
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

OPN203-R

Get started with OpenSearch

Arun Lakshmanan (he/him)

OpenSearch Specialist Solutions Architect
AWS

Jon Handler (he/him)

Sr Principal Solutions Architect
AWS

Agenda

- What is OpenSearch?
- Why OpenSearch?
- Search with OpenSearch
- Operational analytics with OpenSearch
- Hands-on lab – Dashboarding and monitoring

What is OpenSearch?

OpenSearch is a fork

Community-driven, open-source search and analytics suite derived from Apache 2.0–licensed Elasticsearch 7.10.2 and Kibana 7.10.2

It consists of

- A search engine – OpenSearch
- A visualization and user interface – OpenSearch Dashboards
- Tools and plugins adding series functionality

Why OpenSearch?

100% open source

Use, modify, monetize

Built in the open

No contributor license agreement

Community driven

Made with your input

Public road map

Contributions welcome

Full featured

Full text search

Analytics

Fast ingest

Machine learning

Security

And more . . .

Community driven

- Over 40 partners and growing
- Available with multiple service providers on multiple clouds (Oracle, Aiven-Azure, AWS, Bonsai-GCP)
- Hundreds of contributors
- Thousands of pull requests
- 400+ contributions outside of Amazon in 2022
- 100M+ project downloads
- Top 4 search engine on DB-Engines
- Clients (JS, Java, Rust, Go, .Net, Python, Spring, Hadoop & more)
- Growing number of community projects (for example, K8s operator, Terraform module, and more)

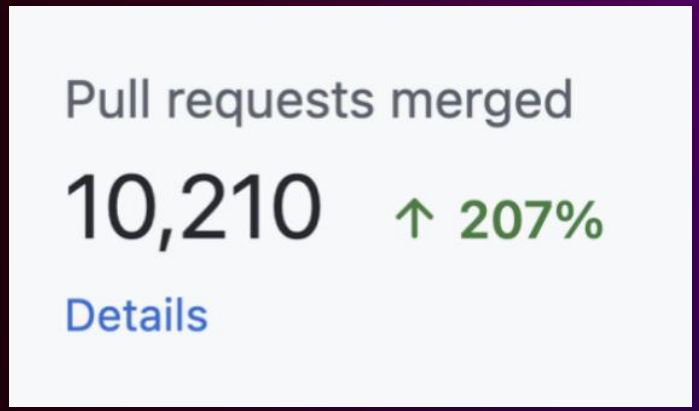

OpenSearch: 7 major releases, 90+ features

OpenSearch 1.2

- Observability Interface
- Feature attribution in anomaly detection
- Shard-level indexing back-pressure
- More efficient k-NN dense vectors

OpenSearch 2.0

- Lucene 9.1 performance optimizations
- Document-level alerting
- RPM Package Manager

OpenSearch 2.2

- Logistic regression
- Lucene HNSW implementation
- Custom GeoJSON
- RCFSummarize algorithm
- Support in region maps

OpenSearch 2.4

- Flat data type support
- Security analytics MVP
- Windows support for k-NN plugin
- Spring Data and Hadoop clients

Q1'22 Q2'22 Q3'22 Q4'22

OpenSearch 1.3

- App Analytics, Trace ID correlation, and Live Tail in Observability
- PPL runtime support
- K-means and Random Cut Forest algorithm support

OpenSearch 2.1

- Automate snapshots with Snapshot Management
- Dedicated resources for ML workloads
- **Multi-terms aggregation**

OpenSearch 2.3

- Segment replication
- Remote-backed storage
- Drag-and-drop visualization
- OpenSearch Playground

aws

How does OpenSearch work?

OpenSearch is a database

1

Send data as JSON via REST APIs

2

Data is indexed – all fields searchable, including nested JSON

3

REST APIs, for fielded matching, Boolean expressions, sorting, and analysis

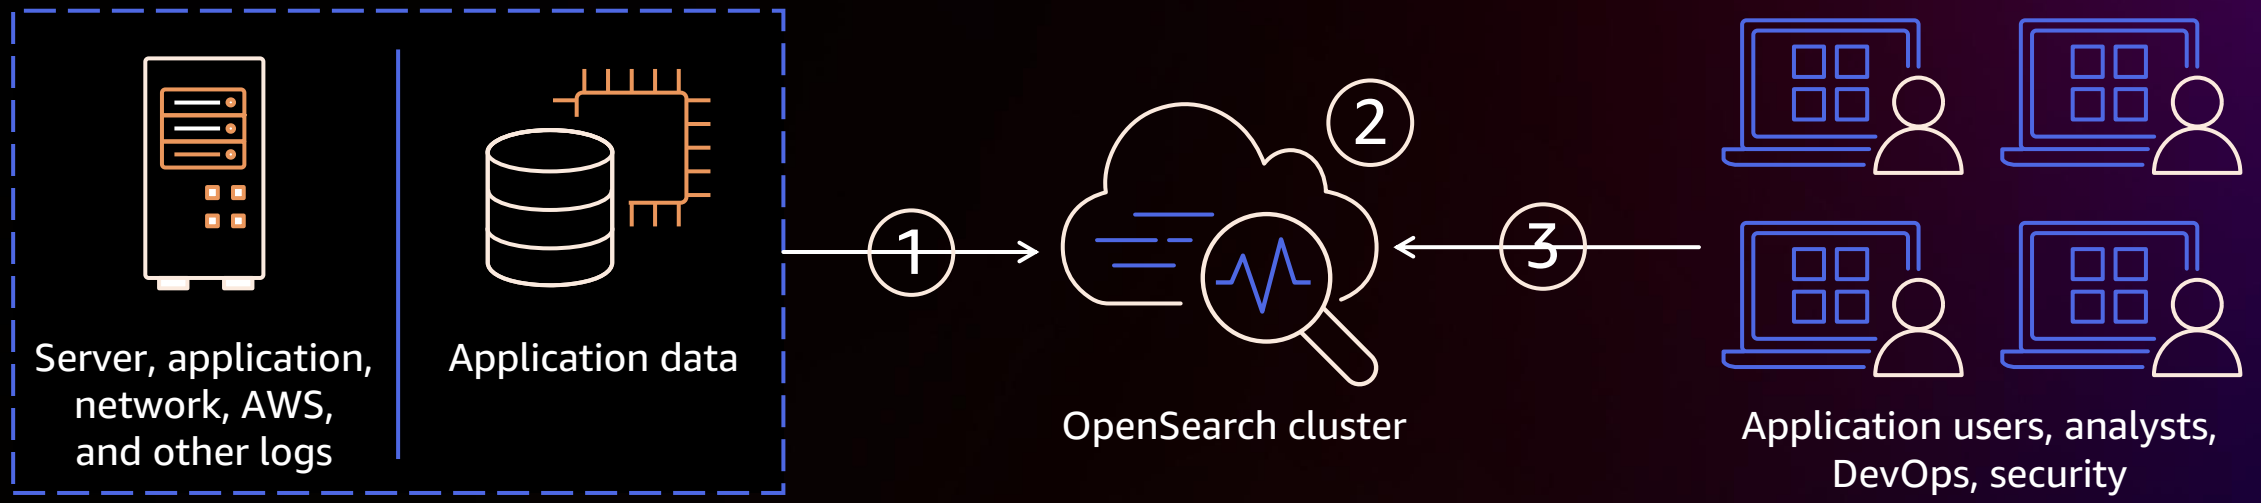

Operational Analytics

IMPROVE INTERNAL BUSINESS EFFICIENCY AND/OR CUSTOMER EXPERIENCE AND VALUE

- Traditional analytics vs operational analytics
  - Operational analytics
    - Descriptive analytics – Observability
    - Predictive analytics
    - Prescriptive analytics

Observability: data drives decisions

Logs

Amazon OpenSearch Service

Metrics

Amazon Managed Service for Prometheus
Amazon Managed Service for Grafana

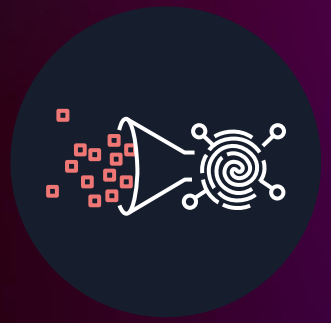

Traces

AWS Distro for OpenTelemetry
Fluent Bit

AWS monitoring and observability services help you maintain SLAs by **detecting, investigating, and remediating** problems to achieve

