

The AI Scaling Challenge

The Solution

Advantages

Key Capabilities & Benefits

Every Organization's Challenge

I built a model prototype on my laptop and on a small server. But when I started to train and scale it on large data, that's when we ran into a host of problems.

AI Developer
Leading SaaS Company

The Anyscale Platform

The **Anyscale Platform™** is a **fully managed scalable compute platform** built on Ray that enables any organization and any AI developer to effortlessly build, tune, train and scale machine learning (ML) and Python workloads.

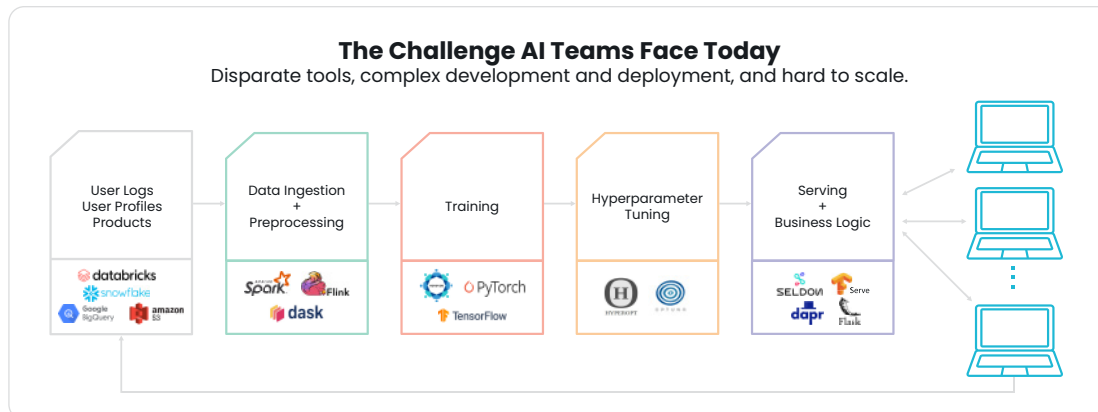
Ray is used by thousands of organizations globally to increase developer velocity and scale AI.

The Anyscale Platform™ offers several key advantages over Ray open source and provides a seamless user experience for developers and AI teams to speed development, improve developer productivity and productionize AI/ML workloads at scale including large data sets. The result is faster time-to-market and faster iterations across the entire AI lifecycle.

The AI Scaling Challenge

Over the past decade, the compute and data requirements of AI/ML workloads has grown dramatically, requiring organizations to build complex infrastructure and set up large, compute-intensive environments to meet their scaling challenges.

Across the AI lifecycle, from development to training to real-time serving, the complexity of scaling AI, including stitching together many AI/ML tools and frameworks, is resulting in project failures, cost overruns and delayed time-to-market. This fundamental challenge has held organizations back from realizing the full potential of AI.



The Solution

Ray is an open-source unified compute framework that simplifies building and scaling AI/ML Applications and Python workloads. The Anyscale Platform™, built on Ray, addresses every aspect of AI scaling, eases development and speeds time-to-market for any AI project and initiative by offering a seamless development and scaling experience, built-in management and observability capabilities, cluster management and comprehensive enterprise-ready features.

The Anyscale Platform™, used across industries today, is built and managed by the Anyscale team, the original creators of Ray.

The Anyscale Platform

Advantages

Anyscale and Ray simplify building and scaling AI workloads for every developer by providing a laptop-like development experience on large clusters. This dramatically accelerates the development of complex applications by enabling developers to easily iterate between development, debugging and testing, and instantly scale and productionize their applications on large data sets – unstructured data included – without any code change or need for complex infrastructure setup. The Anyscale Platform™ provides the first and only scalable unified compute infrastructure for the entire AI lifecycle from data ingest and data preprocessing to training, tuning, model serving and more, all while integrating seamlessly with the rest of the ML ecosystem.

Organizations around the globe use Anyscale and Ray to solve their most pressing AI scaling challenges for both new and existing workloads. In addition to seamless scaling, the Anyscale Platform enables AI teams to bring diverse AI applications into production faster and benefit from vastly increased developer productivity. AI teams also benefit from the ability to use the AI/ML tools and frameworks that they already know and love and avoid refactoring their existing workloads when moving them to Anyscale and Ray.

