Architecture diagrams for manufacturing

Manufacturing on AWS

Manufacturing on AWS: Architecture diagrams for manufacturing

Table of Contents

Manufacturing on AWS

Publication date: **October 28, 2020 (Diagram history)**

With this architectural blueprint, you can see how AWS Cloud enables digital transformation for manufacturers. This architecture highlights how a smart factory, smart products, enterprise applications, and engineering workloads integrate through a central data lake.

The following diagrams describe the architecture:

- *Overview* — How manufacturing workloads connect through a central data lake.
- *Data ingestion* — Bringing data from factory floors and enterprise applications into AWS.
- *Data lake and analytics* — Processing, analyzing, and deriving insights from manufacturing data.
- *Application hosting* — Running enterprise and engineering applications on AWS.
- *Smart products* — Connecting manufactured products and machines to AWS.

Further reading

For additional information, see the following resources:

- AWS Architecture Icons
- AWS Architecture Center
- AWS Well-Architected
- Manufacturing Data Lake on AWS
- AWS Industrial IoT Predictive Maintenance ML Model
- AWS Industrial IoT Predictive Quality
- AWS Industrial IoT Asset Condition Monitoring

Diagram history

To be notified about updates to this reference architecture diagram, subscribe to the RSS feed.

Change	Description	Date

Initial publication	Reference architecture diagram first published.	October 28, 2020

RSS subscription

To subscribe to RSS updates, you must have an RSS plugin enabled for the browser that you are using.

Overview

The overview diagram shows how manufacturing workloads connect through a central data lake on AWS.

