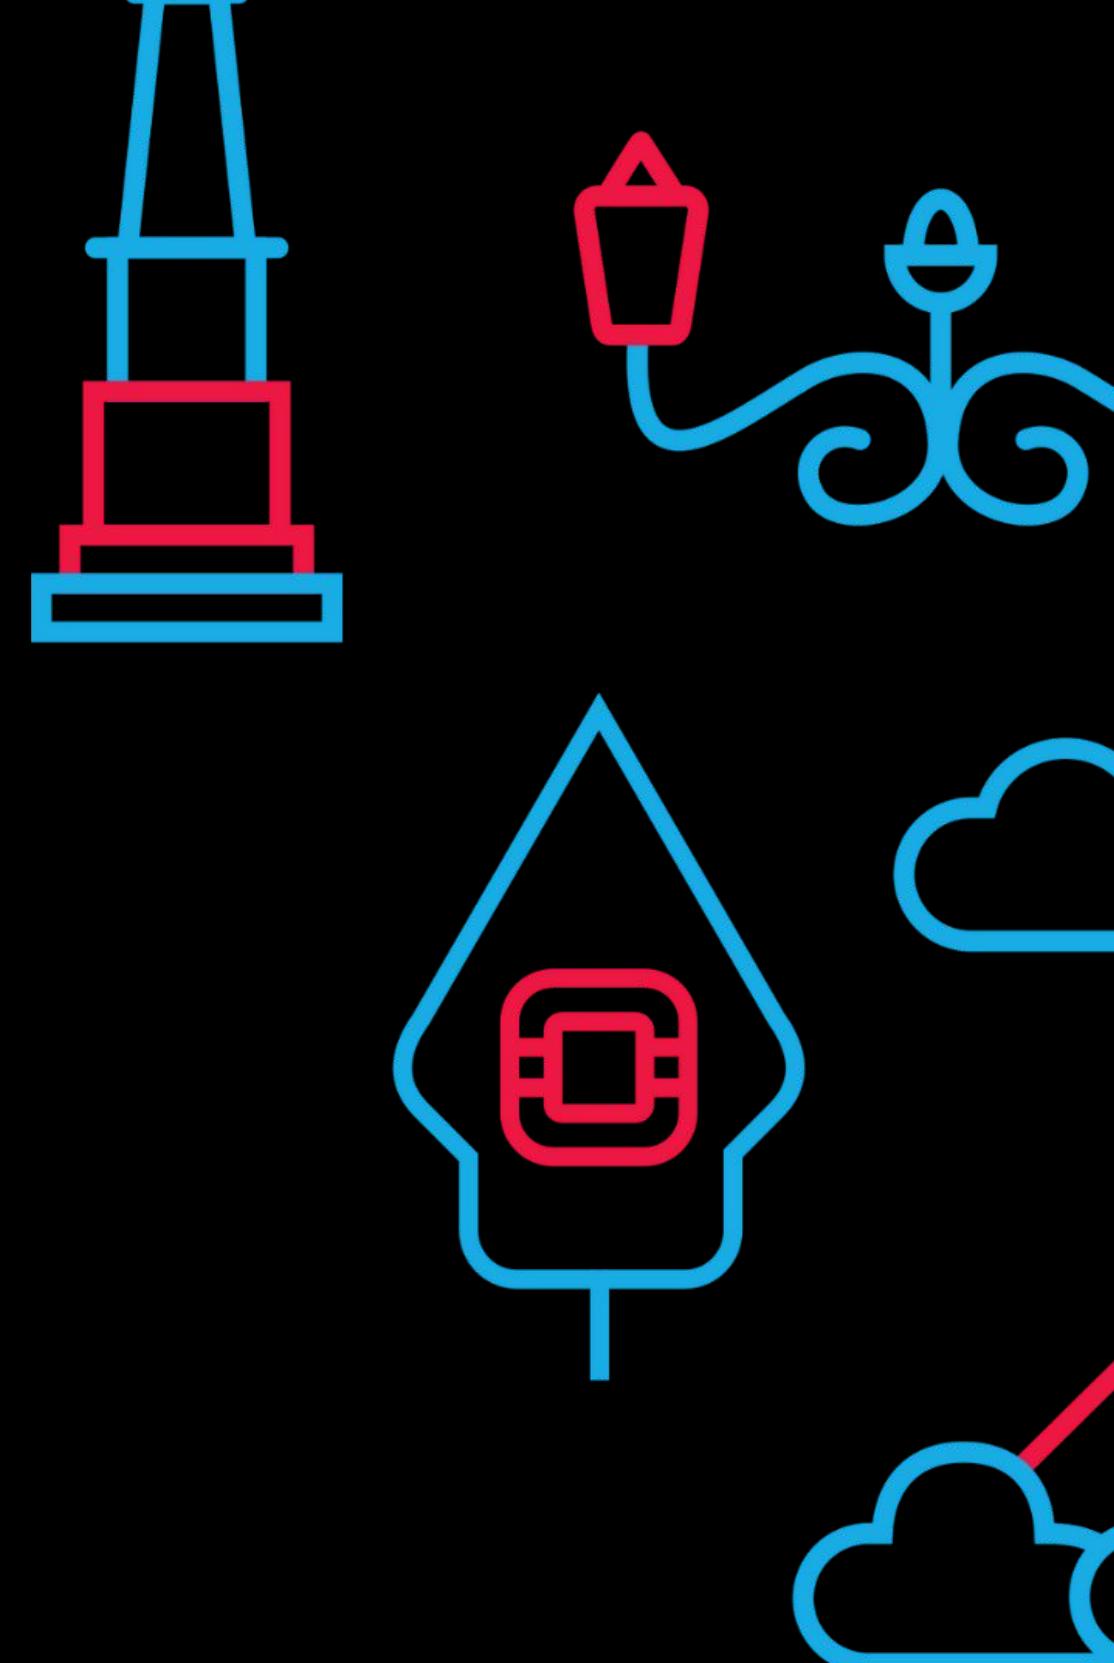


# Building Real-Time Incident Analytics with Apache Flink and Kafka on OpenStack

Ananda Dwi Rahmawati

*Cloud & DevOps Engineer*



Yogyakarta, 19 July 2025



# Ananda Dwi Rahmawati

- Cloud & DevOps Engineer, Singapore
- AWS Container Hero & Google Developer Expert Cloud - Modern Architecture
- Master of Computer Science - University of Texas at Austin
- Alumni of TRPL 2019, Universitas Gadjah Mada
- <https://linktr.ee/misskecupbung>

*"Empowering organizations with real-time analytics transforms incident management from reactive to proactive, unlocking faster decisions, deeper insights, and resilient operations."*

3

“ ”

# Agenda



- Key Takeaways
- Problem Statement
- Apache Flink
- Apache Kafka
- OpenStack
- Architecture Diagram + Data Flow