Availability Reliability Performance

Open source observability architecture

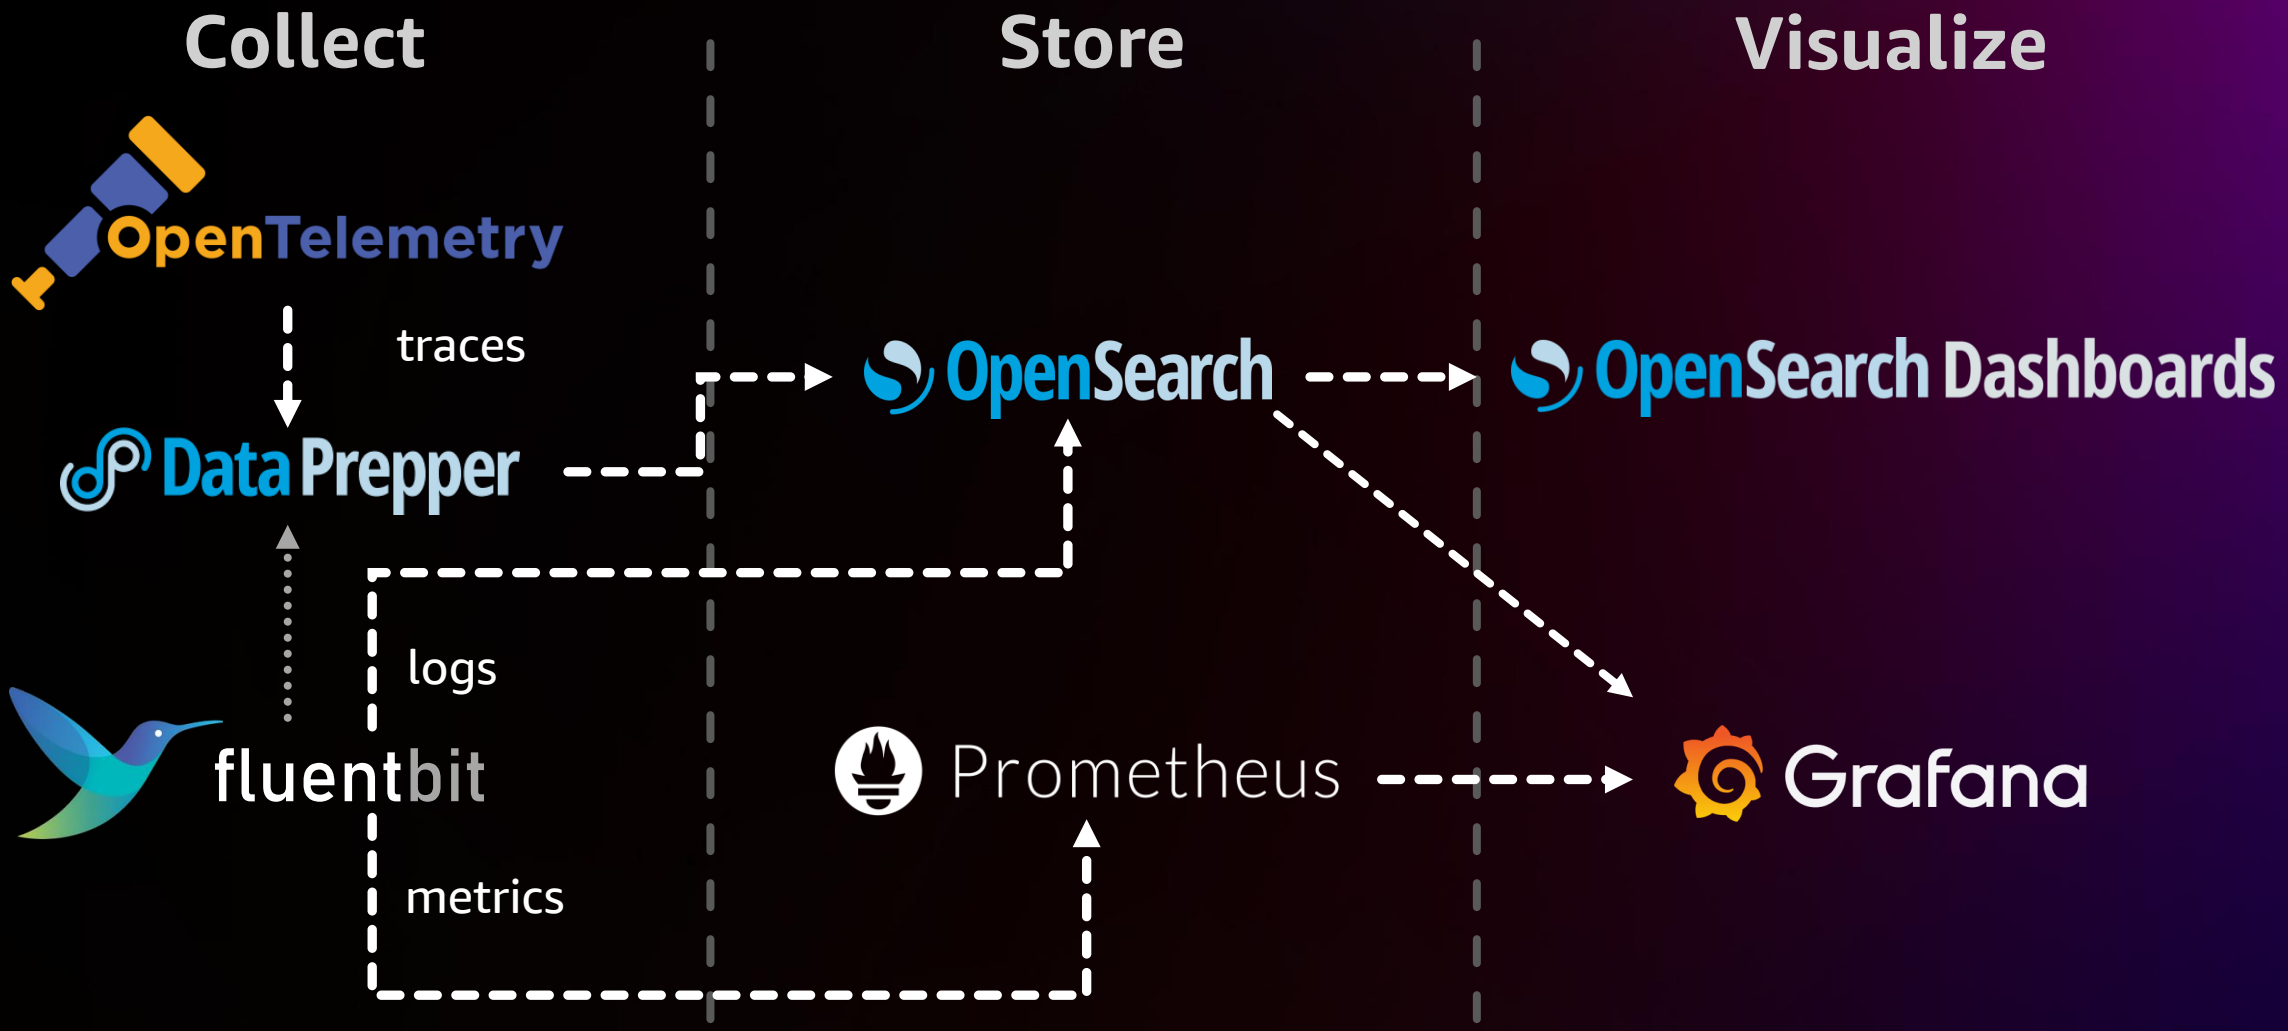

What is distributed tracing?

Identifying problems in cloud applications

A method of observing requests as they propagate through distributed systems

Trace: Hierarchical, end-to-end record of processing a request

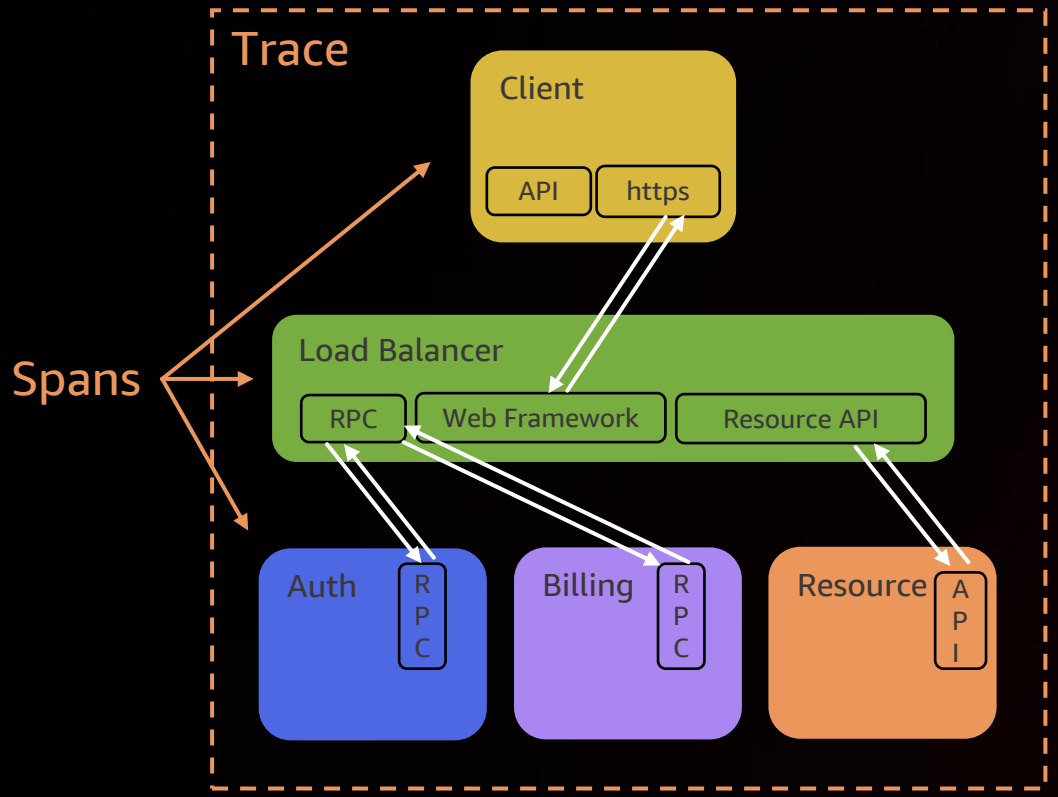

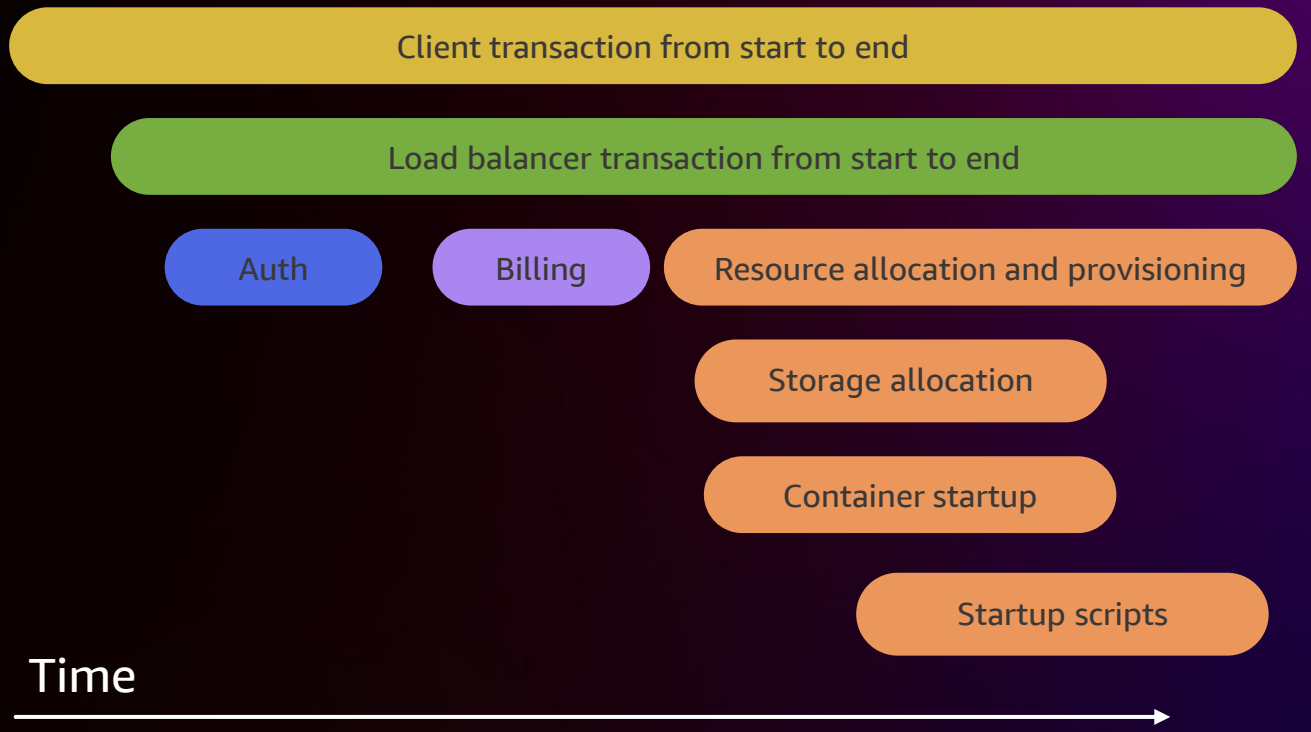

What is trace analytics?