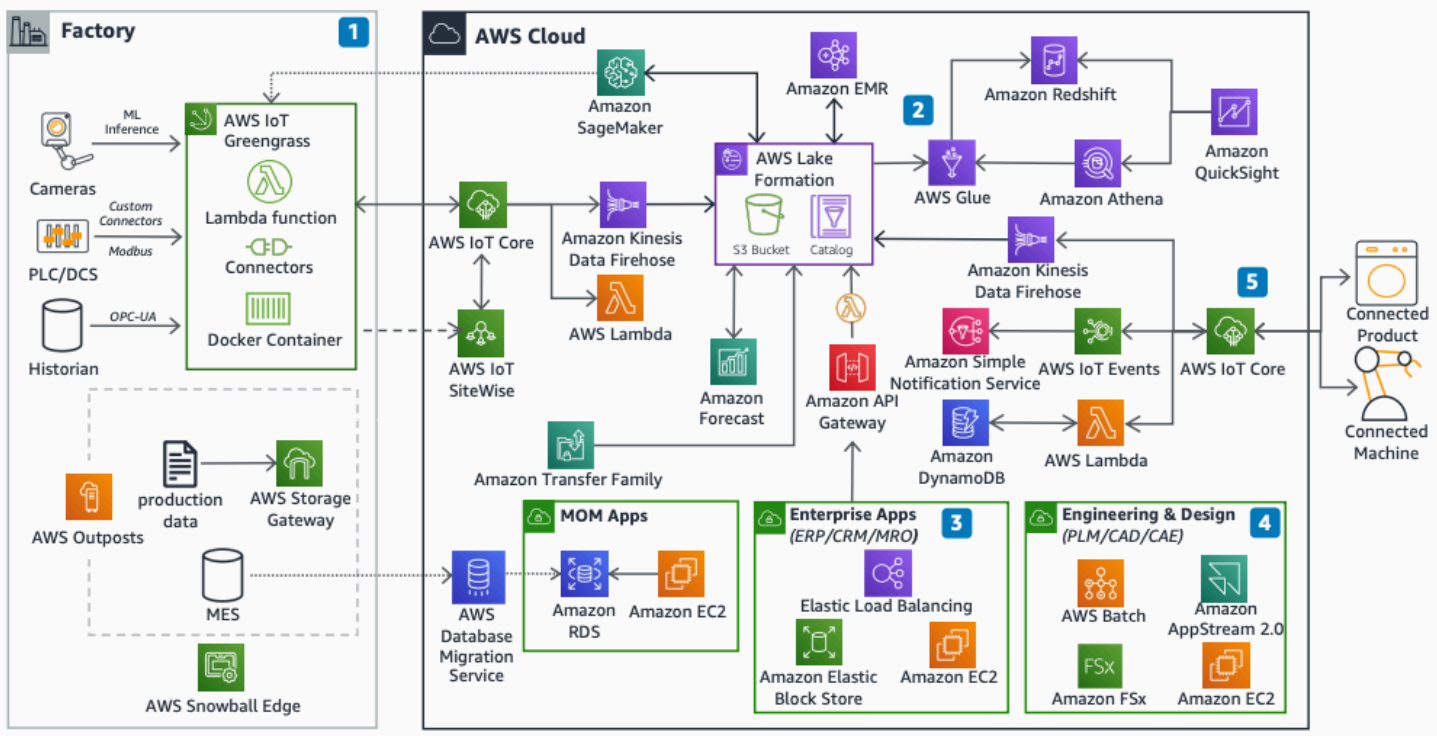

1. Central to the architecture is a manufacturing data lake for analytics and machine learning (ML). Use AWS Lake Formation or Amazon Simple Storage Service (Amazon S3) to structure your data lake.
2. Establish smart factories by connecting Industrial Internet of Things (IIoT) devices to the cloud with hybrid infrastructure.
3. Host enterprise applications with cost-effective, resilient AWS architecture.
4. Use Amazon EC2 Spot Instances and GPU instances for computer-aided design (CAD) and computer-aided engineering (CAE) workloads.
5. Build smart products for connected products and machines that send telemetry to the cloud.

Data ingestion

The data ingestion diagram shows how to bring data from factory floors and enterprise applications into AWS.

