


Guidance for Kelvin AI on AWS

This architecture diagram illustrates how to deploy Kelvin AI on AWS. The solution guidance demonstrates how you can optimize industrial production operations with real-time closed loop control using Kelvin AI Autonomous Operations Platform on AWS.



1 Time-series data from SCADA (Supervisory Control and Data Acquisition) or Programmable Logic Controllers (PLCs) leveraging MQTT, OPC-UA, Modbus or custom protocols are integrated with the Kelvin Node via the Kelvin Connector. Additionally, the bridge supports multi-model video and object (images, files, etc.) data streams for real-time use cases.

2 Kelvin edge nodes, deployed on **AWS Outpost Server** or customer provisioned gateway devices consists of a streaming Data Broker for managing data communication, a Node Historian for local storage and offline operations, a Control Manager for open and closed loop control actions for SmartApps.

By running Kelvin edge nodes on **Amazon EKS on AWS Outposts**, the solution delivers a consistent Kubernetes environment from the cloud to the edge, while unifying management, security and compliance and reducing operational overhead between on-premises and cloud environments.

3 The Node Historian securely synchronizes data between the node and the Kelvin cluster over an encrypted SSL channel. In the event of connectivity loss, the edge node continues to operate fully offline, with data automatically resynchronizing once connectivity is restored as First-In First-Out (FIFO) or Last-In Last-Out (LIFO).

Running Kelvin node at the edge provides real-time data ingestion, local processing and autonomous closed loop control with minimal latency and continuous operation even during intermittent connectivity. Running Kelvin cluster in the cloud enables ML training and long-horizon analytics, centralized orchestration, long-term data storage and enterprise-grade security and scalability.




Reviewed for technical August 18, 2025

© 2025, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Reference Architecture

Guidance for Kelvin AI on AWS

This architecture diagram illustrates how to deploy Kelvin AI on AWS. The solution guidance demonstrates how you can optimize industrial production operations with real-time closed loop control using Kelvin AI Autonomous Operations Platform on AWS.



- 4 The Kelvin Asset Registry stores asset information, user profiles in **Amazon Relational Database Service** and **Amazon Simple Storage Service**. The cloud historian is backed by a Timeseries database (InfluxDB) on **Amazon Elastic Block Store**.
- 5 The Workload Orchestrator deploys and manages Kelvin SmartApps that are built by Kelvin's partner ecosystem or customers. SmartApps observe real-time industrial data and utilize physics, first principle and ML models to process time-series data, detect events, visualize asset performance, and make predictions / recommendations.
- 6 Edge data is contextualized and modeled for continuous training of ML models. Data is exported to **Amazon S3** and **AWS IoT SiteWise**. Customers can train physics based, first principle and AI/ML models are built in **Amazon SageMaker** utilizing the asset and time series data for use cases such as containing monitoring and predictive maintenance.
- 7 Kelvin SmartApps communicate with Enterprise IT systems (ERP, CMMS) via Integration API to update KPIs and take actions in real time. Data can also be visualized in **Amazon QuickSight** for business intelligence.
- 8 Kelvin SmartApps running in the cloud/edge provide recommendations for closed or open loop control changes with intelligent control guardrails considering the physical limitations of the equipment
- 9 Control changes are processed and sent to Industrial Gateway/PLC or SCADA systems leveraging OPC-UA, Modbus or custom protocols supported by the Kelvin Connector



Reviewed for technical August 18, 2025

© 2025, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Reference Architecture