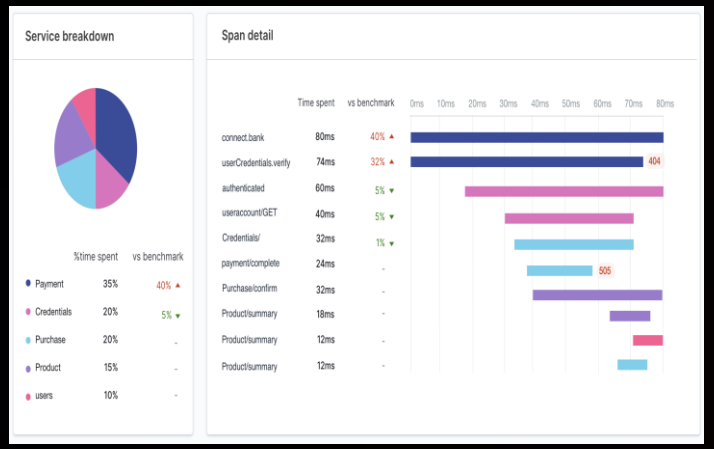

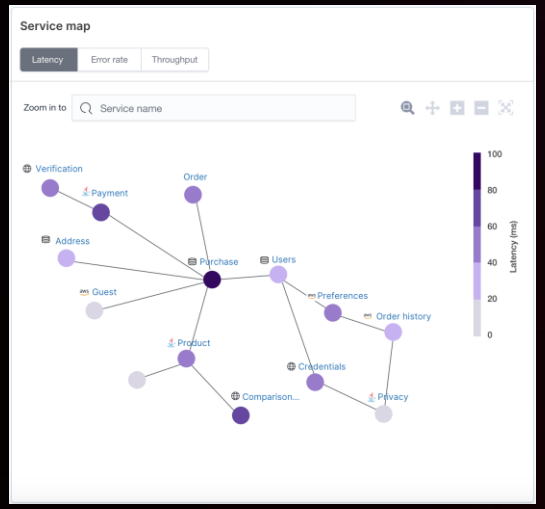

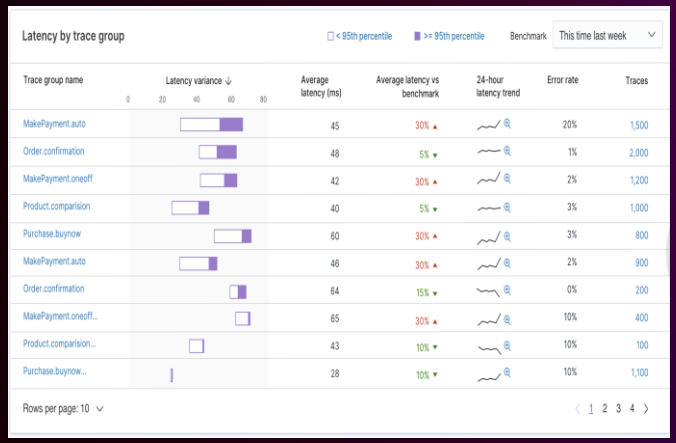

Trace-span details

- Single request performance
- Diagnose root cause

Service maps

- End-to-end view
- Isolate issues to services

Trace groups

- Monitor performance
- Identify issues early

What are you building?

Lab 1: Search application

Prebuilt set of microservices coordinate searching in OpenSearch

- Provision application data (IMDb Top 5000 titles)
- Publish search templates
- Review application monitoring with prebuilt dashboards