# Agenda



- Implementation
- The Challenges
- Use Cases
- Q&A

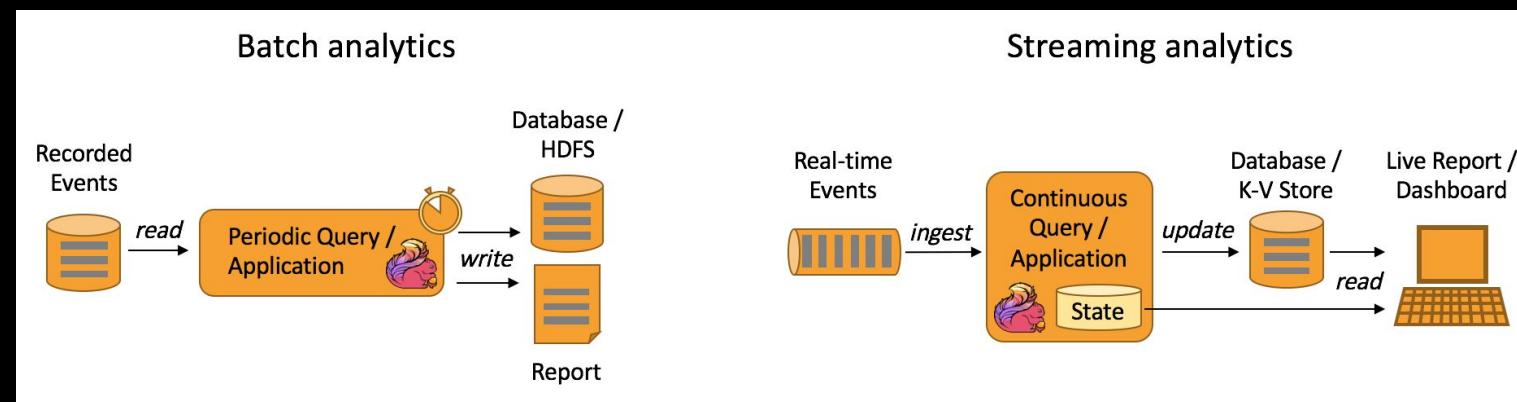
# Key Takeaways

- Real-time analytics is crucial for effective incident management
- Scalable, reliable, and fast data processing is essential for modern operations
- Enables proactive incident detection and mitigation
- Cloud-native technologies streamline operations and reduce costs
- Timely data insights support compliance and audit requirements

# Problem Statement

- Traditional batch processing is **slow** for incident response.
- **Growing** data volumes and **complexity**.
- Need for immediate insights and automated actions.
- Difficulty in correlating incidents across multiple data sources in real time.
- Limited **visibility** into ongoing incidents hampers rapid decision-making.

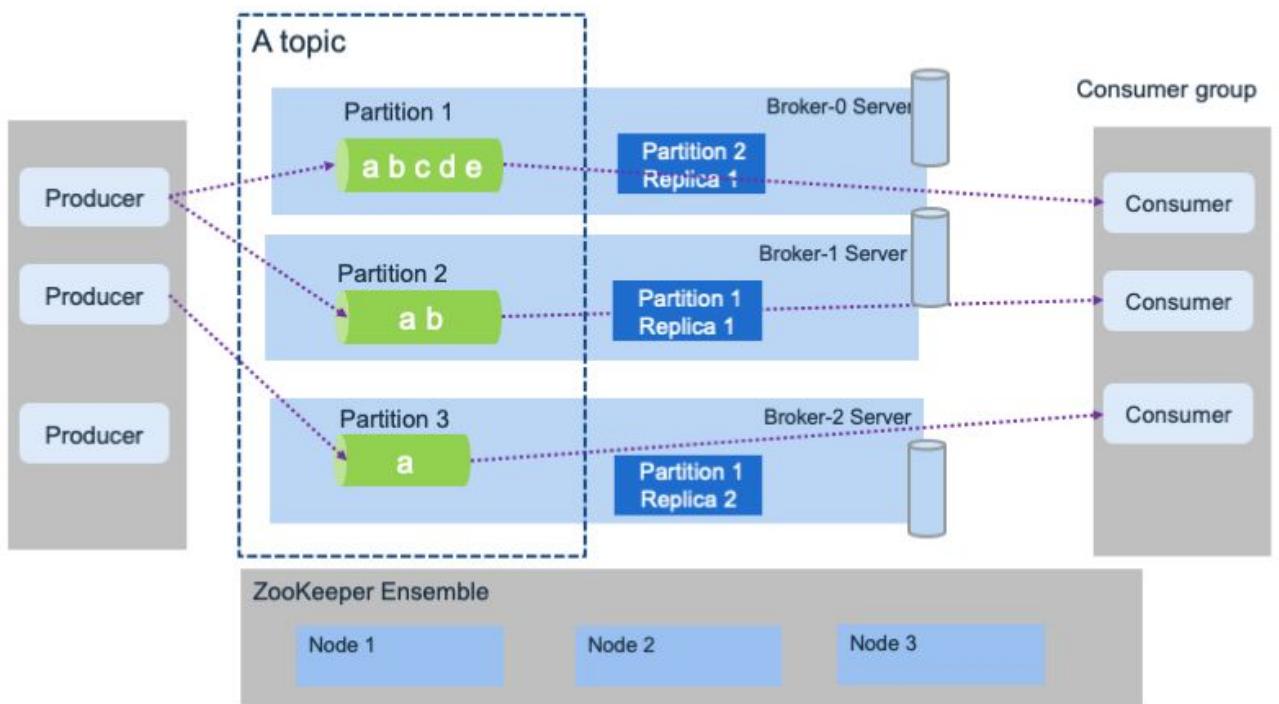
# Apache Flink



- Stream processing framework
- Low latency, high throughput
- Supports complex event processing
- Supports both stream and batch processing for flexible analytics.
- Provides advanced windowing and state management capabilities.
- Integrates easily with various data sources and sinks.