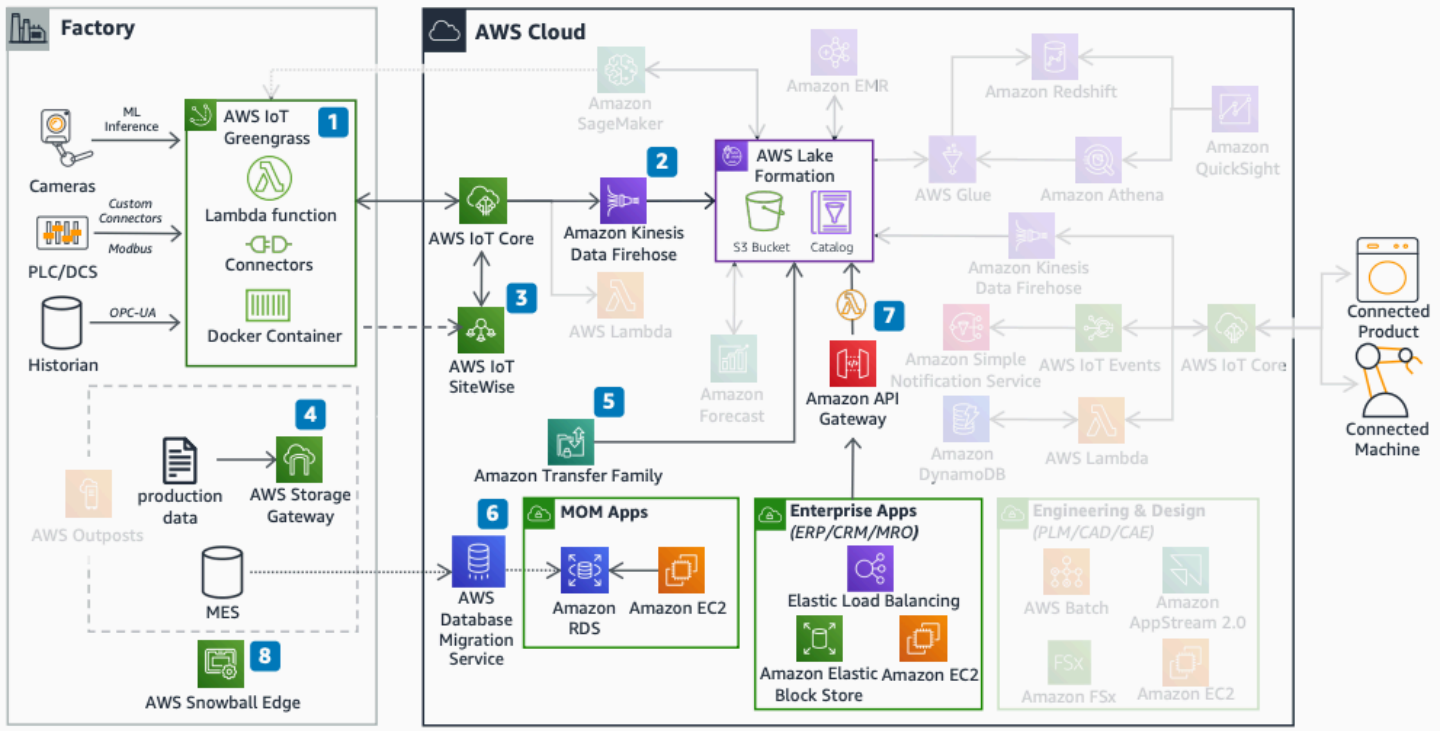

1. Connect industrial devices with AWS IoT Greengrass on an Edge gateway.
2. Stream industrial data with Amazon Data Firehose into the data lake.
3. With AWS IoT SiteWise, model assets and visualize data with AWS IoT SiteWise Monitor.
4. Sync unstructured data with AWS Storage Gateway from on-premises file shares.
5. Use AWS Transfer Family for file transfers into the data lake.
6. Use Amazon API Gateway and AWS Lambda for enterprise application interfaces.
7. Use AWS Snowball Edge for large data set migration from on-premises systems.
8. Use AWS Database Migration Service to synchronize databases into Amazon Relational Database Service (Amazon RDS).

Data lake and analytics

The data lake and analytics diagram shows how to process, analyze, and derive insights from manufacturing data.