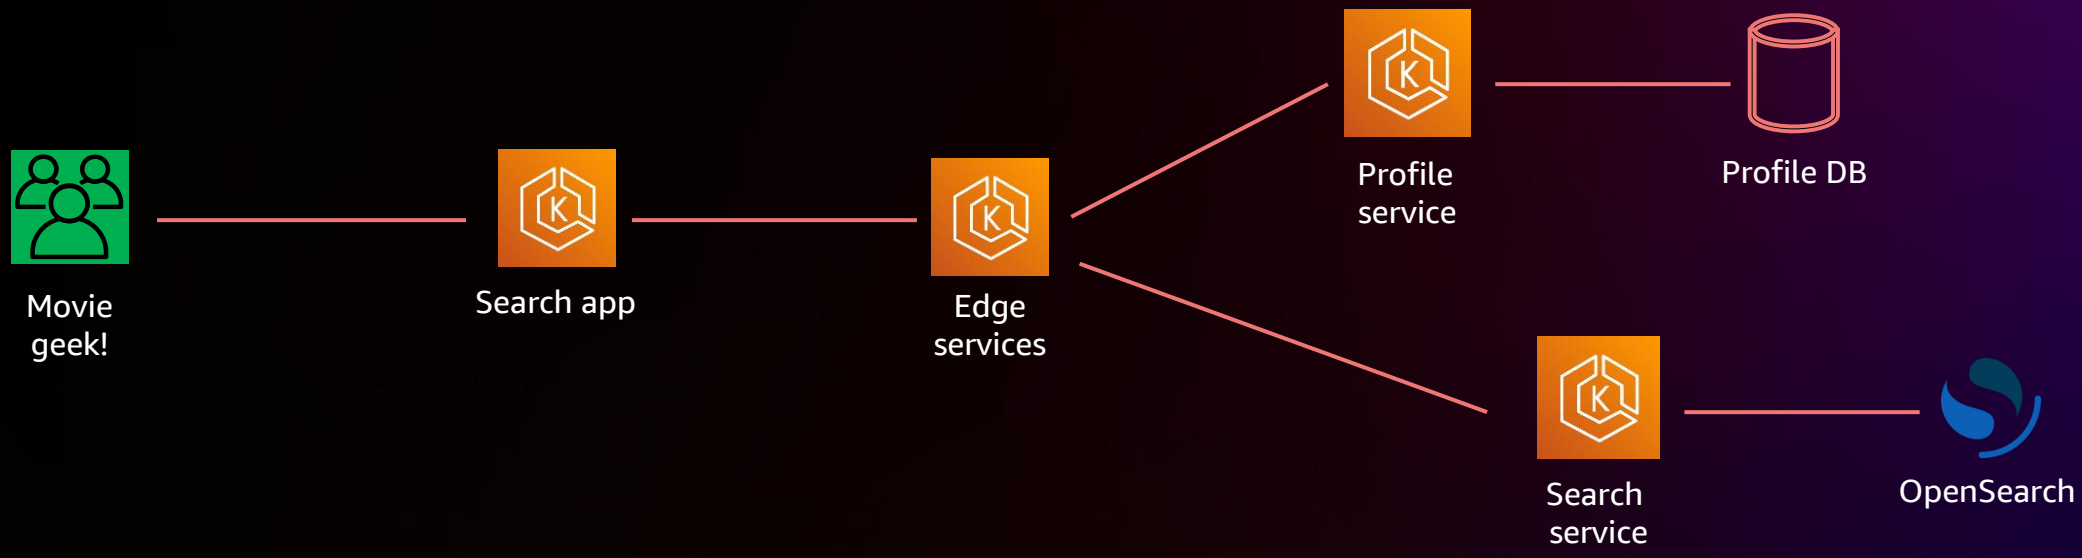

Lab 2: Distributed tracing

Review prebuilt trace data ingestion pipeline

- Simulate errors in services
- Debug and fix them

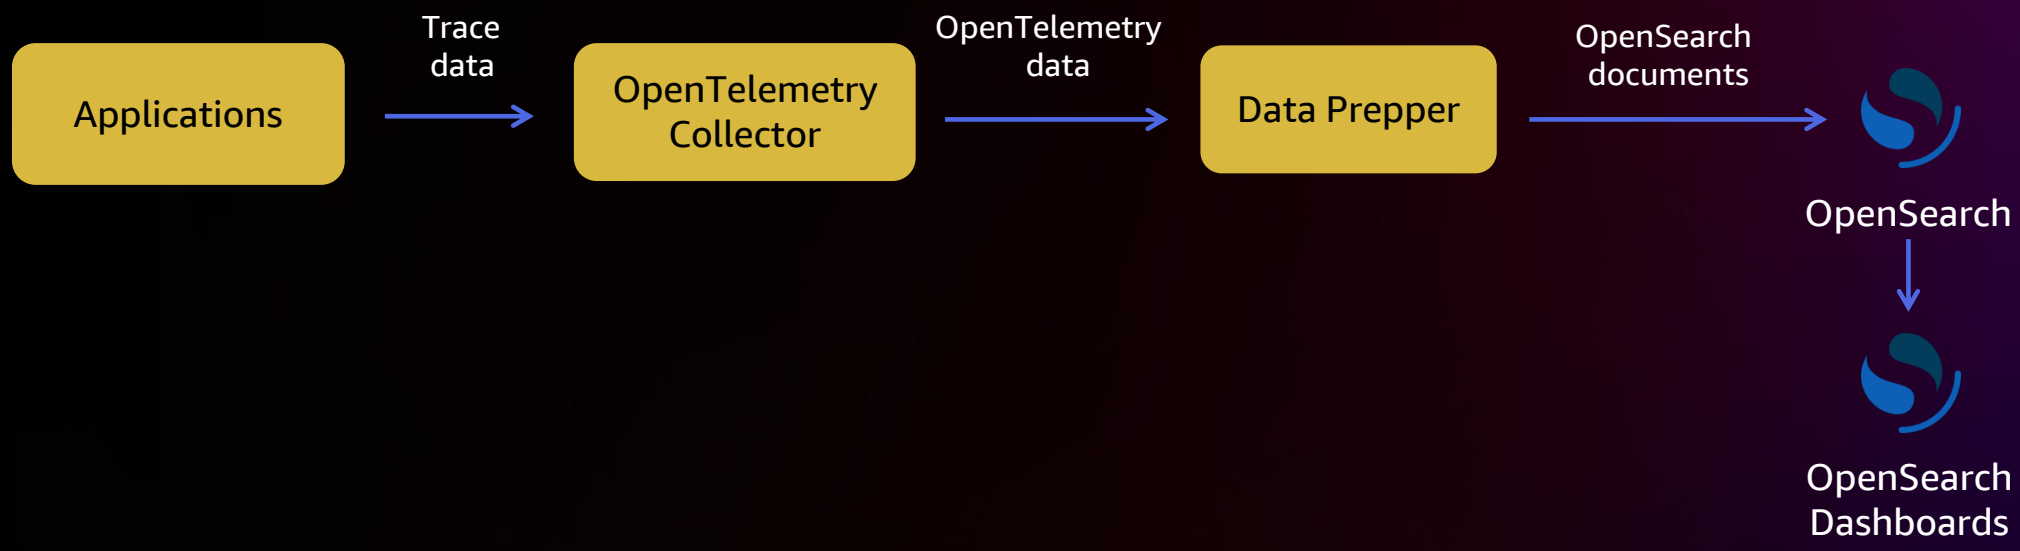

Lab 3: Monitoring Kubernetes cluster

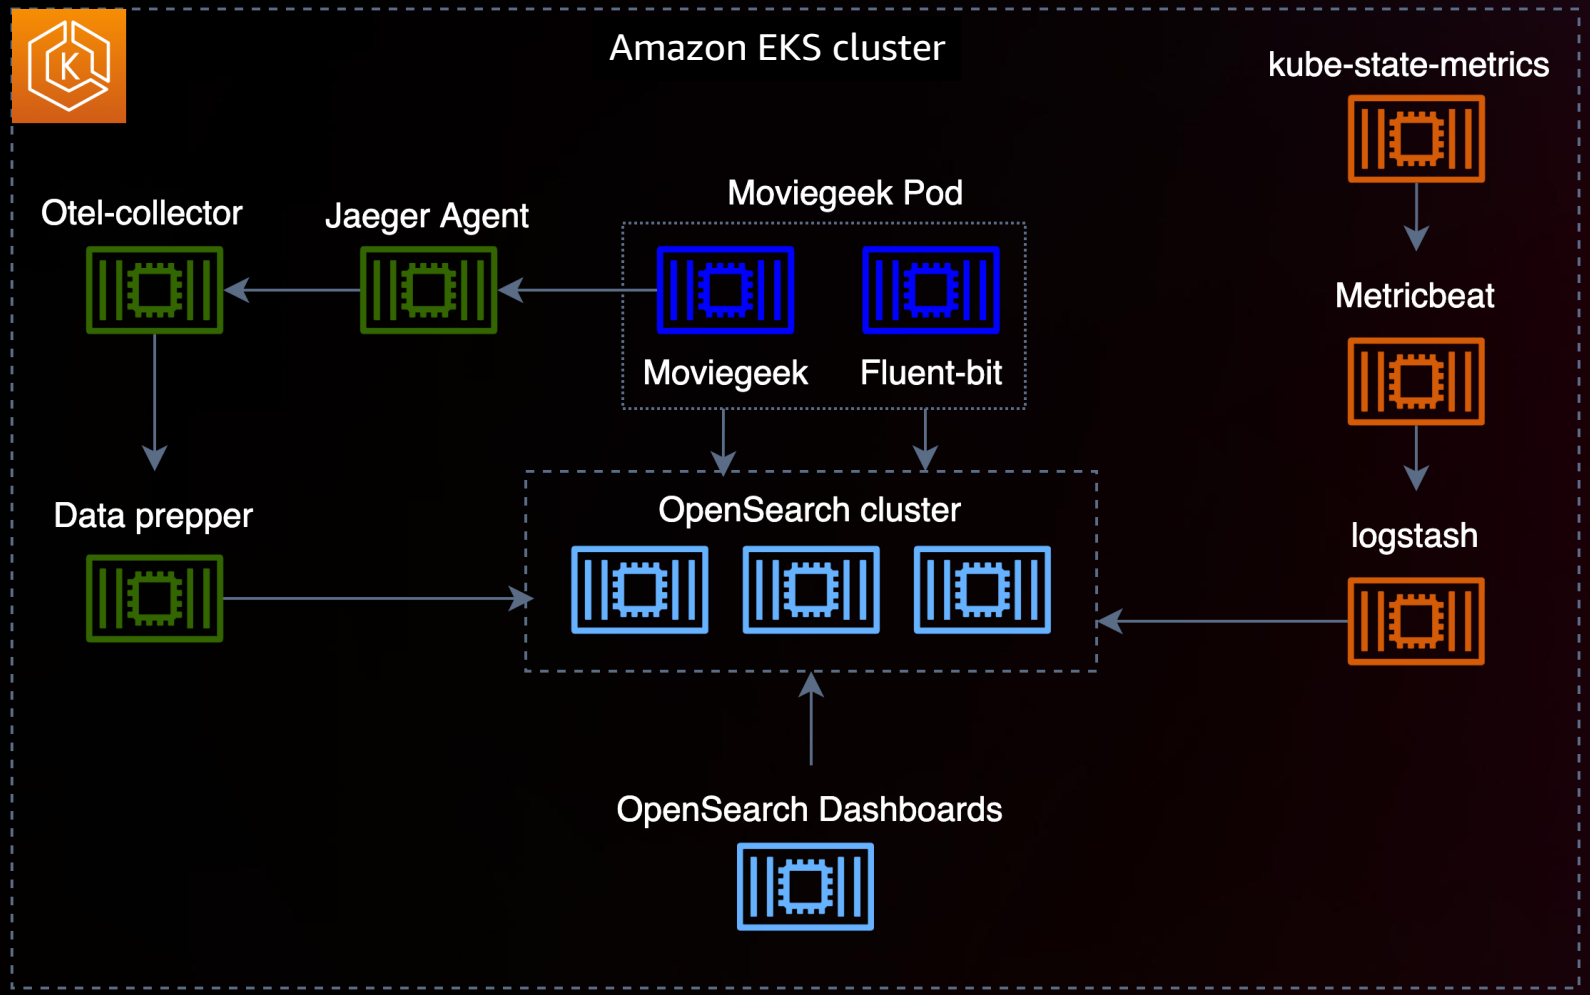

- 3 m5.large worker nodes

Lab 3: Monitoring Kubernetes cluster

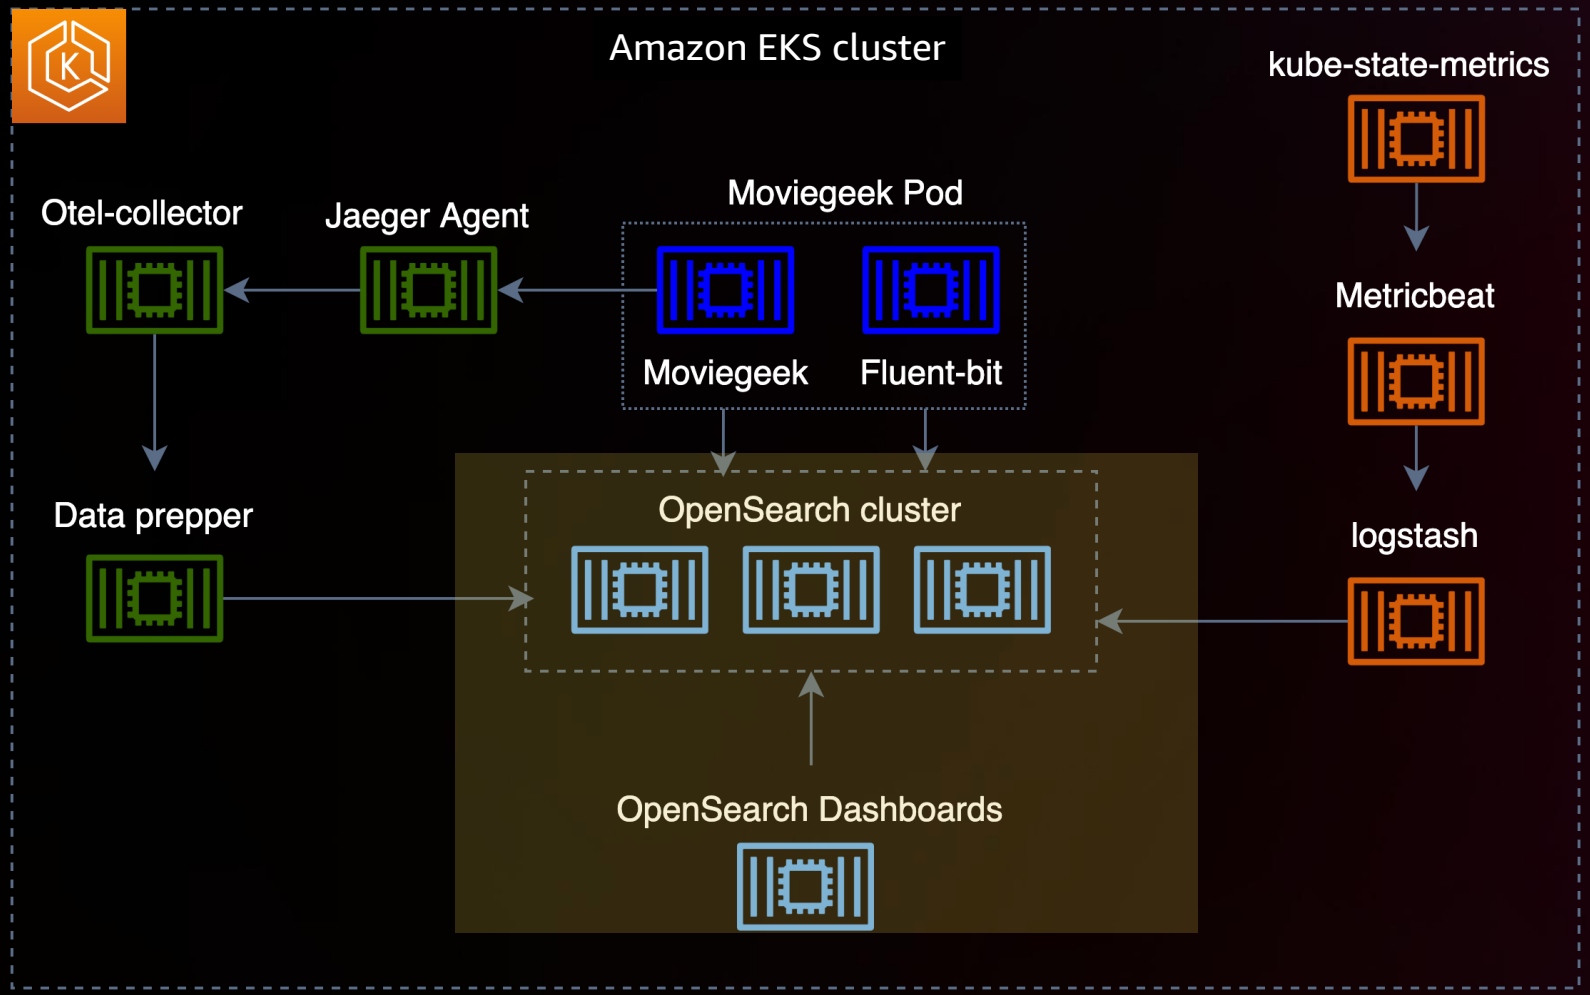

- 3 m5.large worker nodes
- OpenSearch Cluster and Dashboards
  - 4 Pods