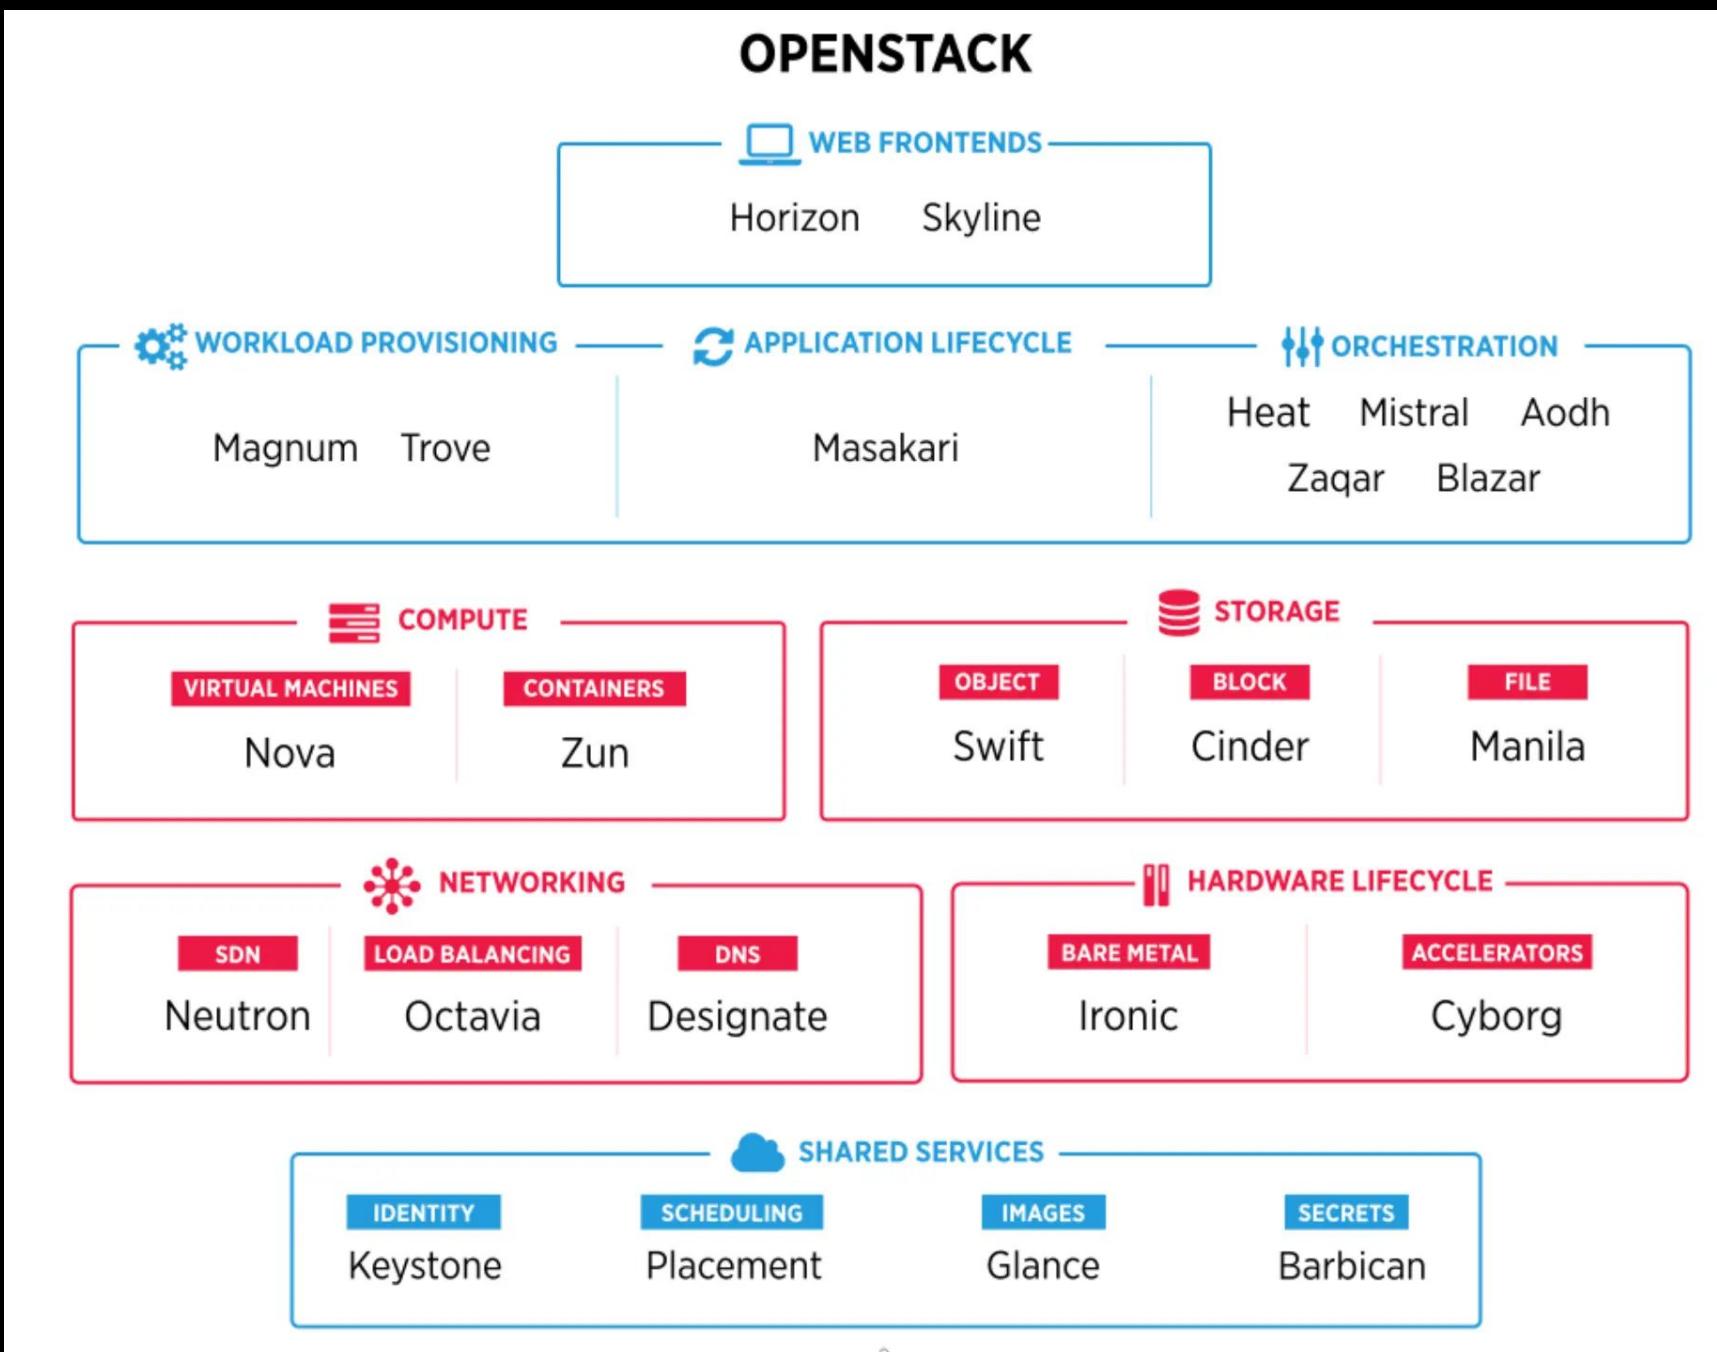
# Apache Kafka

## Kafka Architecture



- Distributed messaging system
- Handles high-velocity data streams
- Decouples data producers and consumers
- Guarantees message durability and fault tolerance
- Enables horizontal scaling for high-throughput workloads
- Offers strong ecosystem support for connectors and monitoring

