

Open Source Ray and the Anyscale Platform™

A Side-by-side Comparison

Open Source Ray and the Anyscale Platform™ solve the most fundamental problem facing AI teams today: scaling AI/ML and Python applications and deploying them in production. Ray is a unified compute framework to scale AI/ML and Python applications and is the fastest growing open-source distributed framework. The Anyscale Platform is the first enterprise-ready, fully-managed Ray platform. It provides advanced capabilities beyond Ray OSS to accelerate the development, deployment and management of Ray applications in production.

The AI Scaling Challenge

The challenge of scaling AI is rooted in the increasing compute demands of machine learning (ML). Since 2010, ML training needs have surged 10x every 18 months. In contrast, the compute power of AI accelerators, like GPUs and TPUs, has less than doubled in the same timeframe. This leaves developers no choice but to build complex distributed applications to scale these workloads. Ray is emerging as a go-to framework to support such workloads. For instance, OpenAI's popular ChatGPT model is trained and scaled on Ray. Yet, beyond the challenges of building distributed applications, AI teams must grapple with the time and resources required to set up, configure, and maintain both development and production infrastructure. Anyscale offers a solution to these problems.

People often ask: **“Why use Ray?”**, and **“Why use the Anyscale Platform™ instead of just using Ray OSS?”** The next eight points expand on these questions and answer them.

1. What is Ray?

Ray is a unified compute framework for scaling AI/ML and Python workloads. Ray enables ML engineers and developers to scale their applications from their laptops to large clusters with no code changes. Ray is Python-native and integrates with the entire machine learning and data ecosystems, including deep learning frameworks (e.g., TensorFlow, PyTorch), traditional machine learning libraries (e.g., scikit-learn), workflow orchestration platforms (e.g., Airflow, Perfect), experimental tracking (e.g., MLflow, Weights & Biases), features stores (e.g., Databricks, Feast, Hopsworks), model registries (e.g., Databricks, SageMaker), and various data formats and systems (e.g., Delta Lake, Snowflake). Ray has a flourishing community that includes leading technology companies like Uber, Spotify, Shopify, and Instacart.

2. What's unique about Ray and what makes it the 'first of its kind' in AI?

There are many popular distributed systems like Apache Spark, Apache Kafka and Apache Flink. These are Java-based systems, and are mostly focused on scaling data processing. In contrast, Ray is a unified, Python-native distributed computing framework that works as effectively on a local machine as it does on hundreds of machines. Ray Core (a part of the overall Ray project) provides primitives to easily parallelize existing Python applications. Ray AI Runtime (AIR) provides machine learning specific tooling for tasks like hyperparameter tuning, reinforcement learning, and production model serving.

3. What Is Anyscale?

Anyscale is the company founded and led by the creators of Ray, Philipp Moritz, Robert Nishihara, and Ion Stoica. Ion Stoica, Executive Chairman of Anyscale is also the co-founder and Executive Chairman of Databricks. Anyscale commercializes the Anyscale Platform, a fully-managed scalable compute platform built on Ray that enables any organization and any AI developer to effortlessly build, tune, train and scale AI/ML and Python workloads.

4. Why do customers use the Anyscale Platform over Ray?

The Anyscale Platform is the best place to create, manage and run Ray applications. It enables users to leverage Ray to solve AI scaling challenges while removing the need to perform complicated infrastructure setup, administration, and cluster management. The Anyscale Platform offers Anyscale Workspaces, an integrated development environment that enables ease of development, fast iteration, collaboration, and seamless deployment of models to production. Second, it offers fully-managed Ray clusters that eliminate the need for organizations to hire a team of several engineers to build and manage a Ray computing infrastructure. Removing this work allows developers to focus on building amazing applications for their users instead of undifferentiated, complex tools to manage Ray clusters. Finally, it provides built-in monitoring and observability capabilities, as well as comprehensive enterprise-ready features, to easily and securely deploy and manage these workloads in production.