Lab 3: Monitoring Kubernetes cluster

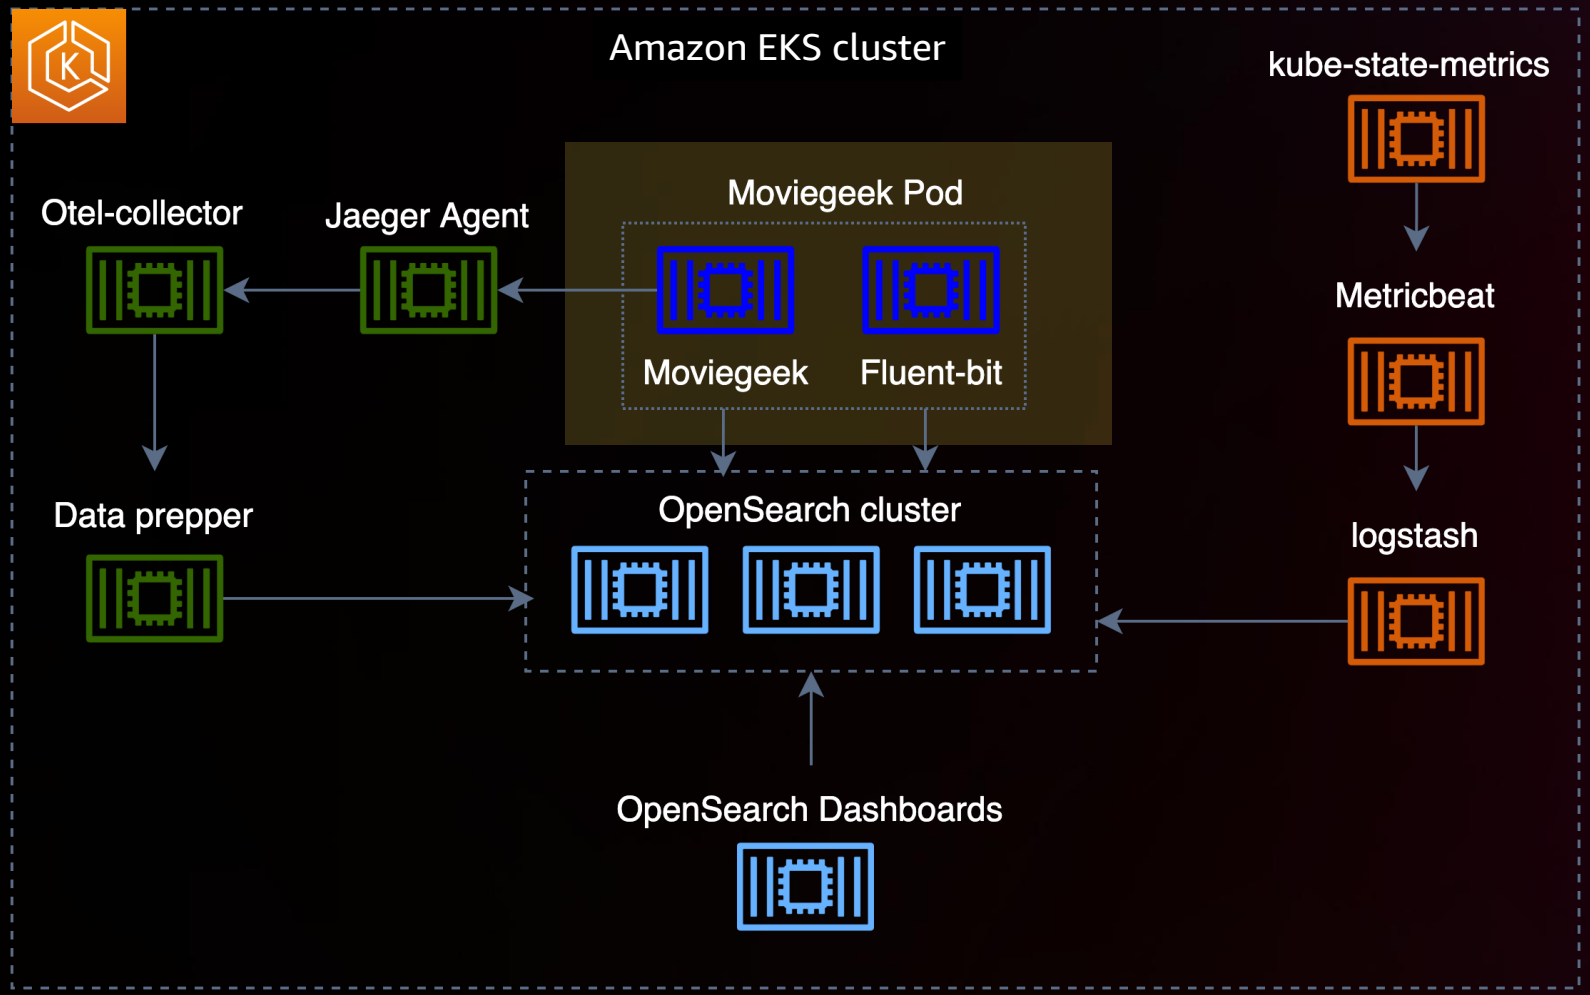

- 3 m5.large worker nodes
- OpenSearch Cluster and Dashboards
  - 4 Pods
- Application with Fluent-bit sidecar

Lab 3: Monitoring Kubernetes cluster

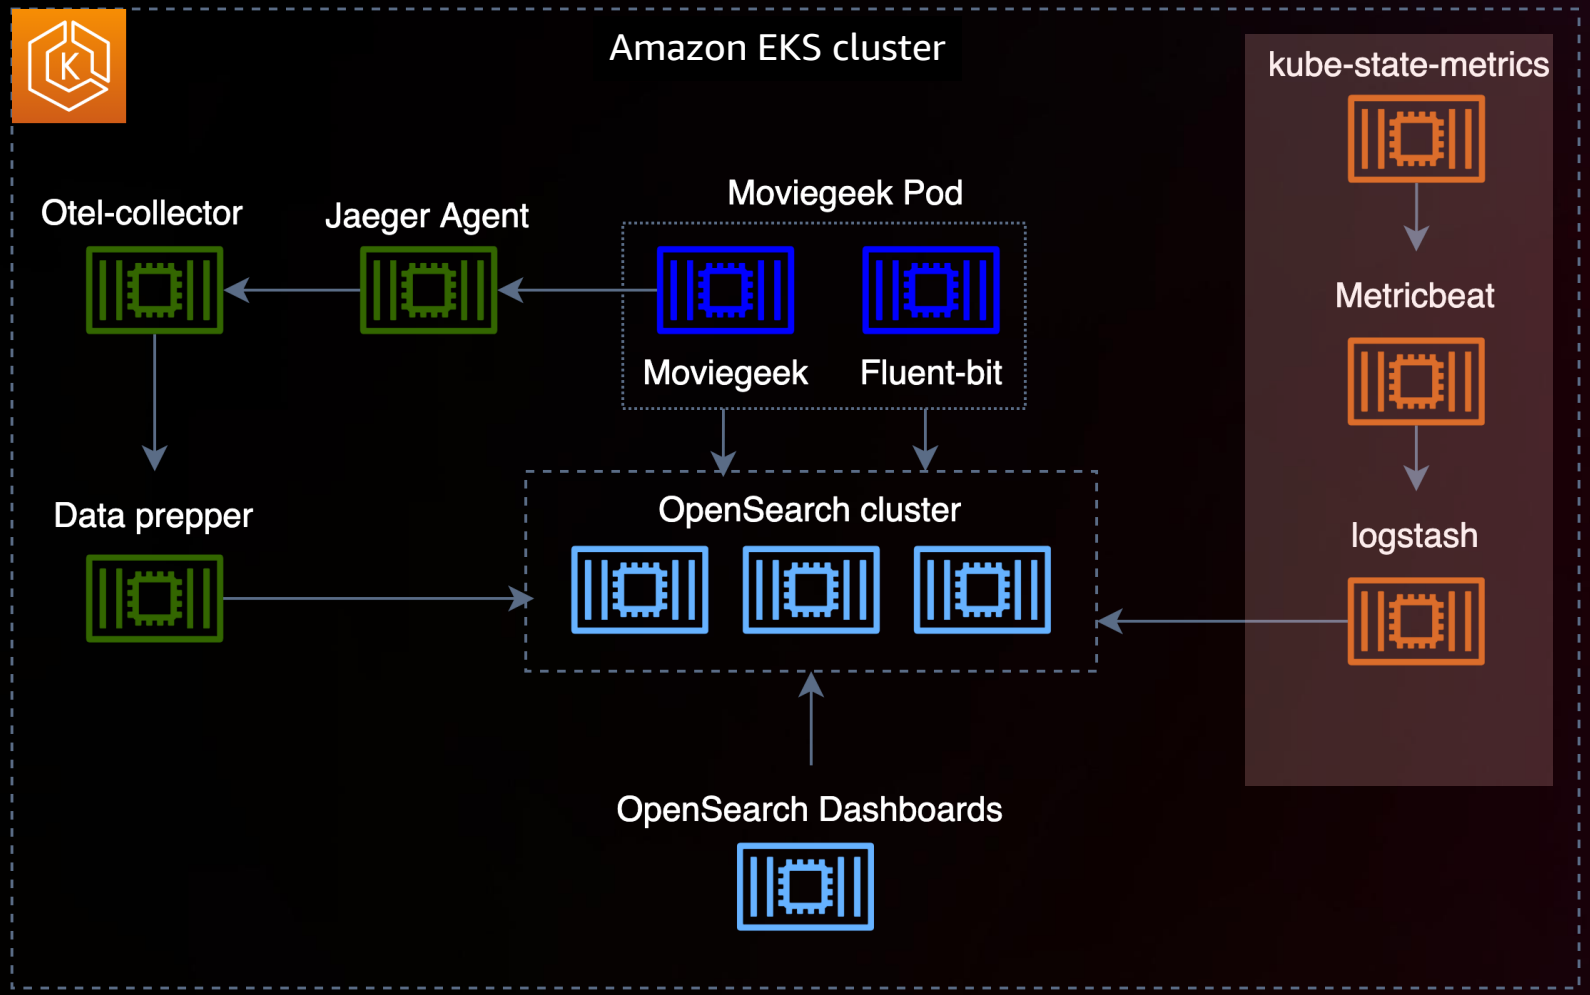

- 3 m5.large worker nodes
- OpenSearch Cluster and Dashboards
  - 4 Pods
- Application with Fluent-bit sidecar
- Cluster monitoring – 6 pods
  - Logstash
  - Metricbeat – daemonset
  - Kube-state-metrics