The AI Scaling Challenge

The Solution

Advantages

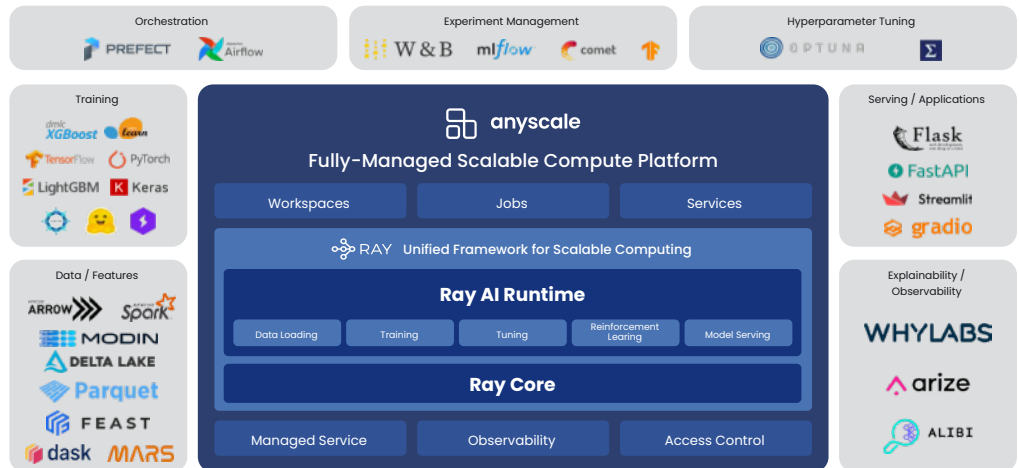
Key Capabilities & Benefits



In the same time that it took to run our original workload (7 days), **we were able to effortlessly migrate over our Python code to Anyscale**, fine tune our job for scaling, and **move to production effortlessly.**

Jake Carter
Director Data, ML, and Technology
Biolexis Therapeutics

Unified Framework for Scaling AI/ML and Python Workloads



- 1** Production-grade, unified scalable compute platform
- 2** Accelerate development and experimentation of ML and Python workloads
- 3** Seamlessly move between development and production

The AI Scaling Challenge

The Solution

Advantages

Key Capabilities & Benefits

The Anyscale Platform

Key Capabilities & Benefits

1 Production-grade, unified scalable compute platform

- 1. Fully managed service** Anyscale operates clusters of machines on demand so AI teams don't have to operate the cluster. ML practitioners get access to an interactive, scalable compute environment. It accelerates application development irrespective of the scale of the workloads.
- 2. Bring your own cloud** Anyscale is built from the ground up with customer data security in mind. It runs in an organization's infrastructure and cloud account, while still providing an exceptional managed experience.
- 3. Optimize compute costs** Anyscale's autoscaling and auto-suspend features, and spot instance support allow teams to reduce the compute costs of running workloads. Anyscale also allows teams to leverage their existing agreements with public cloud providers (i.e. AWS Reserved Instances or Saving plans).
- 4. Governance and compliance** Anyscale provides user access controls for projects, workspaces and clusters as well as cost tracking mechanism. Anyscale has SOC 2 Type 1 attestation.
- 5. Expert support** With Anyscale, teams get direct access to the Ray creators. In contrast to OSS, where support is provided on a best-effort basis by the community and occasionally by committers, Anyscale offers dedicated support provided by Ray and Anyscale engineers that help teams develop and move applications to production much faster.

The AI Scaling Challenge

The Solution

Advantages

Key Capabilities & Benefits



Anyscale Workspaces allows me to go from development, to experimenting at scale, all the way to production all within the same environment.

Workspace **reduces context switching for us by 50%**, and integrates with other tools we use.

