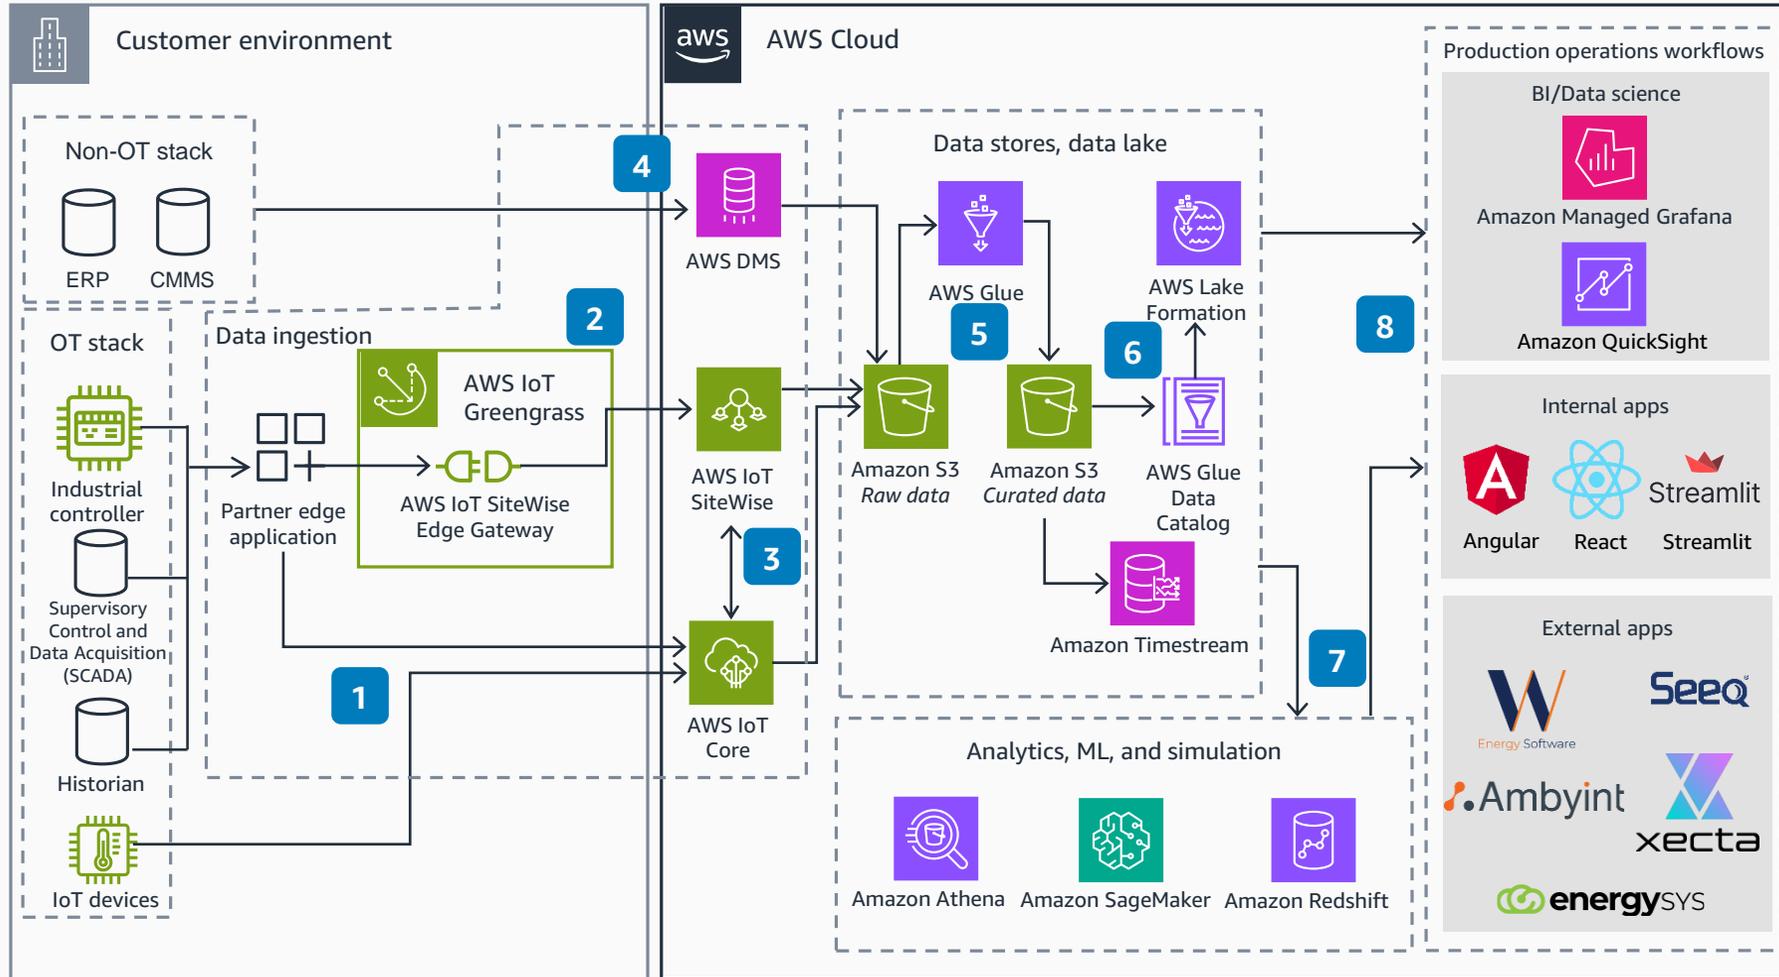# Guidance for Production Monitoring & Surveillance on AWS