Lab 3: Monitoring Kubernetes cluster

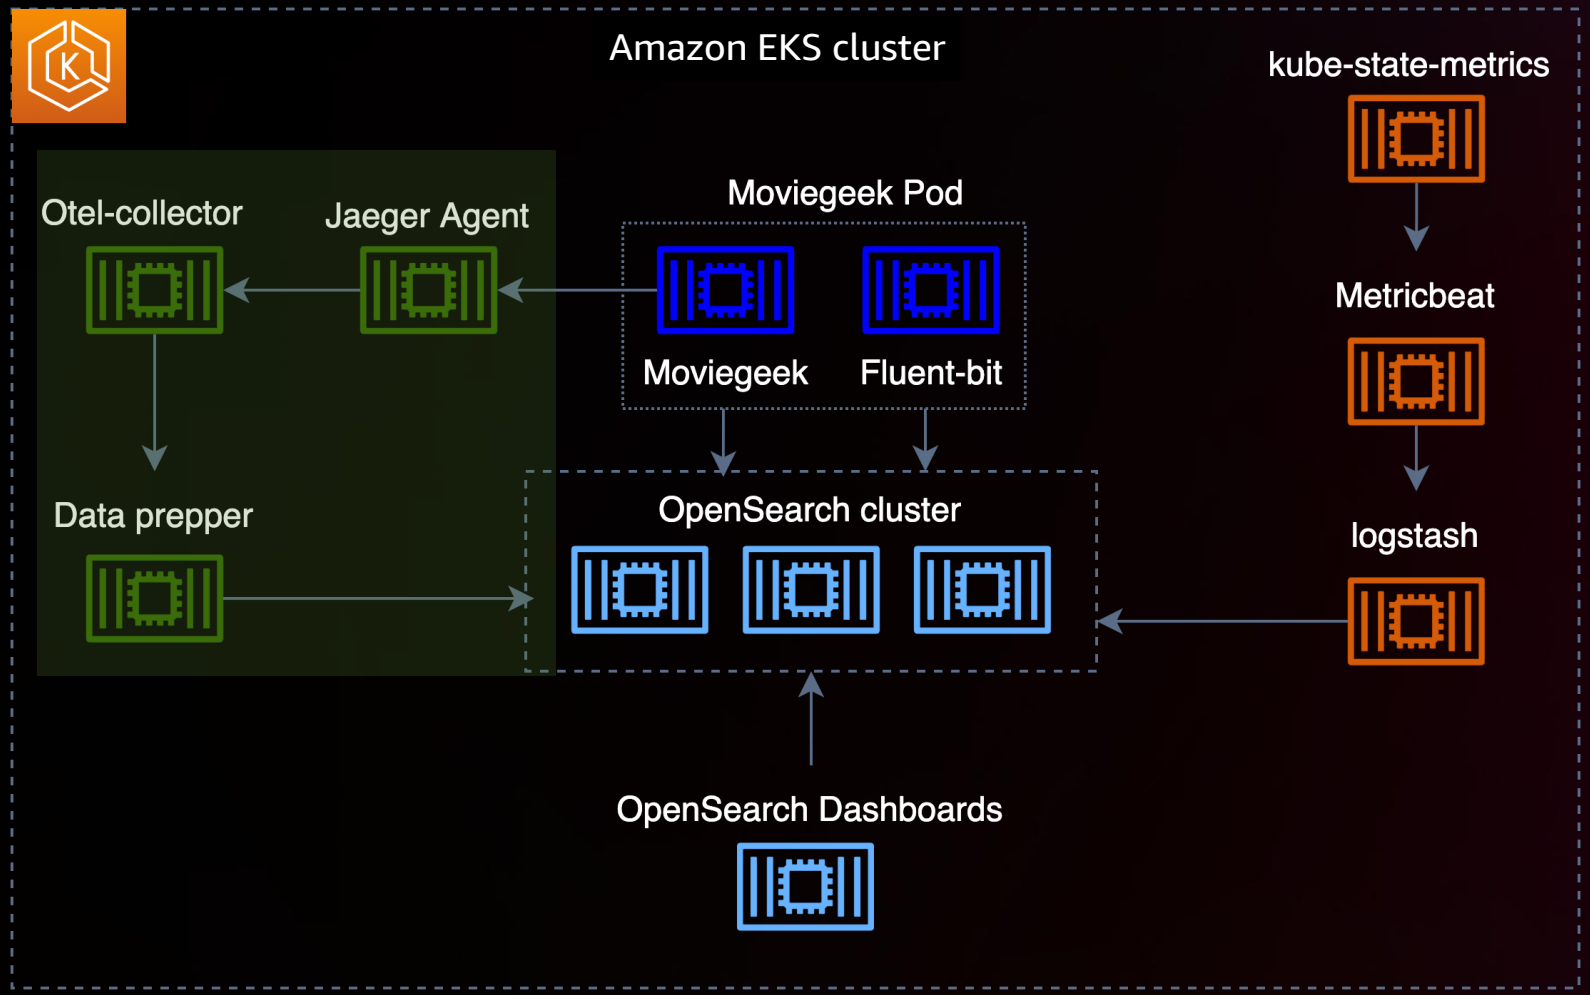

- 3 m5.large worker nodes
- OpenSearch cluster and dashboards
  - 4 Pods
- Application with Fluent-bit sidecar
- Cluster monitoring – 6 pods
  - Logstash
  - Metricbeat – daemonset
  - Kube-state-metrics
- Trace analytics – 3 pods
  - Jaeger Agent
  - Otel-collector
  - Data prepper

Lab architecture

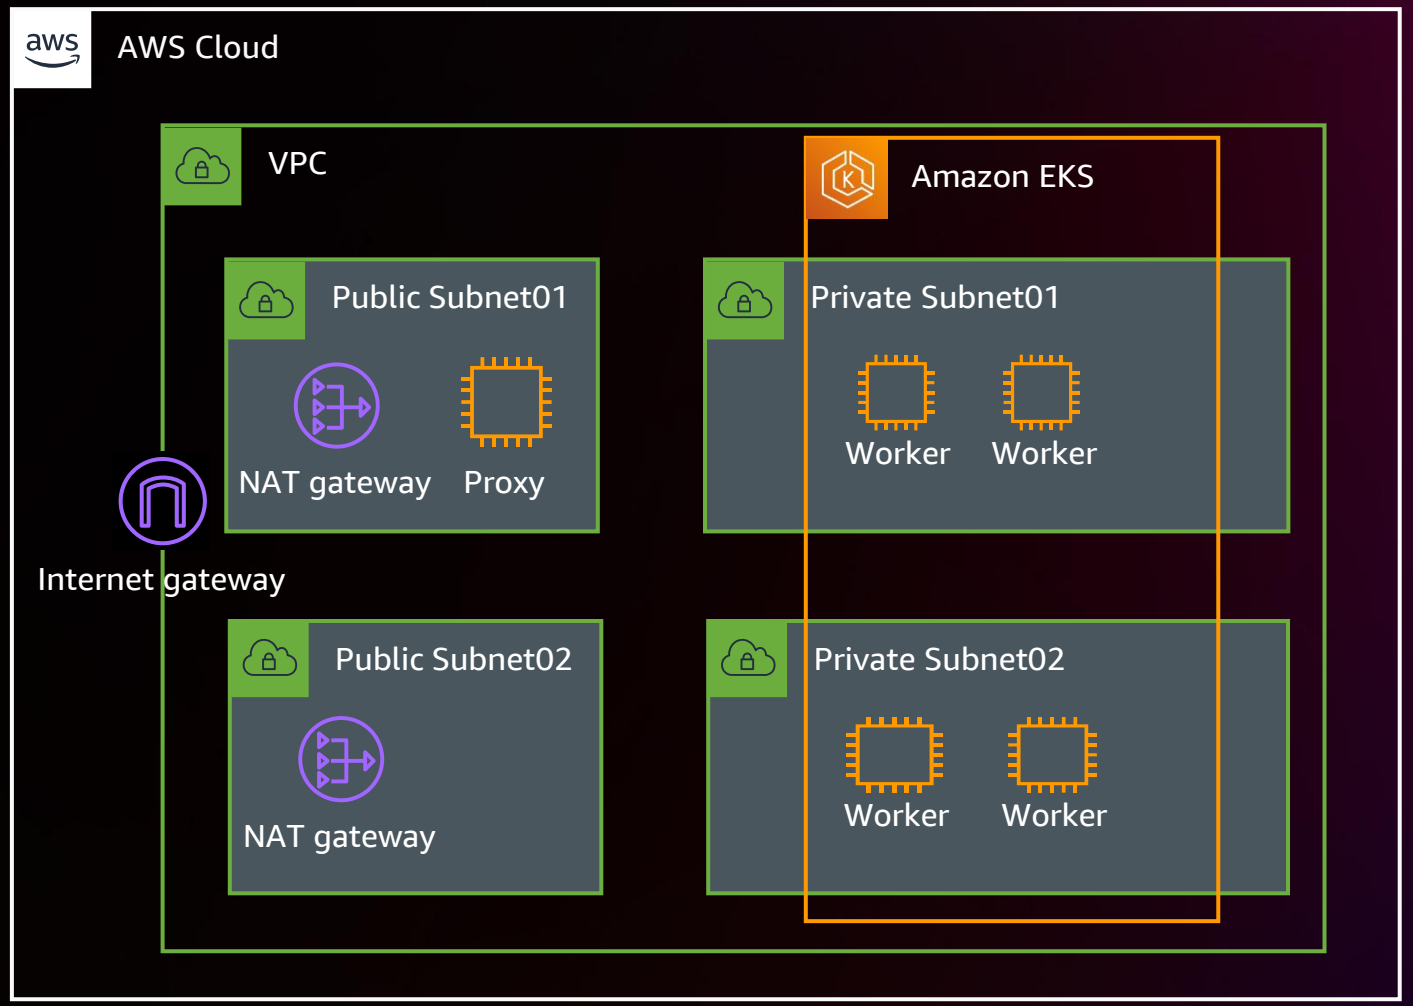

Get started with this workshop

- As a participant, you will have access to an AWS account with any optional pre-provisioned infrastructure and IAM policies needed to complete this workshop
- The AWS account will only be available for the duration of this workshop; you will lose access to the account thereafter
- The pre-provisioned infrastructure are deployed to the **us-east-1** Region
- Be sure to review the terms and conditions of the event; do not upload any personal or confidential information in the account