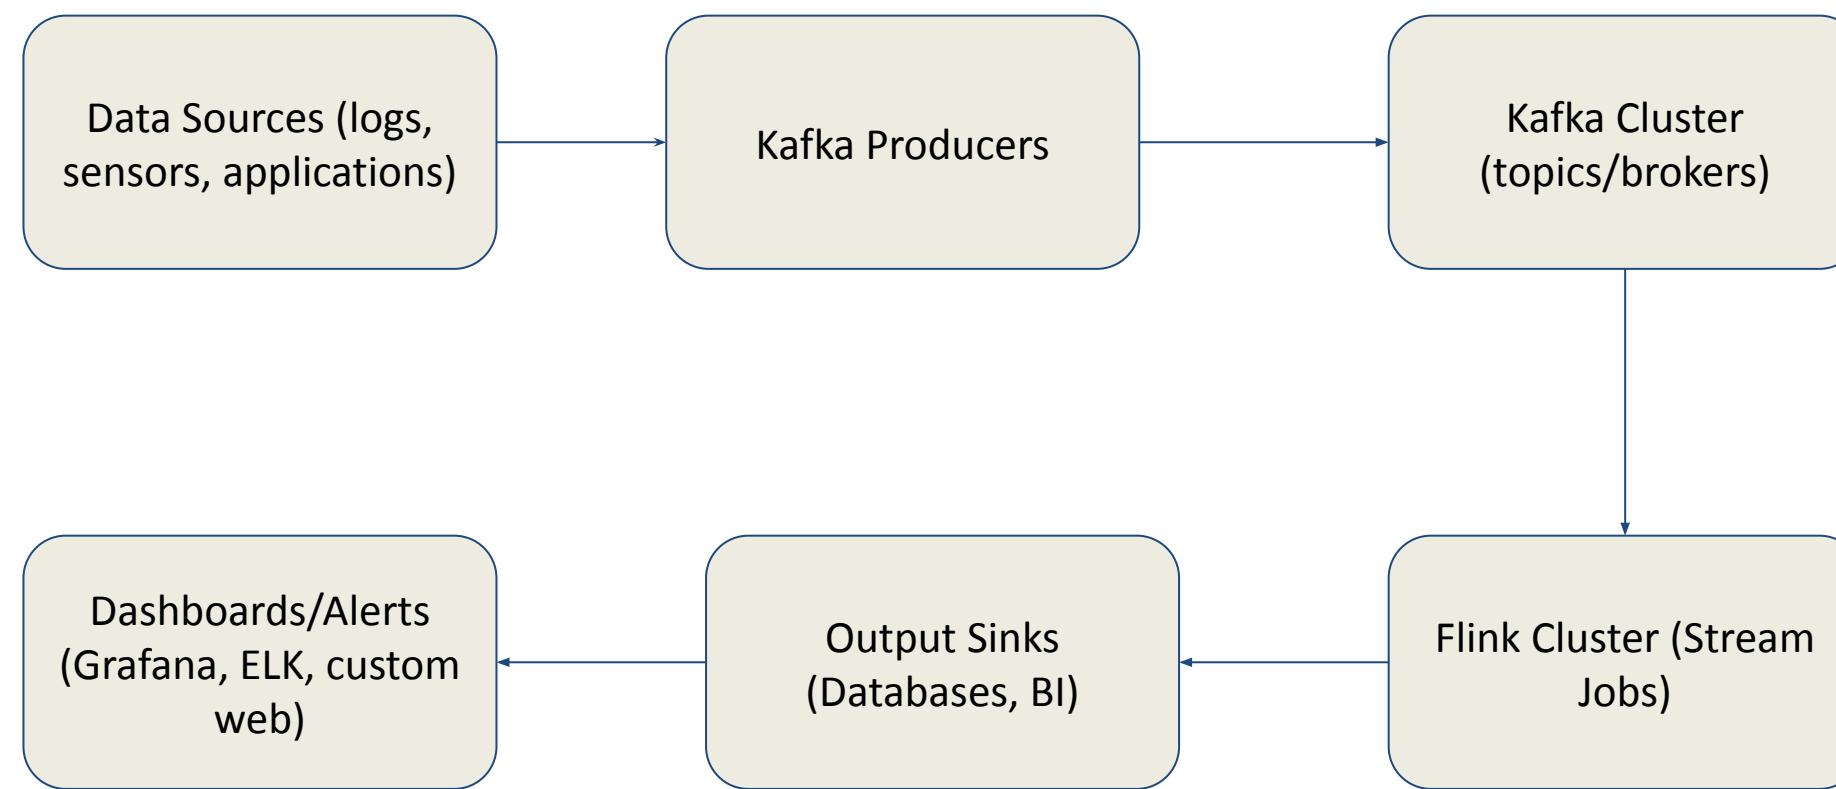
# OpenStack Overview



- Open-source cloud platform
- Provides compute, storage, and networking resources
- Enables dynamic scaling of analytics infrastructure
- Supports automation and orchestration for rapid resource provisioning.
- Enables multi-tenancy and isolation for secure analytics environments.



