

Ray.io

Course outline

Module 1: Introduction to Ray

Module 1: Introduction to Ray module for Ray.io course provides an overview of the Ray framework and its components. It covers topics such as the Ray API, Ray actors, Ray tasks, and Ray clusters. It also provides an introduction to the Ray programming model and how to use Ray to build distributed applications.

Lessons

- Overview of Ray and its features
- Setting up Ray on your local machine
- Introduction to Ray's distributed execution engine
- Ray's distributed storage system
- Ray's distributed task scheduling system
- Ray's distributed resource management system
- Ray's distributed debugging system
- Ray's distributed monitoring system
- Ray's distributed logging system
- Ray's distributed optimization system

After completing this module, students will be able to:

- Understand the fundamentals of Ray, including its architecture, components, and use cases.
- Utilize Ray to parallelize and scale Python applications.
- Create Ray clusters and deploy applications to them.
- Monitor and debug Ray applications.

Module 2: Ray Basics

Module 2: Ray Basics is a course designed to introduce users to the basics of Ray.io, a distributed computing platform. This module covers topics such as Ray's architecture, how to set up a Ray cluster, and how to use Ray to run distributed applications. It also provides an introduction to Ray's API and how to use it to write distributed applications.

Lessons

- Introduction to Ray
- Ray Architecture

- Ray Components
- Ray API
- Ray Schedulers
- Ray Actors
- Ray Resources
- Ray Debugging
- Ray Performance Tuning
- Ray Deployment Strategies

After completing this module, students will be able to:

- Understand the fundamentals of Ray, including its architecture, components, and use cases.
- Utilize Ray to parallelize and scale Python applications.
- Create Ray actors and tasks to distribute workloads.
- Monitor and debug Ray applications using the Ray dashboard.

Module 3: Ray API

Module 3 of the Ray.io course introduces the Ray API, a powerful set of tools for building distributed applications. It covers topics such as creating and managing Ray clusters, using the Ray API to create distributed applications, and deploying applications to the cloud. It also provides an introduction to Ray's advanced features, such as distributed hyperparameter tuning and distributed reinforcement learning.

Lessons

- Introduction to Ray API
- Ray API Architecture
- Ray API Components
- Ray API Security
- Ray API Performance Tuning
- Ray API Debugging
- Ray API Deployment Strategies
- Ray API Best Practices
- Ray API Integrations
- Ray API Troubleshooting

After completing this module, students will be able to:

- Understand the fundamentals of the Ray API and its components.
- Develop distributed applications using Ray.
- Utilize Ray's distributed execution engine to scale applications.
- Deploy and manage Ray clusters in production environments.

Module 4: Ray Schedulers

Module 4 of the Ray.io course introduces Ray Schedulers, a powerful tool for managing distributed workloads. This module covers the basics of Ray Schedulers, including how to set up and configure a Ray Scheduler, how to use the Ray API to submit tasks to the scheduler, and how to monitor and debug tasks. Additionally, this module covers advanced topics such as fault tolerance, resource allocation, and scheduling strategies.

Lessons

- Introduction to Ray Schedulers
- Understanding Ray Scheduler Components
- Configuring Ray Schedulers
- Implementing Ray Schedulers
- Troubleshooting Ray Schedulers
- Advanced Ray Scheduler Strategies
- Optimizing Ray Schedulers
- Integrating Ray Schedulers with Other Systems
- Best Practices for Ray Schedulers
- Security Considerations for Ray Schedulers

After completing this module, students will be able to:

- Understand the fundamentals of Ray Schedulers and how to use them to manage workloads.
- Create and deploy Ray Schedulers to manage workloads in a distributed environment.
- Monitor and debug Ray Schedulers to ensure optimal performance.
- Utilize Ray Schedulers to optimize resource utilization and improve scalability.

Module 5: Ray Actors

Module 5 of the Ray.io course introduces Ray Actors, a powerful tool for distributed computing. Ray Actors allow users to create and manage distributed applications with ease. This module covers the basics of Ray Actors, including how to create and manage them, how to use them to scale up applications, and how to debug them. Additionally, this module provides an overview of the Ray API and how to use it to interact with Ray Actors.

Lessons

- Introduction to Ray Actors
- Building and Deploying Ray Actors
- Working with Ray Actors
- Debugging Ray Actors
- Optimizing Ray Actors
- Advanced Ray Actors Concepts
- Ray Actors in Production
- Best Practices for Ray Actors
- Troubleshooting Ray Actors
- Integrating Ray Actors with Other Technologies

After completing this module, students will be able to:

- Understand the fundamentals of Ray Actors and how to use them to build distributed applications.
- Create Ray Actors and use them to communicate with other Ray Actors.
- Utilize Ray Actors to scale applications across multiple machines.
- Debug and troubleshoot Ray Actors applications.

Module 6: Ray Optimization

Module 6 of the Ray.io course focuses on Ray Optimization. This module covers topics such as how to optimize Ray applications, how to use Ray's distributed scheduling algorithms, and how to use Ray's distributed resource management system. Additionally, this module provides hands-on exercises to help students gain practical experience in optimizing Ray applications.

Lessons

- Introduction to Ray Optimization
- Ray Autoscaler
- Ray Tune Hyperparameter Optimization
- Ray RLlib for Reinforcement Learning
- Ray SGD for Distributed Training
- Ray Actors for Parallel Computing
- Ray Serve for Model Serving
- Ray Debugger for Debugging
- Ray Metrics for Monitoring
- Ray Deploy for Deployment

After completing this module, students will be able to:

- Understand the fundamentals of Ray optimization and how to use it to improve the performance of distributed applications.
- Utilize Ray's optimization features to reduce the latency of distributed applications.
- Implement strategies to optimize the resource utilization of Ray clusters.
- Analyze the performance of Ray applications and identify areas for improvement.