Step 1: Sign-in via OTP

https://catalog.workshops.aws/join

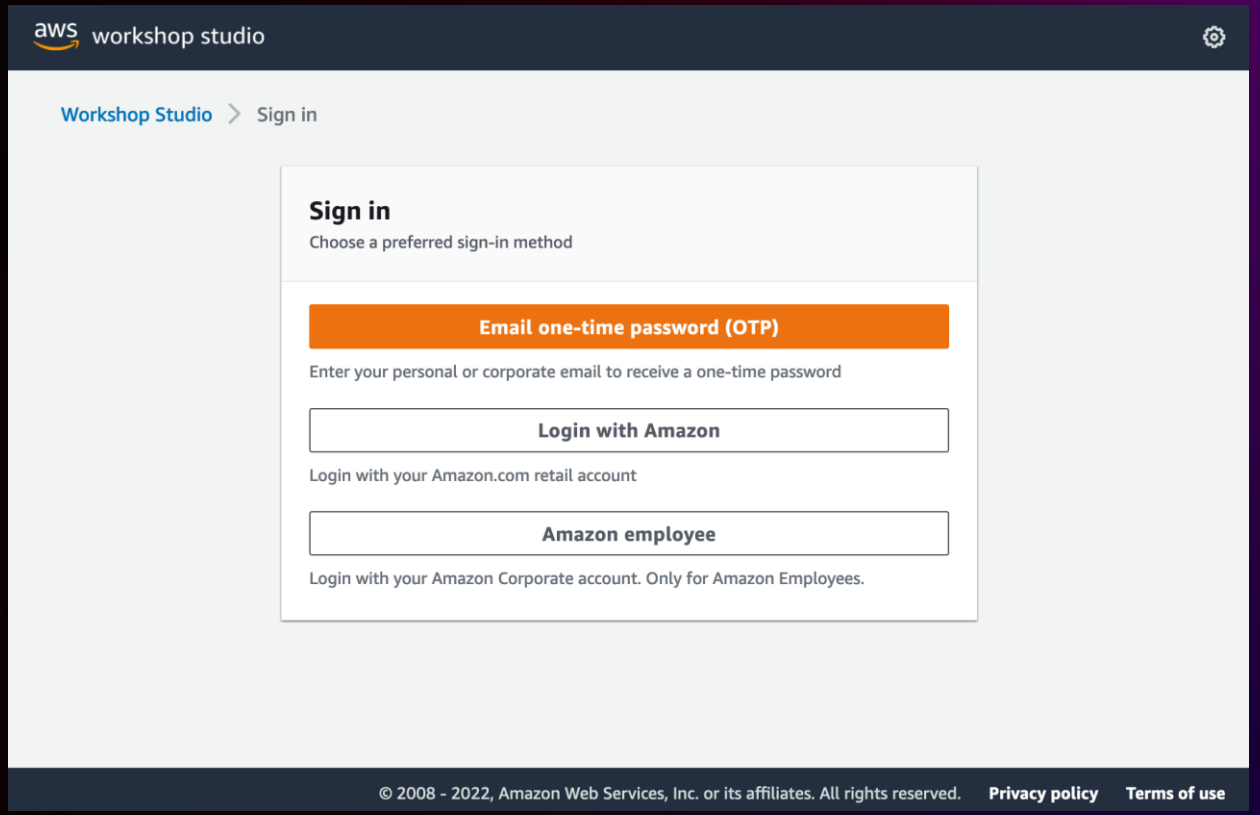

Step 2: Enter event access code

Enter 12-digit event access code

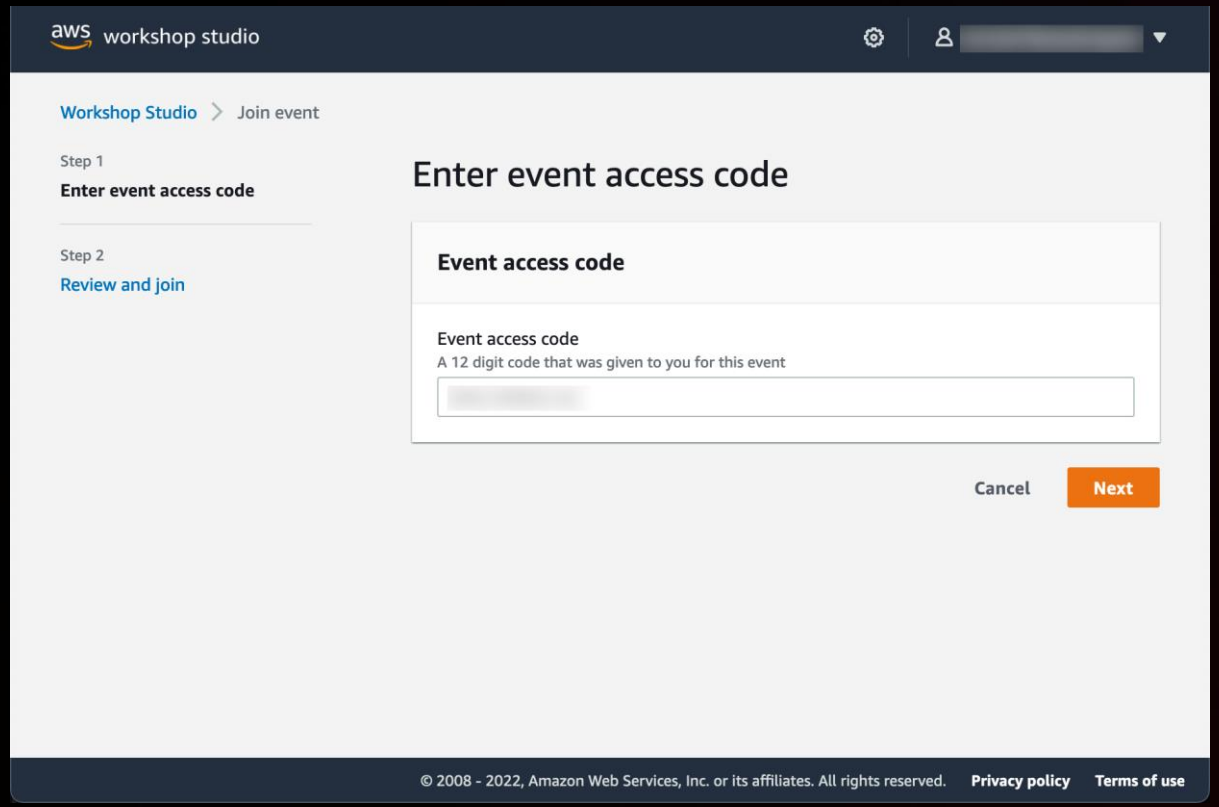

767a-037d58-84

Step 3: Review terms and join event

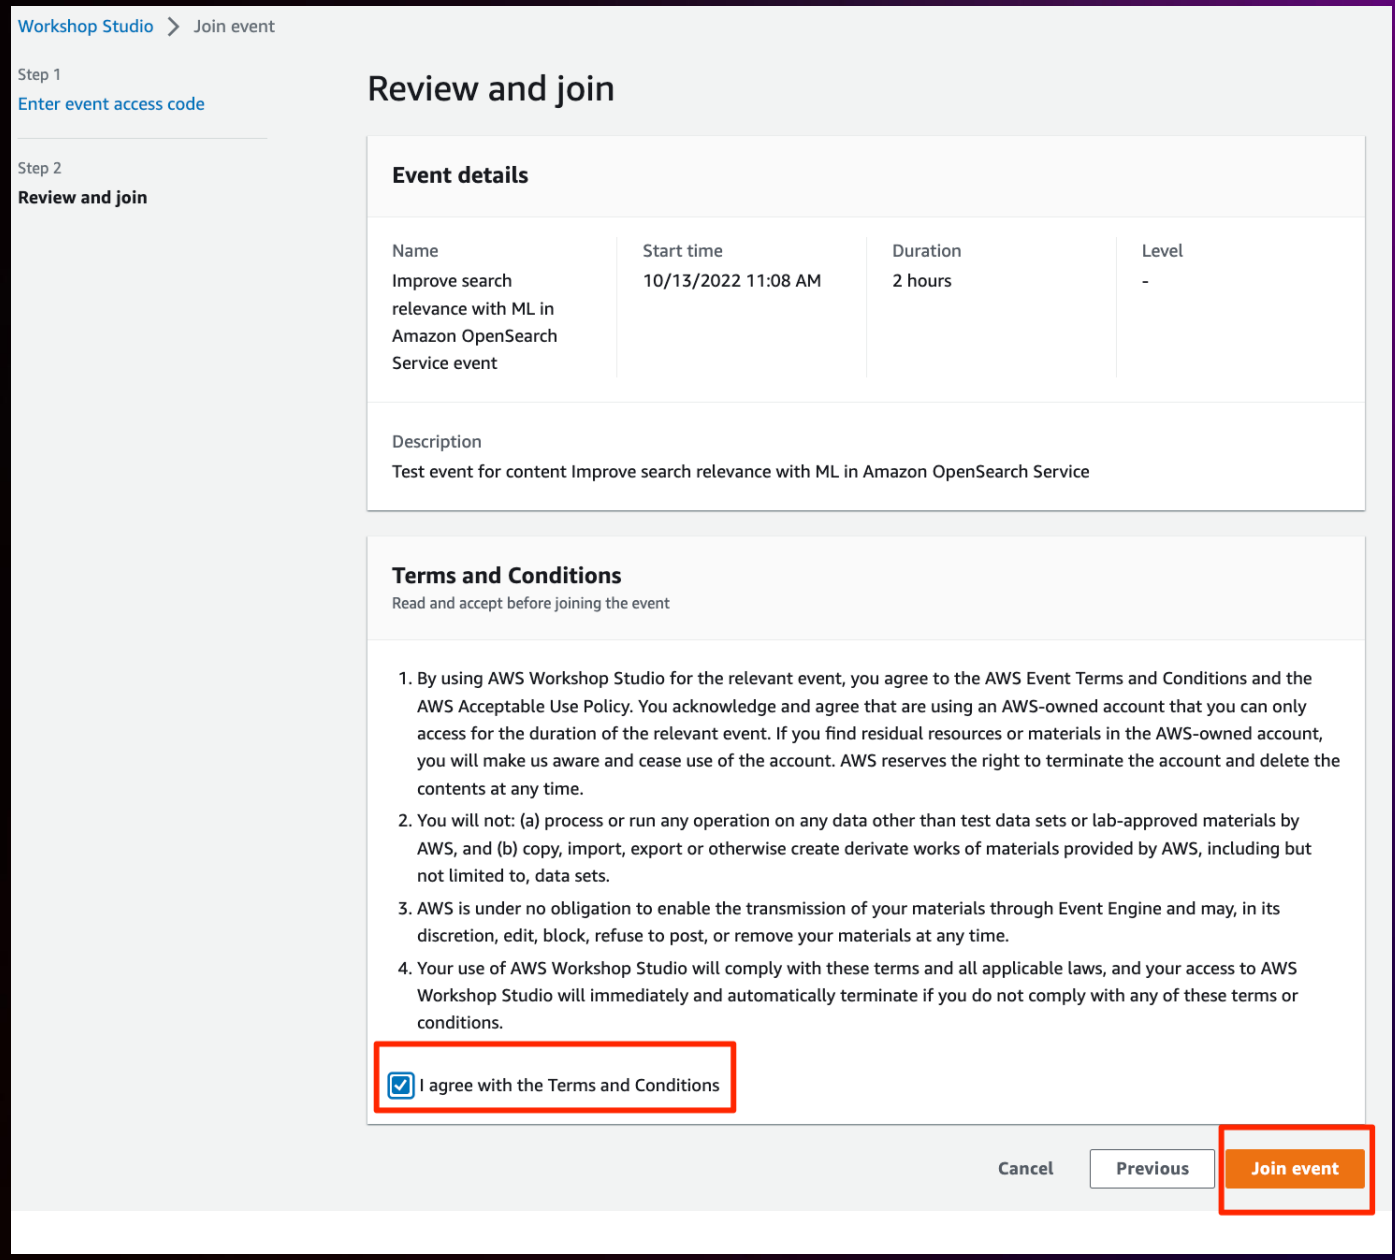

Step 4: Get started with the workshop

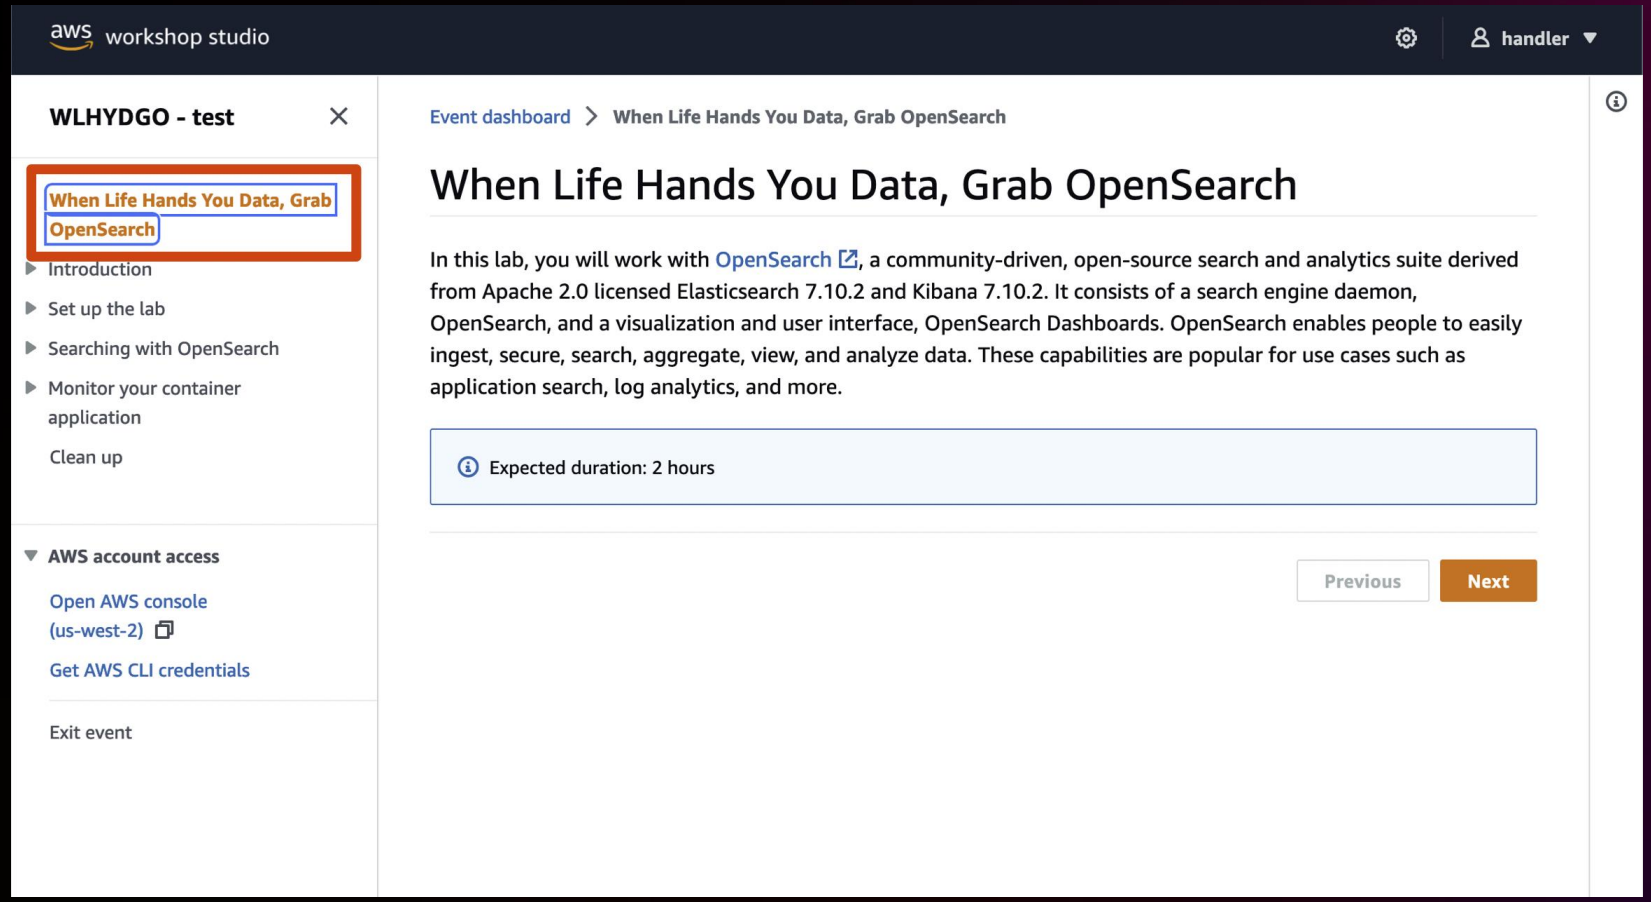