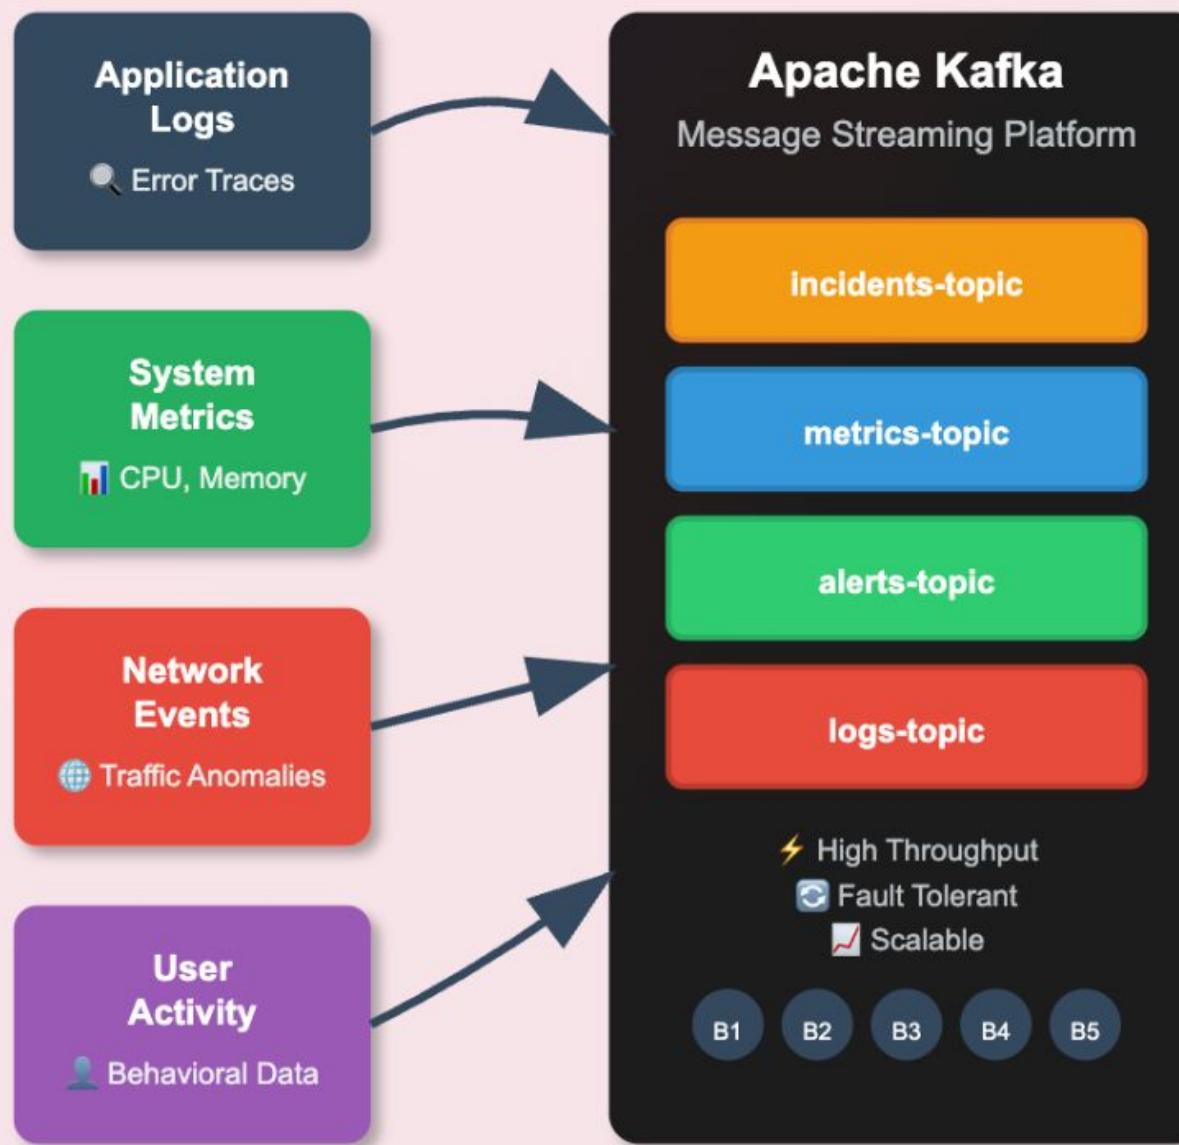
# Architecture Diagram



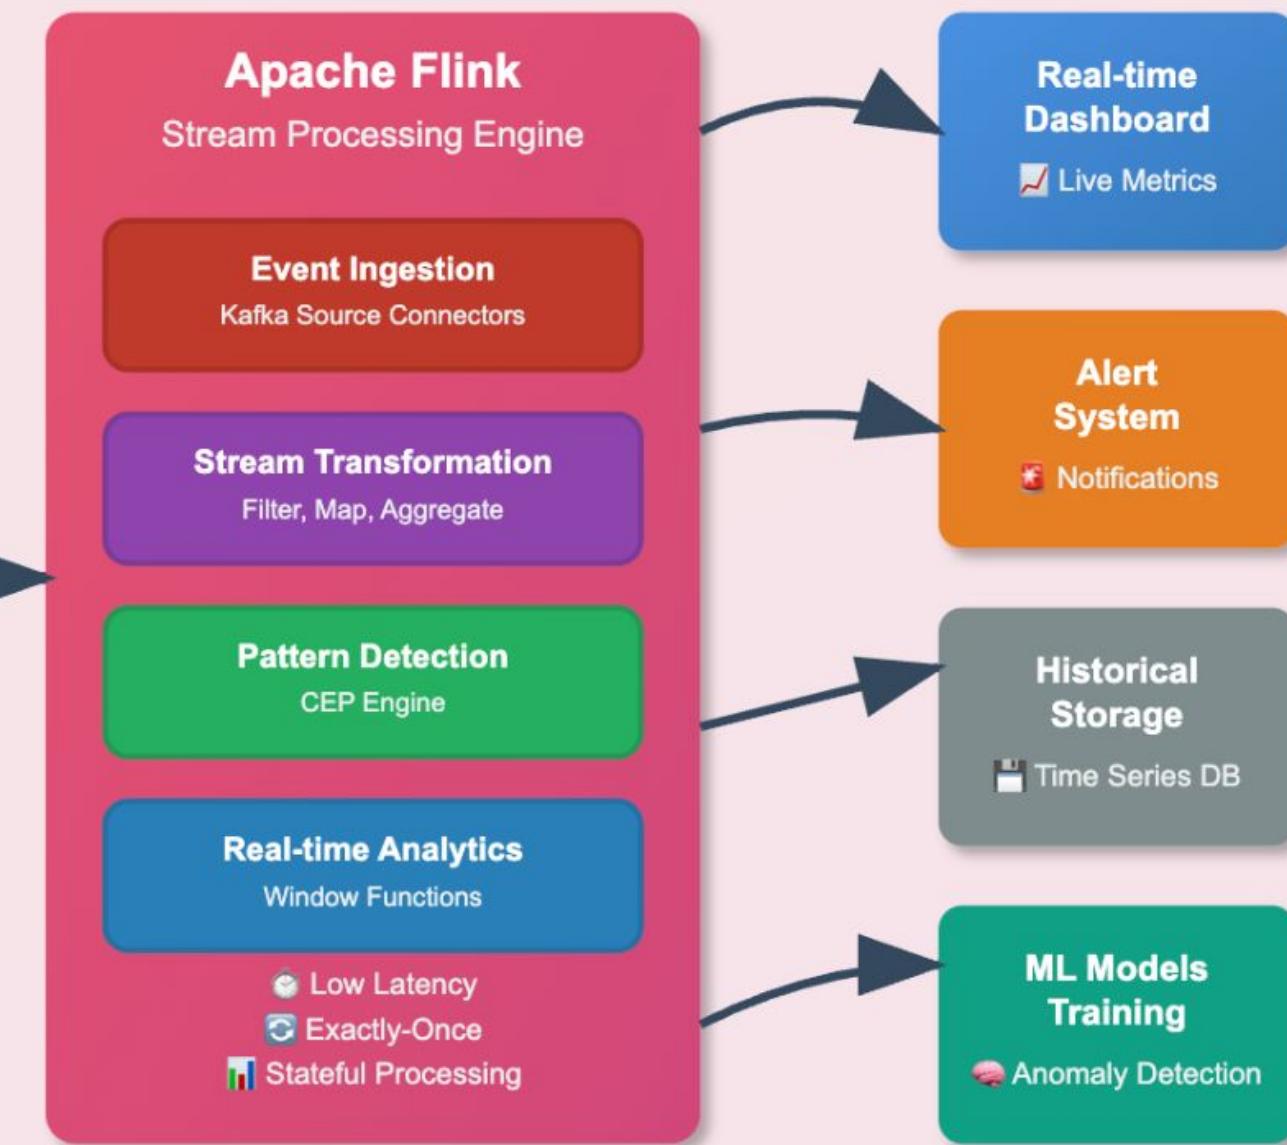
- **Incident Data Generation**
  - Various sources (applications, sensors, logs) generate incident data in real time.
- **Kafka Producers**
  - Producers collect and format incident data.
  - Data is published to specific Kafka topics.
- **Kafka Cluster**
  - Kafka brokers store and manage incoming data streams.
  - Ensures durability, ordering, and availability.
- **Flink Consumers**
  - Flink jobs subscribe to Kafka topics.
  - Flink processes data streams, applies analytics, filtering, and event correlation.
  - Supports windowing, stateful computations, and complex event processing.
- **Output Sinks**
  - Processed results are sent to:
  - Dashboards (e.g., Grafana, Kibana)
  - Alerting systems (e.g., email, SMS, Slack)
  - Databases for storage and further analysis
- **Feedback/Monitoring**
  - Monitoring tools track data flow, system health, and performance.