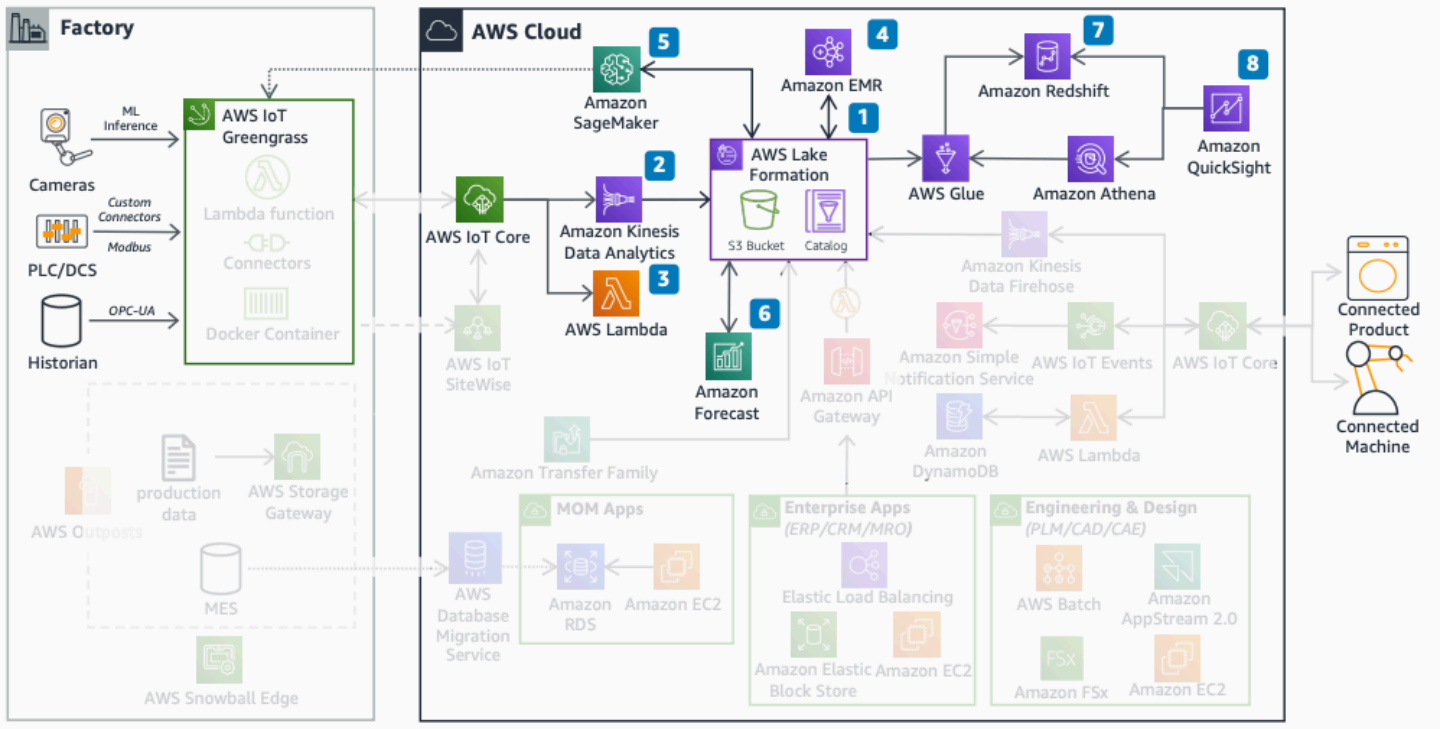

1. Use Lake Formation or Amazon S3 for the data lake structure to store raw and processed manufacturing data.
2. Use Amazon Managed Service for Apache Flink for streaming analytics on real-time factory data.
3. Use Lambda for near real-time analytics and event-driven processing.
4. Use Amazon SageMaker AI to build, train, and deploy ML models for predictive maintenance and quality.
5. Use Amazon EMR for large-scale data processing and transformation.
6. Use Amazon Forecast for demand forecasting.
7. Use Amazon Redshift for structured data warehousing and analytics.
8. Use Amazon Quick Sight with Amazon Athena and Amazon Redshift for business intelligence (BI) dashboards.

Application hosting

The application hosting diagram shows how to run enterprise and engineering applications on AWS.

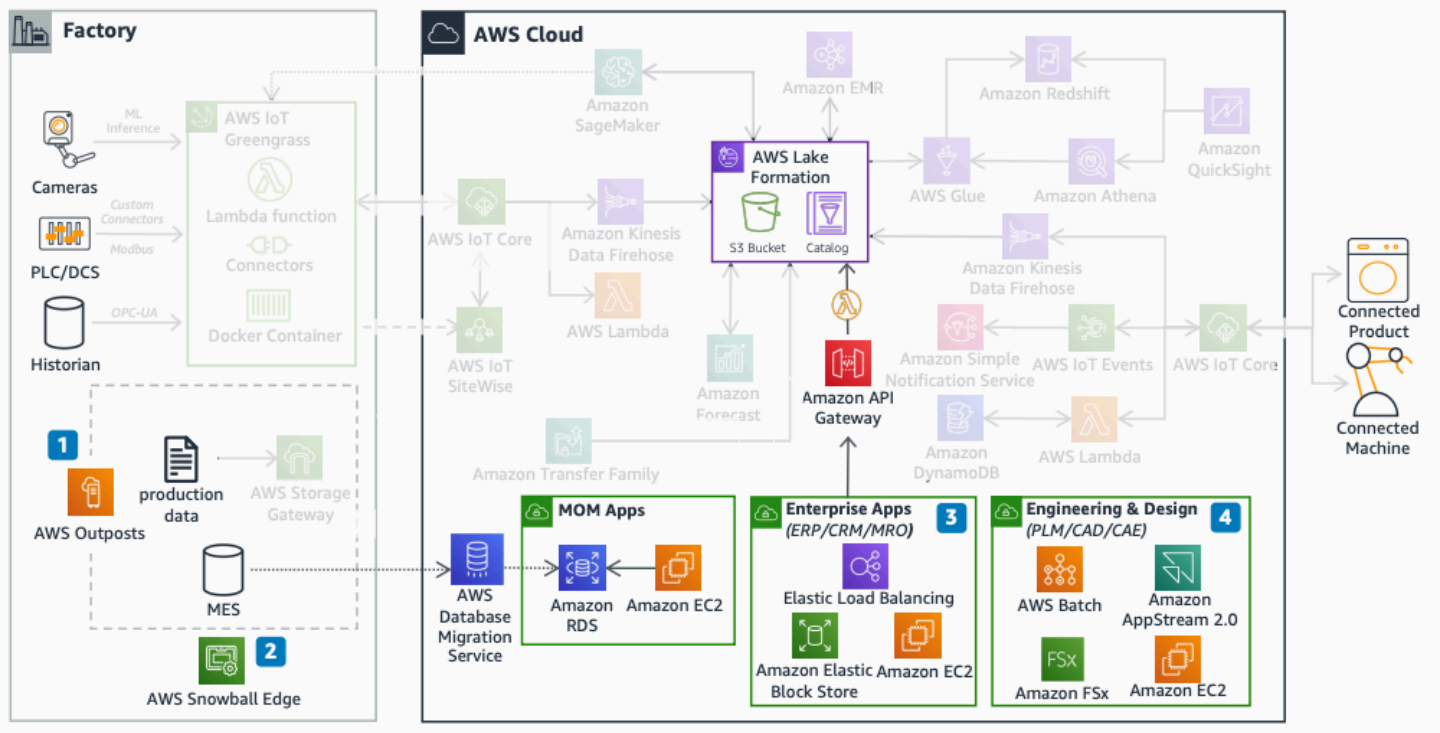

1. Use AWS Outposts for latency-sensitive workloads that must run close to on-premises systems.
2. Use AWS Snowball Edge for disconnected or intermittent connectivity use cases.
3. Host enterprise applications on Amazon EC2 with Auto Scaling and Elastic Load Balancing.
4. Use Amazon FSx and Amazon EC2 with WorkSpaces Applications for high performance computing (HPC), CAD, and CAE workloads.