Access AWS account

Access the AWS Management Console or generate AWS CLI credentials as needed

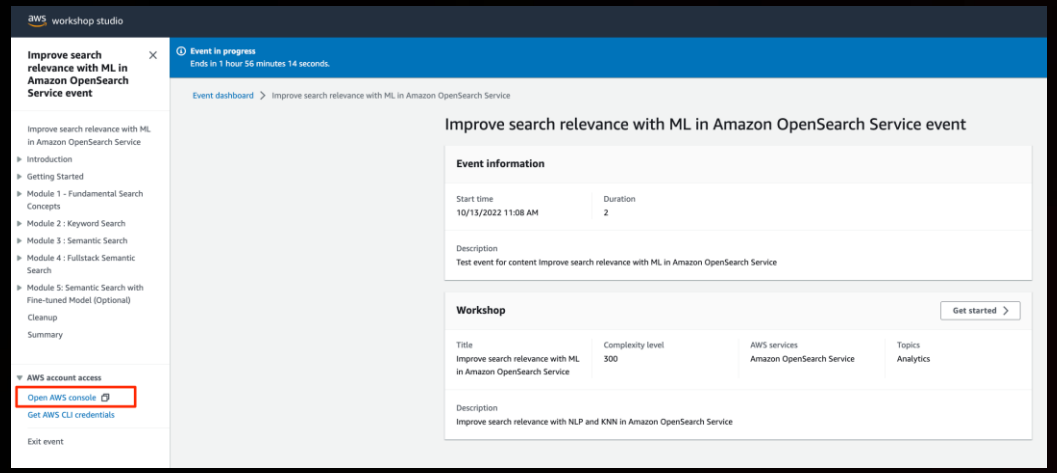

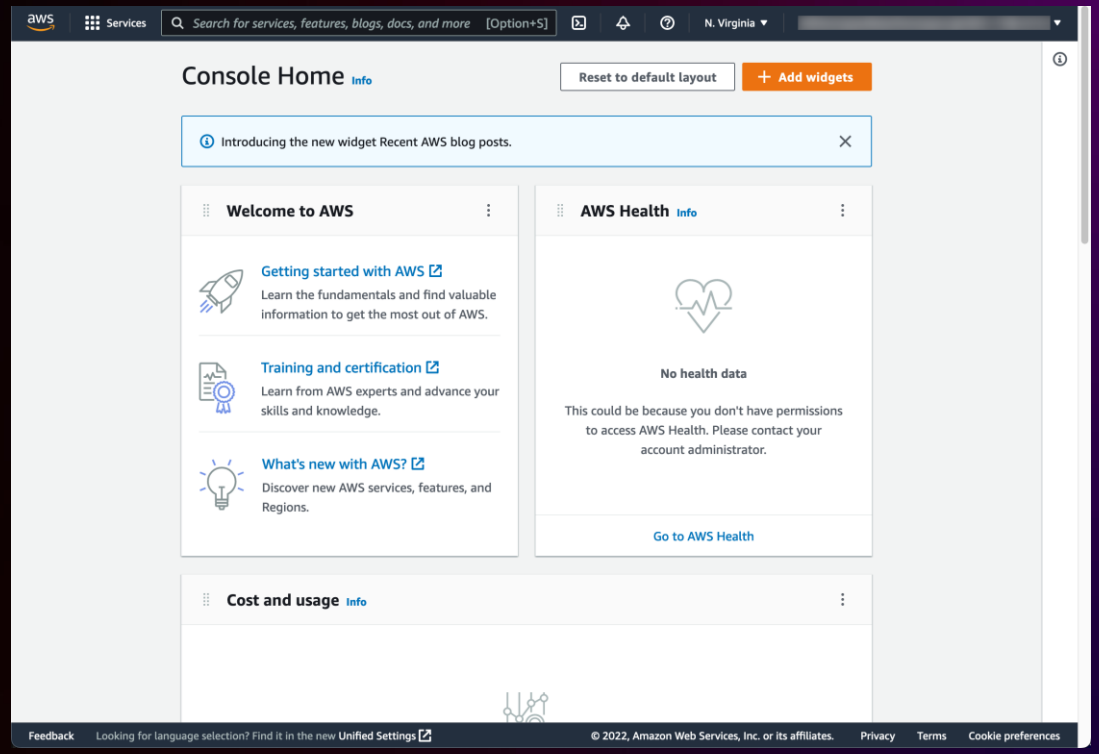

Get started

Event access for today

https://tinyurl.com/26f4j4cd

Event access code

767a-037d58-84

Workshop permanent link

https://tinyurl.com/2mwmx9n2

Thank you!

Please complete the session survey in the **mobile app**