# OpenStack Infrastructure

## Data Sources



## Output Systems



## Key Performance Metrics

- Event Processing Rate: 1M+ events/sec
- End-to-End Latency: < 100ms
- Availability: 99.9% uptime
- Fault Tolerance: Auto-recovery
- Scalability: Horizontal scaling
- Data Retention: 7-day rolling window

## Architecture Benefits

- Real-time incident detection and response
- Predictive analytics for proactive monitoring
- Elastic scaling based on workload demands
- Complex event pattern matching
- Multi-tenant isolation and security
- Integration with existing OpenStack services

## Common Use Cases

- Security breach detection and response
- Application performance monitoring
- Infrastructure health monitoring
- User behavior analysis
- Compliance and audit trail generation
- Capacity planning and resource optimization

# Implementation Step(s)

- Set Up OpenStack Environment
- Deploy Kafka Cluster
  - Install Kafka on provisioned VMs.
  - Configure broker settings and start Kafka services.
  - Example (Kafka start):

## # On each Kafka node

```
wget https://downloads.apache.org/kafka/3.6.0/kafka_2.13-3.6.0.tgz
tar -xzf kafka_2.13-3.6.0.tgz
cd kafka_2.13-3.6.0
bin/zookeeper-server-start.sh config/zookeeper.properties &
bin/kafka-server-start.sh config/server.properties &
```

# Implementation Step(s)

- Deploy Flink Cluster
  - Install Flink on provisioned VMs.
  - Start Flink JobManager and TaskManager.
  - Example (Flink start):

```
wget https://archive.apache.org/dist/flink/flink-1.17.1/flink-1.17.1-bin-scala_2.12.tgz
tar -xzf flink-1.17.1-bin-scala_2.12.tgz
cd flink-1.17.1
./bin/start-cluster.sh
```

# Implementation Step(s)

- Develop Flink Jobs for Incident Analytics (Python)
  - Write a Flink job using **PyFlink** to consume from Kafka and process incidents.
  - Example (PyFlink job):

```
from pyflink.datastream import StreamExecutionEnvironment
from pyflink.datastream.connectors import FlinkKafkaConsumer
from pyflink.common.serialization import SimpleStringSchema

env = StreamExecutionEnvironment.get_execution_environment()
properties = {'bootstrap.servers': 'localhost:9092', 'group.id': 'incident-group'}
consumer = FlinkKafkaConsumer(
    topics='incident-topic',
    deserialization_schema=SimpleStringSchema(),
    properties=properties
)
stream = env.add_source(consumer)
alerts = stream.filter(lambda data: 'CRITICAL' in data)
alerts.print()
env.execute('Incident Analytics Job')
```

# The Challenges (the Next Step!)

- Managing state and fault tolerance in Flink
- Ensuring data consistency in Kafka
- Resource allocation and scaling on OpenStack
- Security and access control
- Integrating with legacy systems and diverse data sources.
- Monitoring and troubleshooting distributed components.
- Balancing performance with cost efficiency.