This architecture shows the various components that allow operational data to flow from industrial sources into AWS and partner applications, helping oil and gas customers to monitor, forecast, and increase production.

Customer environment

Non-OT stack
- ERP
- CMMS

OT stack
- Industrial controller
- Supervisory Control and Data Acquisition (SCADA)
- Historian
- IoT devices

Data ingestion
- Partner edge application
- AWS IoT Greengrass
- AWS IoT SiteWise Edge Gateway

AWS Cloud

- AWS DMS
- AWS IoT SiteWise
- AWS IoT Core

Data stores, data lake
- AWS Glue
- Amazon S3 *Raw data*
- Amazon S3 *Curated data*
- AWS Lake Formation
- AWS Glue Data Catalog
- Amazon Timestream

Analytics, ML, and simulation
- Amazon Athena
- Amazon SageMaker
- Amazon Redshift

Production operations workflows

BI/Data science
- Amazon Managed Grafana
- Amazon QuickSight

Internal apps
- Angular
- React
- Streamlit

External apps
- W Energy Software
- Seeq
- Ambyint
- Xecta
- energySYS

---

1 Use partner applications, such as Embassy of Things, TwinTalk, and Element Unify to consume historical and real-time telemetry data. You can also use asset metadata from industrial sources using industrial protocols, such as Modbus, and OPC. Proprietary formats, such as AVEVA PI, can be integrated as well. Connect devices enabled as Internet of Things (IoT) to **AWS IoT Core** through secure sessions in which the devices are authenticated and data is encrypted in transit with x.509 certificates.

2 Deploy **AWS IoT Greengrass** if edge processing is desired for data pre-processing, batching, or to run code at the oil field. Install **AWS IoT Greengrass** in an industrial computer on site.

3 Process incoming messages from the field for unmodeled data, or from partner applications for modeled data with **AWS IoT Core** and **AWS IoT SiteWise**.

4 Enrich data from operational sources with records from non-operational technology (OT) systems. For example, computerized maintenance management system (CMMS), or enterprise resource planning (ERP) for production records. Extract these systems with **AWS Database Migration Service** (AWS DMS).

5 Ingest data from OT and non-OT sources into **Amazon Simple Storage Service** (Amazon S3), where **AWS Glue** can further process data (such as compression, aggregation, and calculated records). Ingest results into a modern data architecture composed by **Amazon S3** as the data lake. These results can be loaded to **Amazon Timestream** if fast querying of time-series data is required.

6 Load metadata from the data lake to enable analytics with AWS Glue Data Catalog. Similarly, **AWS Lake Formation** allows fine-grained governance and access control for the data lake.

7 Run specialized queries with **Amazon Athena** to calculate production aggregates, and load them into **Amazon Redshift**. Or use machine learning (ML) to forecast future production with **Amazon SageMaker**.

8 Leverage production data and results from analytics and ML with business intelligence (BI) internal applications, such as **Amazon Managed Grafana** or **Amazon QuickSight**. Leverage external applications such as Seeq for analytics, or EnergySys for hydrocarbon accounting, and either Ambyint or Xecta for production optimization.

AWS Reference Architecture