Smart products

The smart products diagram shows how to connect manufactured products and machines to AWS for telemetry and event processing.

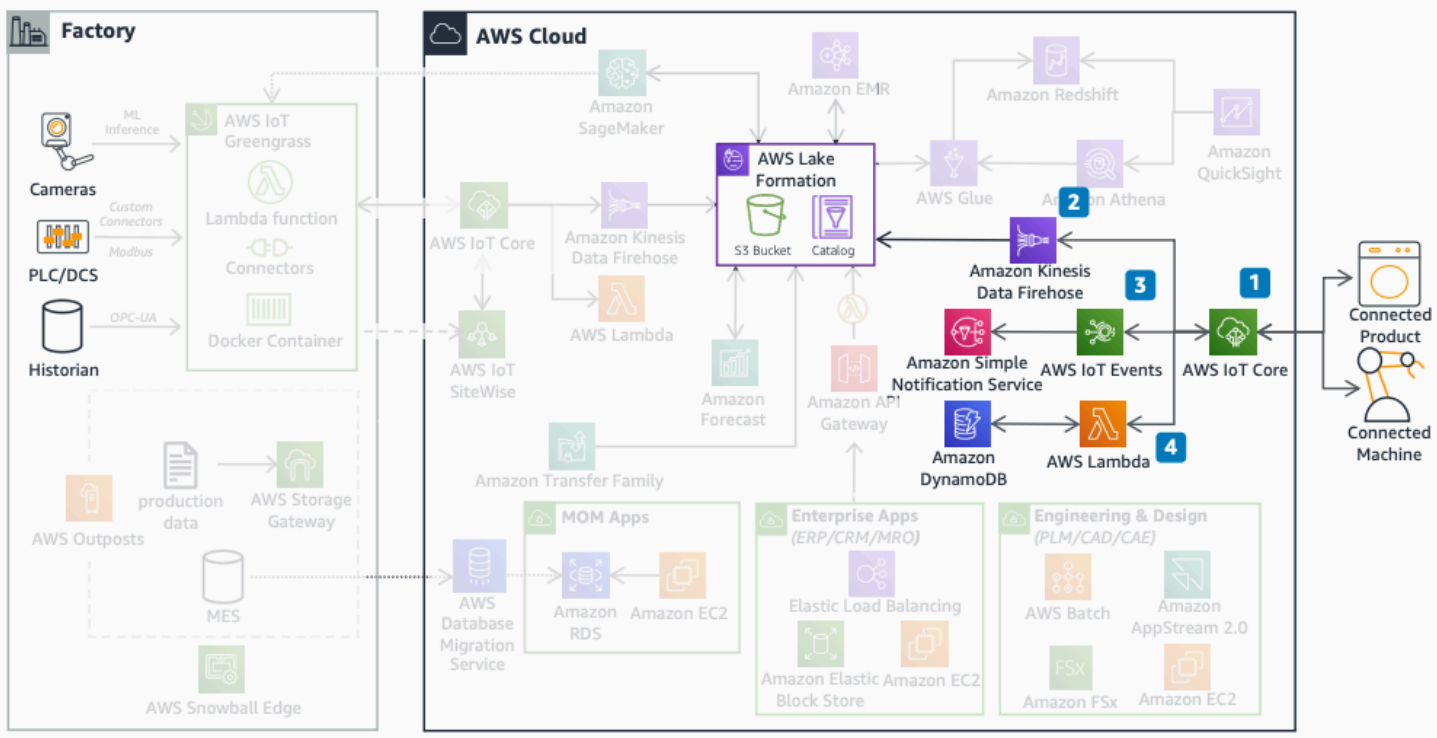

1. Use AWS IoT Core to connect products by using MQTT (Message Queuing Telemetry Transport).
2. Ingest telemetry to Amazon S3 with Amazon Data Firehose for durable storage.
3. Define event logic with AWS IoT Events and Amazon SNS for alerting.
4. Build microservices with Lambda and Amazon DynamoDB to process telemetry and serve applications.