

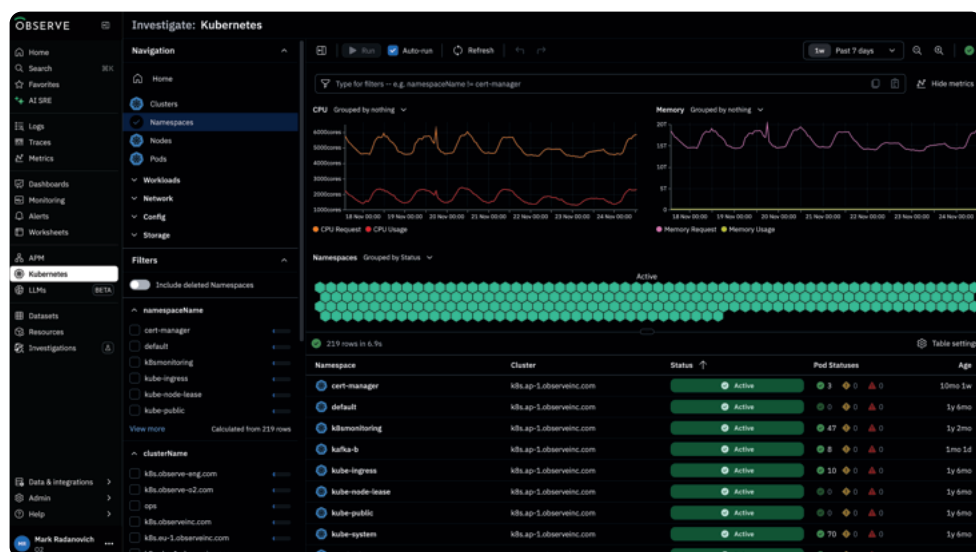
# Observability Engineered for Scale

AI-powered observability for 10x faster troubleshooting at 60% lower cost

## Scalable Observability for Logs, Metrics, and Traces

Modern systems generate more telemetry than legacy observability platforms can cost-effectively ingest. Observe solves the challenges of scale through a data lake architecture, which stores telemetry at dramatically lower cost, and a knowledge graph that tracks relationships across services and infrastructure. Pairing agentic workflows with context from the knowledge graph, Observe's chat-based AI SRE accelerates root cause analysis and resolution.

### TRUSTED BY



Log Management  
& Analytics



Infrastructure  
Monitoring



Application Performance  
Monitoring



LLM Observability

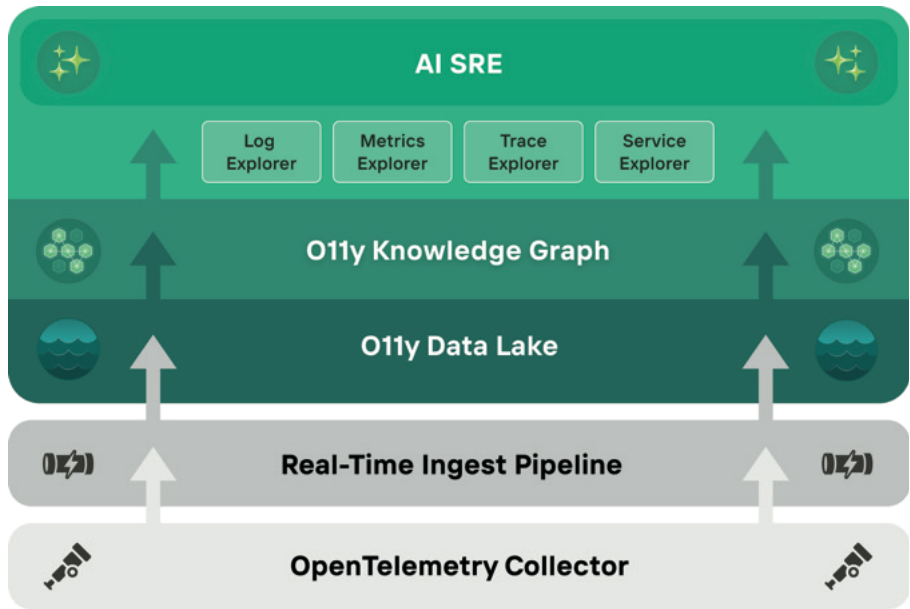


During a period of explosive growth at Tekion, we realized our existing observability tools weren't going to scale with us. We had tried major commercial and open-source tools, but both resulted in escalating costs and constant tuning efforts that drained engineering resources. Observe gave us a cost-effective unified platform for logs, metrics, and traces, with the ability to correlate across all of them. That visibility has been critical to maintaining service reliability while scaling aggressively."

**Binu Mathew**  
CTO at Tekion

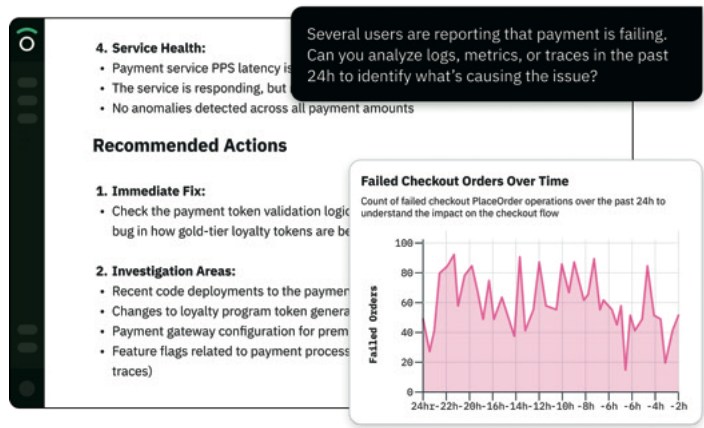
# Modern Observability Architecture

Observe’s cloud-native architecture maintains context across logs, metrics, and traces with a unified data lake and knowledge graph. Built on open standards like OpenTelemetry and Apache Iceberg, it ensures interoperability and avoids lock-in.



## Observe AI SRE

The AI SRE applies context, identifies root causes, and suggests fixes, enabling teams to reduce toil and troubleshoot 10x faster.



**Pinpoint failures and fix issues faster** using contextual understanding of logs, metrics, and traces

**Cut observability costs** with a single low-cost data lake backend

**Operate at enterprise scale** with built-in data privacy, compliance, and customizable agentic workflows via MCP Server

# Platform Features

## AI SRE

Automates investigative heavy lifting to accelerate incident response and deliver precise root cause analysis.

## O11y Knowledge Graph

Structures and correlates telemetry for fast, context-aware AI and analytics.

## O11y Data Lake

Keeps telemetry searchable longer at 60% lower cost. Supports Apache Iceberg for data sharing.

## Explorers

Interfaces for common data analysis use cases. Includes Log, Metrics, Trace, Service, and K8s Explorers.

## Elastic Compute

Scales compute as needed, without runtime interference from other workloads.