to control access to resources, ensuring compliance and efficient utilization of AI infrastructure.



Consistent UX Regardless of Location

Ensure that data scientists and developers can enjoy a consistent user experience, regardless of their physical location or the underlying cloud platform. The platform delivers a unified user interface, abstracting the complexity of managing multiple clouds. This eliminates the need to learn and navigate different platforms, enhancing productivity and collaboration.



Resource Management and Optimization

By pooling all compute regardless of location and automating resource allocation, organizations can achieve optimal resource utilization and effectively manage their AI workloads. Easily move workloads to the cloud when on-premises resources are fully utilized, run training and inference at different cloud providers or run workloads on different instances or clouds based on their priority.

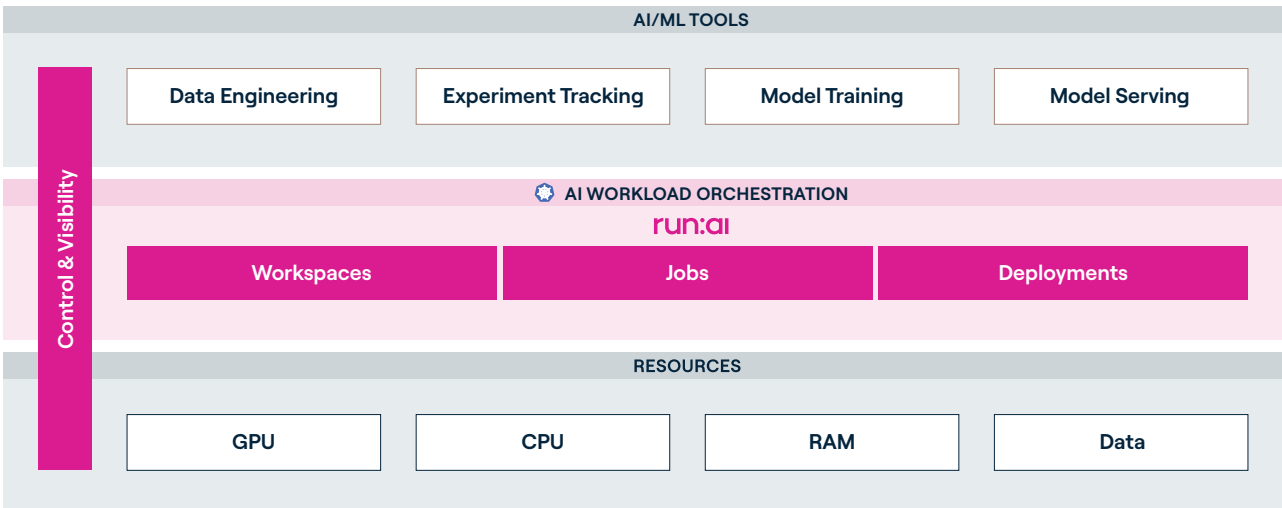
ANY CLOUD

Run:ai offers support for all major public cloud vendors and their Managed Kubernetes offering.



Platform Overview

The Run:ai Atlas platform sits in between the infrastructure and the AI workloads that require access to these valuable resources. Platform teams gain centralized control and visibility across all AI infrastructure, whether on-premises or cloud. AI/ML teams get streamlined and self-service access to all the compute they need, when they need it, using the tools they prefer.



Feature Highlights



Multi-Cluster Dashboard

Provides a holistic, aggregated view across clusters and sites, including information about Cluster and Node utilization, available resources, and allocated resources.



Workspaces

Workspaces enable data scientists to self-provision the model development tools (like Jupyter Notebook, W&B, MLflow) together with the compute and data they need, in a simplified, streamlined and secured manner.



Cross-Cloud Control

Define Policies that determine which users can access data and the different resources across clouds, specify quotas per cloud or cluster, or support the use of heterogeneous clusters.

Customers Accelerating AI with Run:ai Atlas

SONY

BNY MELLON

ZEBRA

XIAOMI