## Whole Cluster

Maximum Cluster CPU Usage

**5.34**

Maximum EP + Flink + EEM CPU Usage

**0.101**

Maximum EP + Flink CPU Usage

**0.0974**

Maximum CPU Usage

**0.0752**

Maximum Flink CPU Usage

**0.0282**

Maximum SP CPU Usage

**0.00601**

Maximum EEM CPU Usage

**0.293**

Maximum Cluster Memory Usage

**50.2 GiB**

Maximum EP + Flink + EEM Memory Usage

**3.39 GiB**

Maximum EP + Flink Memory Usage

**2.97 GiB**

Maximum CPU Usage

**2.21 GiB**

Maximum Flink Memory Usage

**781 MiB**

Maximum SP Memory Usage

**420 MiB**

Maximum EEM Memory Usage

**3.72 GiB**

## Flink Task Managers

Number of Task Managers

**1**

Taskslots total

**10**

Number of Running Jobs

**10**

Taskslots Available

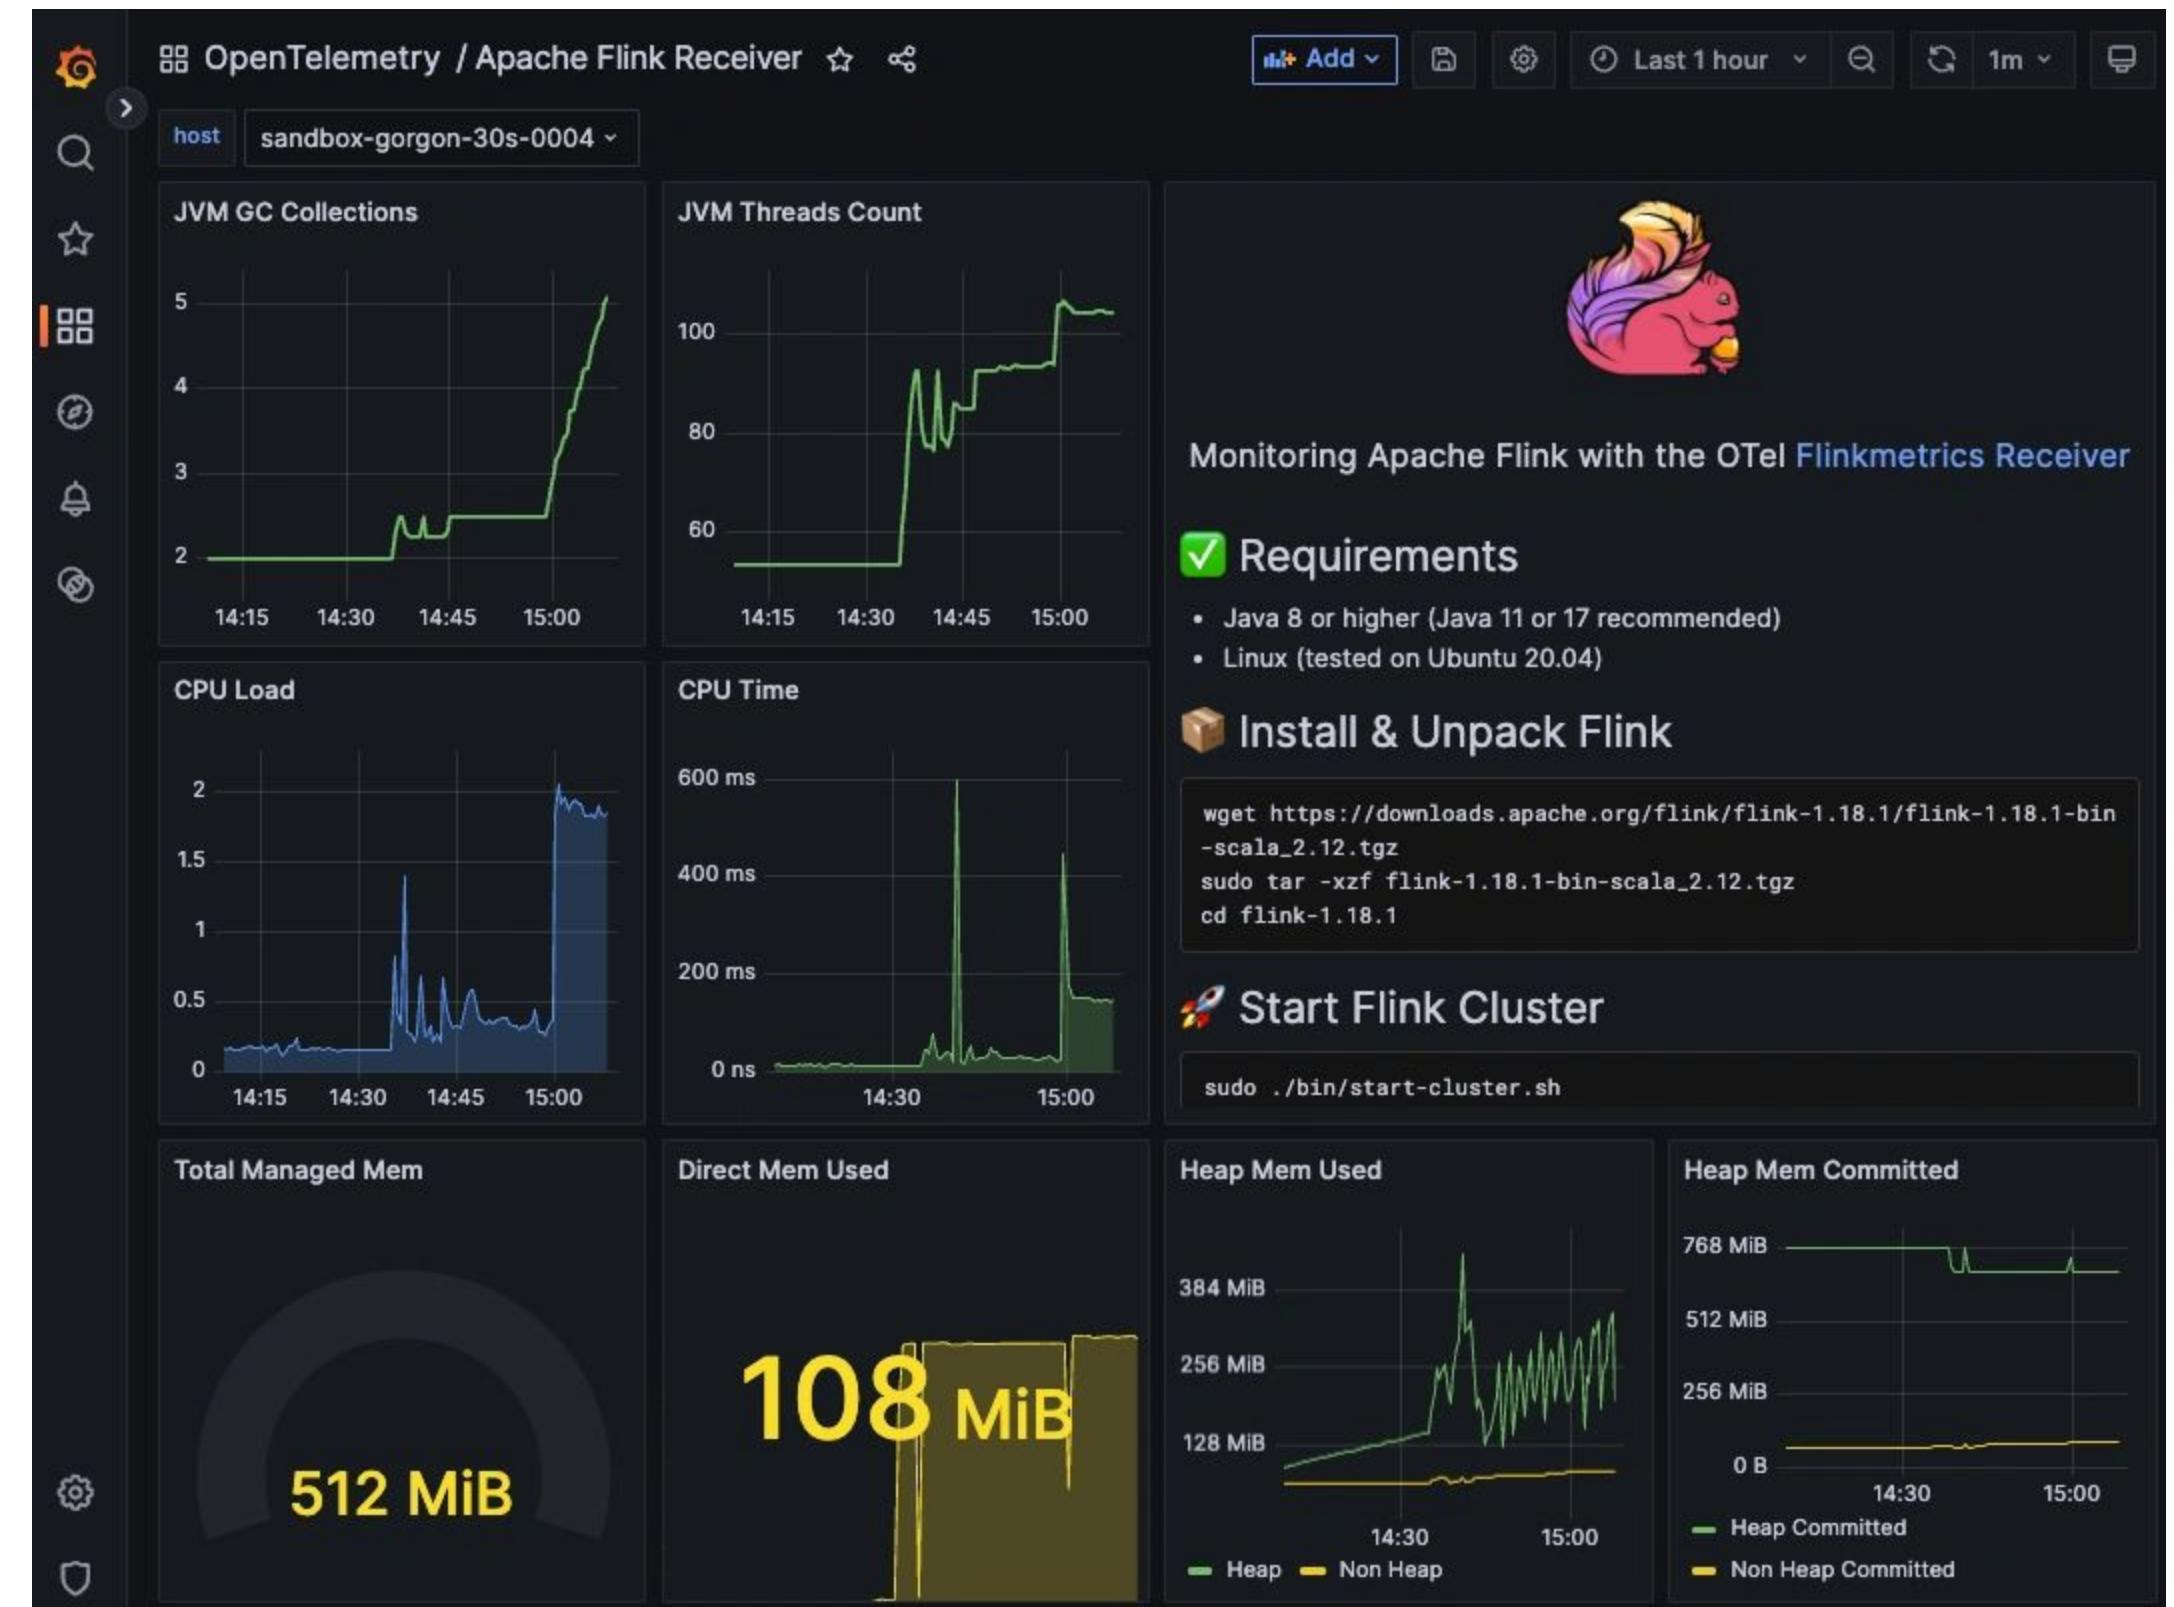
**0**

Number records in per second



Number records out per second

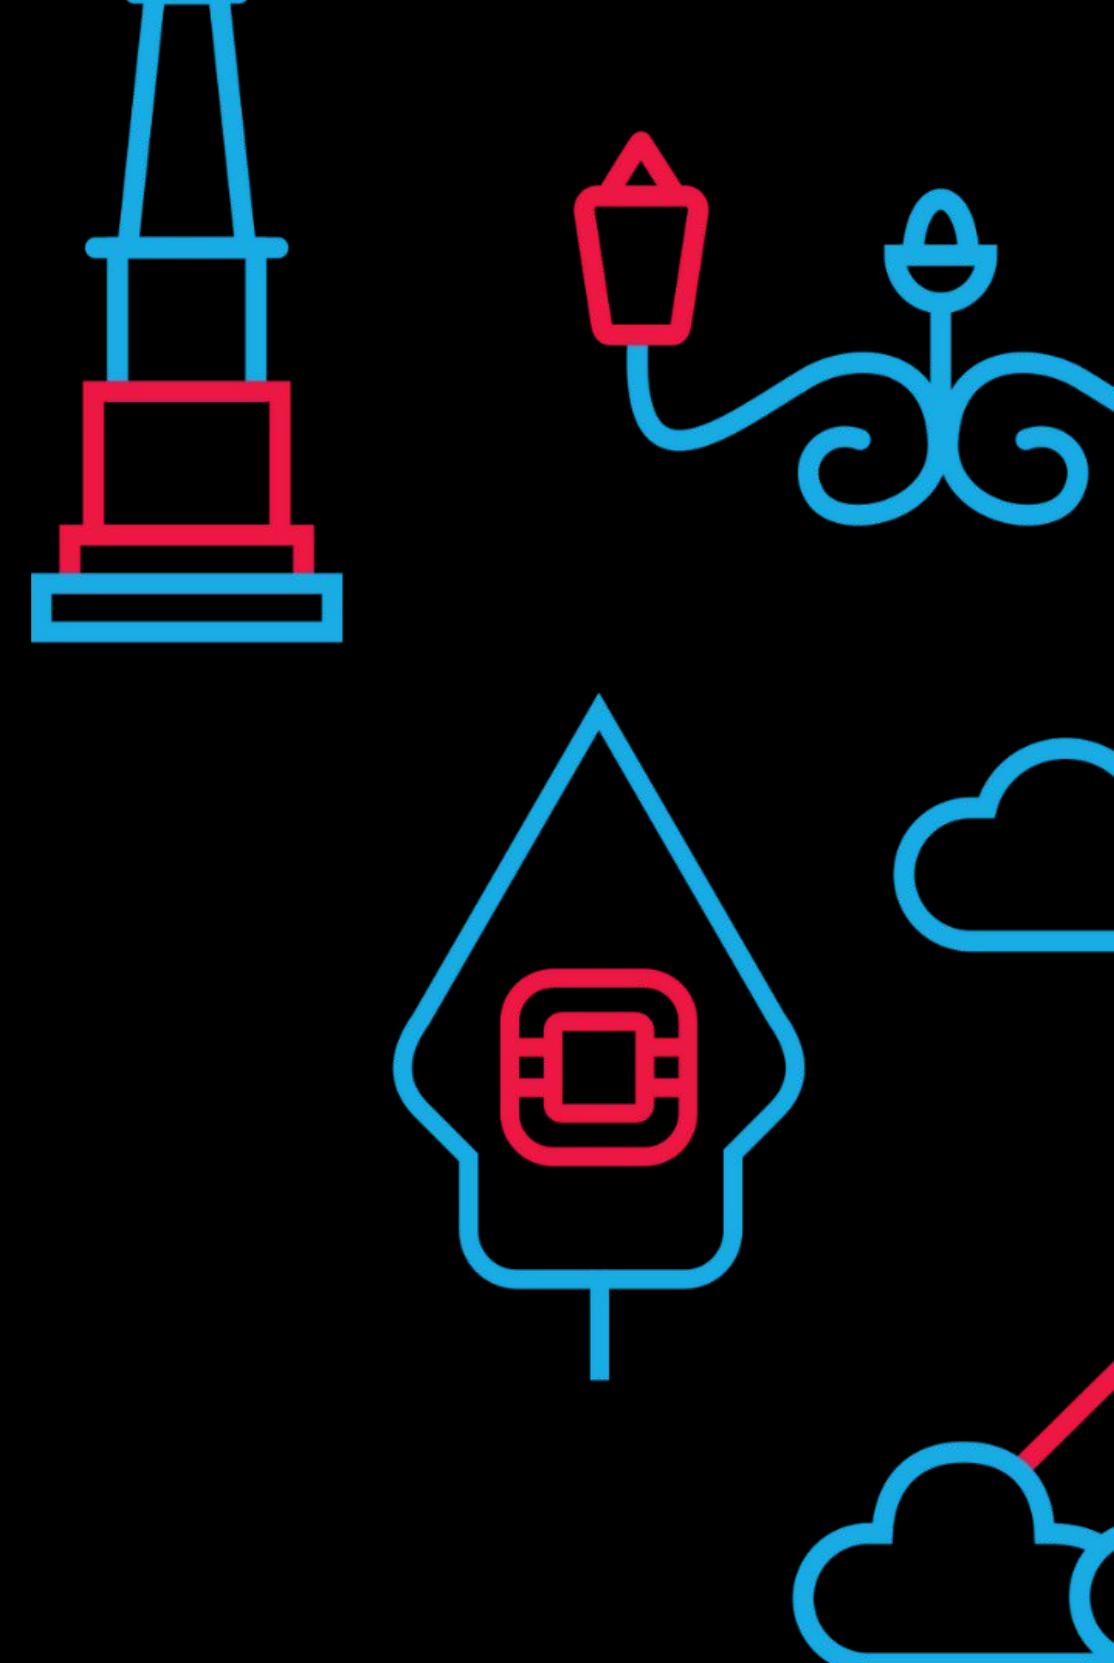






# QnA

(Answers not Guarantee)



Yogyakarta, 19 July 2025

# THANK YOU



Yogyakarta, 19 July 2025