5. Is the Anyscale Platform easier to use than Ray?

While Ray, as a framework, makes it easy to scale ML/AI and Python applications, the Anyscale Platform provides all the tooling around Ray applications critical for users to succeed putting those applications into production. The Anyscale Platform makes it easier for developers to collaborate to speed up development and deployment of Ray applications. For development, Anyscale provides Workspaces, a fully-managed development experience that integrates with popular development tools like Visual Studio Code and Jupyter Notebooks while providing simple access to logs for debuggability. For production, Anyscale provides native jobs and services, supporting batch and online serving applications respectively. Jobs and services include features like cron scheduling, automatic failure recovery, retries, and more. Critically, the fully managed Anyscale Platform frees

developers from setting up and managing the Ray clusters for development or production. This allows developers to accelerate experimentation and reduce time to market, focusing on their business problems instead of infrastructure

6. Do Ray and the Anyscale Platform require any special skills from developers?

There are no special skills outside of Python knowledge to use Ray or Anyscale. Ray enables developers to scale their existing Python and ML workloads by adding just a few lines of code or by just using the Ray AIR libraries. The Anyscale Platform enables users to get the most out of Ray by providing fully managed clusters, a hosted development experience called Workspaces and simple production tooling for batch jobs and online applications.

7. Is there any difference between Ray OSS and Ray as provided by the Anyscale Platform?

The Anyscale Platform provides access to Ray OSS, so existing Ray applications run without any changes. In addition, the Anyscale Platform provides unique capabilities for existing Ray applications, such as cluster management, faster cluster startup times, and powerful development and observability tools. As a result, the Anyscale Platform dramatically accelerates development, and reduces time to market by enabling developers to move their Ray workloads to production with no code changes.

8. Do Ray and the Anyscale Platform integrate with existing machine learning and data tools, libraries, and frameworks? Is there any difference between how it works with Ray vs. the Anyscale Platform?

Ray integrates seamlessly with the existing data and ML ecosystems. Being Python-native, Ray enables developers to use their favorite Python libraries and tools, including deep learning tools (TensorFlow, PyTorch), machine learning and statistical modeling libraries (scikit-learn, statsmodel), NLP tools (spaCy, Hugging Face, AllenNLP), and model tuning frameworks (Hyperopt, Tune). In addition, developers can read and write data from a variety of data sources, including Databricks, Snowflake, Kafka, and more. Finally, developers can use existing workflow frameworks such as Airflow and Prefect to orchestrate Ray jobs, and ML tracking and observability tools such as MLFlow, Weight & Biases or Arize. Anyscale makes these integrations even easier by providing one-click integrations with integration authentication and bi-directional integration between popular tools.

Key features provided by the Anyscale Platform™ in comparison to building your own on Ray open source.

CATEGORY	FEATURE	OPEN SOURCE RAY	ANYSCALE PLATFORM™	DESCRIPTION
Development	Development experience	Basic	Workspaces with hosted VSCode and Jupyter Notebooks	Managed VSCode and Jupyter Notebooks with Git integration. Workspaces provide a similar development experience on a large cluster as on a laptop.
Development	Sharing artifacts	None	Easily share application's code and environment.	Provide the ability to duplicate the application's code and environment from Workspaces, Jobs, and Services for development and debugging. Share and collaborate with other users.
Development	Dependency management	None	Dependency management	Anyscale manages dependencies for the latest Python release with CUDA. Anyscale also provides various dependency management capabilities for users to integrate with their existing infrastructure simply and effectively.
Development	Cluster Startup Time	Basic	Faster cluster startup time.	Anyscale optimizes cluster startup time, allowing for a quicker development cycle.
Production	Job Submission	Basic	Job submission	Anyscale provides major enhancements over Ray jobs, including cron support, retries, and persistent outputs along with email notifications for retries & failures.
Production	Services	Basic	Services	Anyscale provides high availability service deployments with autoscaling and observability capabilities for workloads such as ML inference.
Production	Observability	Ray Dashboard	Grafana + Enhanced Ray Dashboard + Persistent logs	With OSS, the developer needs to instrument their own logging and monitoring stack using Prometheus or Grafana and build additional infrastructure for logs persistence and access controls. This is all managed for you on Anyscale.
Cost saving	Cluster auto suspension	None	Built-in	Anyscale provides automatic cluster termination to reduce wasted spend.
Support	Support	Community Support	Enterprise support	Anyscale supports SSO integration and provides capabilities for access control
Governance	Authentication & Access Control	None	SSO integration and access controls	Dedicated engineer for one month and prioritized support.
Compliance	Production	None	SOC 2 Type 2	Anyscale upholds rigorous security and privacy controls.

For more on capabilities and benefits of the Anyscale Platform and Open Source Ray, [read our latest blogs.](#)