*Data Scientist
Manufacturing
Conglomerate*

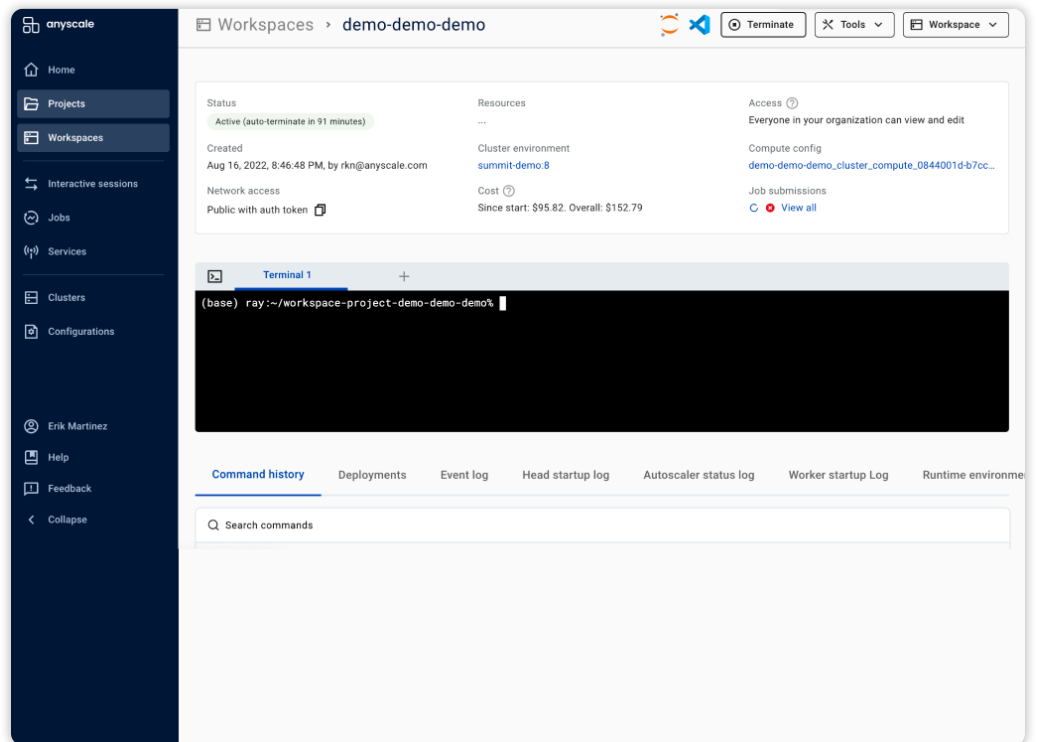
The Anyscale Platform

Key Capabilities & Benefits

2

Accelerate development and experimentation of ML and Python workloads

- Anyscale Workspaces** Anyscale Workspace provides an integrated IDE experience that allows teams to edit and run code, install dependencies, monitor jobs and resources across a scalable cluster just like you would on your laptop.
- Use the tools you know and already love** Anyscale provides integration and instant setup for popular tools such as VSCode and Jupyter notebook, with Github, Weights & Biases, and more.
- Collaborate** Share or clone experiments with a click of a button. Different users can access a workspace with all the same configuration and environments and be productive instantaneously.



The AI Scaling Challenge

The Solution

Advantages

Key Capabilities & Benefits

Want to get started with the Anyscale Platform?

Contact us at info@anyscale.com

Request a Trial at www.anyscale.com

The Anyscale Platform

Key Capabilities & Benefits

3

Seamlessly move between development and production

- 1. A unified environment for development and production** Anyscale focuses the experience on a unified environment - run, debug, and test your code at scale on the same cluster configuration with the same software dependencies for both development and production.
- 2. Flexible and extensive dependency management** Anyscale provides different options to manage your dependencies across your cluster. You can use existing base images, bring your own Docker, or use Ray's runtime environments for faster iteration.
- 3. Jobs and services** Jobs and services API and SDK provide an easy interface to operationalize your workloads and integrate with your existing deployment tools. Anyscale Jobs support cron jobs, ephemeral cluster creation and retry capabilities, while Anyscale Services provides replica management, no downtime upgrades and high availability.
- 4. Managed logs, monitoring, and observability** Anyscale provides a production-grade monitoring and observability stack with a managed Grafana and Ray dashboard. Additionally, Anyscale provides production monitoring and notifications for added trust that pipelines are running well.