Module 7: Ray Debugging

Module 7 of the Ray.io course introduces students to Ray Debugging, a powerful tool for debugging distributed applications. This module covers topics such as how to use Ray Debugging to identify and fix errors in distributed applications, how to use the Ray Debugging API, and how to use the Ray Debugging dashboard to visualize and analyze distributed application performance.

Lessons

- Introduction to Ray Debugging
- Ray Debugging Tools and Techniques
- Debugging Ray Applications
- Debugging Ray Actors
- Debugging Ray Tasks

- Debugging Ray Services
- Debugging Ray Libraries
- Debugging Ray Clusters
- Debugging Ray Performance
- Troubleshooting Ray Issues

After completing this module, students will be able to:

- Understand the fundamentals of debugging with Ray.io
- Utilize the debugging tools available in Ray.io to identify and resolve errors
- Develop strategies for debugging complex Ray.io applications
- Identify and troubleshoot common issues with Ray.io applications

Module 8: Ray Deployment

Module 8 of the Ray.io course covers the deployment of Ray applications. It covers topics such as deploying Ray applications to the cloud, deploying Ray applications to Kubernetes, and deploying Ray applications to a local cluster. It also covers topics such as debugging and monitoring Ray applications.

Lessons

- Introduction to Ray Deployment
- Setting up a Ray Cluster
- Deploying Ray Applications
- Managing Ray Clusters
- Scaling Ray Clusters
- Monitoring Ray Clusters
- Troubleshooting Ray Clusters
- Best Practices for Ray Deployment
- Advanced Ray Deployment Strategies
- Security Considerations for Ray Deployment

After completing this module, students will be able to:

- Understand the fundamentals of Ray and its components
- Deploy Ray clusters on cloud providers such as AWS, GCP, and Azure
- Utilize Ray for distributed training and hyperparameter optimization
- Monitor and debug Ray applications using the Ray dashboard and other tools

Module 9: Ray Security

Module 9 of the Ray.io course focuses on security, providing an overview of the security features of Ray and how to use them to protect your applications. It covers topics such as authentication, authorization, encryption, and secure communication. It also provides guidance on how to use Ray's security features to protect your applications from malicious actors.

Lessons

- Introduction to Ray Security
- Authentication and Authorization
- Securing Ray Applications
- Encryption and Data Protection
- Network Security
- Auditing and Logging
- Best Practices for Ray Security
- Troubleshooting Ray Security Issues
- Ray Security Tools and Resources

After completing this module, students will be able to:

- Understand the fundamentals of Ray Security and how to use it to protect applications and data.
- Implement authentication and authorization mechanisms to secure Ray applications.
- Utilize Ray Security to protect data in transit and at rest.
- Monitor and audit Ray Security configurations to ensure compliance with security policies.

Module 10: Ray Performance Tuning

Module 10 of the Ray.io course focuses on performance tuning for Ray applications. It covers topics such as profiling and debugging, optimizing Ray applications, and using Ray's distributed scheduler to improve performance. Additionally, the module provides guidance on how to use Ray's distributed resource management system to manage resources and improve scalability.

Lessons

- Introduction to Ray Performance Tuning
- Profiling Ray Applications
- Optimizing Ray Applications
- Strategies for Improving Ray Performance
- Debugging Ray Applications
- Tuning Ray Schedulers
- Tuning Ray Workers
- Tuning Ray Actors
- Tuning Ray Resources
- Monitoring Ray Performance

After completing this module, students will be able to:

- Understand the fundamentals of Ray performance tuning and how to apply them to optimize Ray applications.
- Utilize Ray's built-in performance tuning tools to identify and address performance bottlenecks.
- Analyze and debug Ray applications to identify and address performance issues.
- Implement best practices for Ray performance tuning to ensure optimal performance.

Module 11: Ray Best Practices

Module 11 of the Ray.io course provides an overview of best practices for using Ray, a distributed computing framework. It covers topics such as how to optimize Ray for performance, how to debug Ray applications, and how to use Ray for distributed training. It also provides guidance on how to use Ray for distributed hyperparameter tuning and distributed reinforcement learning.

Lessons

- Understanding Ray Architecture
- Designing Efficient Ray Applications
- Optimizing Ray Performance
- Debugging Ray Applications
- Deploying Ray Applications
- Scaling Ray Applications
- Integrating Ray with Other Frameworks
- Using Ray for Distributed Machine Learning
- Using Ray for Distributed Reinforcement Learning
- Using Ray for Distributed Hyperparameter Tuning

After completing this module, students will be able to:

- Understand the fundamentals of Ray, including its architecture, components, and features.
- Utilize Ray's distributed computing capabilities to scale up applications.
- Implement best practices for deploying and managing Ray clusters.
- Develop applications using Ray's APIs and libraries.

Module 12: Ray in Production

Module 12: Ray in Production is a module in the Ray.io course that covers the practical aspects of deploying Ray applications in production. It covers topics such as scaling Ray applications, monitoring and debugging, and deploying Ray applications to the cloud. It also covers best practices for deploying Ray applications in production.

Lessons

- Introduction to Ray in Production
- Deploying Ray on Cloud Platforms
- Scaling Ray with Kubernetes
- Monitoring Ray Applications
- Debugging Ray Applications
- Optimizing Ray Performance
- Integrating Ray with Other Services
- Security and Compliance for Ray
- Automating Ray Deployments
- Advanced Ray Features and Use Cases

After completing this module, students will be able to:

- Understand the fundamentals of Ray and its components
- Deploy Ray clusters on cloud providers and local machines
- Utilize Ray for distributed training and hyperparameter optimization
- Monitor and debug Ray applications in production environments