


Guidance for Industrial Data Fabric on AWS


This architecture diagram illustrates how to effectively support Smart Manufacturing use cases on AWS. It shows the key components and their interactions, providing an overview of the architecture's structure and functionality.



- 1 Identify information related to industrial activities from on-premise equipment.
- 2 Collect real-time data from edge devices and transmit data streams securely to **AWS IoT SiteWise** in the cloud. Leveraging partners like Litmus, Domatica's EasyEdge, Siemens, and Belden accelerate integration with **AWS IoT SiteWise Edge**.
- 3 Connect edge devices through the Shop Floor Connectivity Framework using industrial protocols. Stream data securely to AWS Cloud services. Deploy cloud-developed machine learning models at the edge using **AWS IoT Greengrass** for defect detection and anomaly inference.
- 4 Leverage Siemens Industrial Edge, a centrally managed solution to connect data from assets and IT systems to **AWS IoT SiteWise**, **AWS IoT Core**, and **Amazon Simple Storage Service (Amazon S3)**. Deploy ML models and industrial applications.
- 5 Connect edge devices and IT systems using AWS partner solutions like HighByte. Contextualize and stream data to AWS storage and analytics services including **AWS IoT SiteWise**, **Amazon S3**, **Amazon S3 Tables**, **AWS TimeStream** for Influx DB, **Amazon Redshift**, **Amazon RDS**, **AWS IoT Core**, **Amazon Kinesis Data Streams**, **Amazon Kinesis Data Firehose**, and **Amazon Managed Service for Apache Flink (Amazon MSK)**.
- 6 Connect on-premises applications to **Amazon S3** through **AWS Storage Gateway** using NFS and SMB file shares.

Guidance for Industrial Data Fabric on AWS

This architecture diagram illustrates how to effectively support Smart Manufacturing use cases on AWS. It shows the key components and their interactions, providing an overview of the architecture's structure and functionality.



7 Extend AWS infrastructure and services to premises with **AWS Outposts**, a fully managed service. Run manufacturing-specialized applications on AWS services locally at the plant and integrate with AWS cloud infrastructure.

8 Ingest diverse data through AWS services. Stream real-time data via **AWS IoT SiteWise**, **AWS IoT Core**, **Amazon Kinesis**, and **Amazon MSK**. Transfer structured data using **AWS Database Migration Service (AWS DMS)** and **AWS Glue**. Process unstructured and semi-structured data with **Amazon S3**. Extract, create, and update ERP data with **Amazon AppFlow**.


9 Access AWS Analytics and AI/ML services through **Amazon SageMaker** Unified Studio. Find and query data and AI assets across the organization. Collaborate on projects to build analytics and AI artifacts. Share data, models, and generative AI applications securely. **Amazon SageMaker** Unified Studio is integrated with **Amazon Q** Developer, which provides AI-powered assistance for code development, data analysis, and ML workflows.

10 Process OT and IT data through AWS analytics services via **Amazon SageMaker** Unified Studio Portal. Catalog data with **AWS Glue** Catalog, transform through Visual ETL, and analyze using **Amazon Athena**. Process streaming data with **Amazon Managed Service for Apache Flink** and generate insights using **Amazon EMR**.



Guidance for Industrial Data Fabric on AWS

This architecture diagram illustrates how to effectively support Smart Manufacturing use cases on AWS. It shows the key components and their interactions, providing an overview of the architecture's structure and functionality.



- 11 Unify data across **Amazon S3** data lakes, including **S3 Tables**, and **Amazon Redshift** data warehouses with **Amazon SageMaker** Lakehouse. Build analytics and AI/ML applications on a single data copy using Apache Iceberg-compatible tools and engines.
- 12 Organize assets, users, and projects within **Amazon SageMaker** Unified Studio domains. Create multiple domains to match enterprise structure. Collaborate in projects to manage data assets, analyze data, develop ML models, and build generative AI applications.
- 13 Enrich technical catalog metadata with business context using **Amazon SageMaker** Catalog. Discover and access approved data and models through generative AI semantic search. Monitor data quality, track lineage, and enforce access policies in **Amazon SageMaker** Unified Studio.
- 14 Build, train, and deploy machine learning models and generative AI capabilities using **Amazon SageMaker AI** and **Amazon Bedrock**. Leverage agentic AI to improve manufacturing, optimize supply chain, get digital twins agents for Engineering and Design, all impacting sustainability.
- 15 Integrate with cloud-hosted ERP, Supply Chain, Maintenance, and WMS/TMS manufacturing solutions, including Model Context Protocol (MCP) servers for industrial knowledge. Exchange data with enterprise platforms like Snowflake and Databricks to enhance manufacturing analytics.
- 16 Visualize data with **Amazon Managed Grafana** from **Amazon Redshift** or **Amazon S3** via **Amazon Athena**. Build dashboards using **Amazon QuickSight** and **Amazon Athena**.




Reviewed for technical accuracy October 2025
© 2025, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Reference Architecture

Guidance for Industrial Data Fabric on AWS

Edge Services


This architecture diagram illustrates how to effectively collect data from the factory and send to AWS services in the cloud.



Guidance for Industrial Data Fabric on AWS

Cloud Services

This architecture diagram illustrates how data collected from the factory can be used with AWS cloud services.



1 Process diverse data types through AWS services. Stream real-time data through **AWS IoT SiteWise**, **AWS IoT Core**, **Amazon Kinesis**, and **Amazon MSK**. Transfer structured data using **AWS DMS** and **AWS Glue**. Process unstructured data with **Amazon S3**. Extract and update ERP data using **Amazon AppFlow**.

2 Access AWS Analytics and AI/ML services through **Amazon SageMaker** Unified Studio. Find and query data and AI assets organization-wide. Collaborate on projects to build analytics and AI artifacts. Share data, models, and generative AI applications securely. **Amazon SageMaker** Unified Studio is integrated with **Amazon Q** Developer, which provides AI-powered assistance for code development, data analysis, and ML workflows.

3 Process OT and IT data through AWS analytics services via **Amazon SageMaker** Unified Studio. Catalog data using **AWS Glue** Catalog, transform with **AWS Glue** Visual ETL, and analyze with **Amazon Athena**. Process streams with **Amazon Managed Service for Apache Flink** and generate insights using **Amazon EMR**.

4 Unify data across **Amazon S3** data lakes, including **S3** Tables, and **Amazon Redshift** data warehouses with **Amazon SageMaker** Lakehouse. Build powerful analytics and AI/ML applications on a single copy of data using all Apache Iceberg-compatible tools and engines


5 Leverage **Amazon Bedrock** and agentic AI to improve manufacturing and optimize supply chain. Use **Amazon Bedrock AgentCore**, a comprehensive set of enterprise-grade services that help securely deploy and operate AI agents at scale, with components like **AgentCore Runtime** for low-latency serverless environments.



Guidance for Industrial Data Fabric on AWS

Cloud Services with Partners and Agentic AI

This architecture diagram illustrates how data collected from the factory can be used with AWS cloud services.



- 1 Access functionality and tools from AWS Analytics and AI/ML services through **Amazon SageMaker** Unified Studio's single development environment. Query data and AI assets organization-wide. Collaborate on projects to build and share analytics and AI artifacts, including data, models, and generative AI applications, in a secure environment
- 2 Leverage **Amazon Bedrock** and agentic AI to improve manufacturing and optimize supply chain. Use **Amazon Bedrock AgentCore**, a comprehensive set of enterprise-grade services that help securely deploy and operate AI agents at scale, with components like **AgentCore Runtime** for low-latency serverless environments. Access production documentation and policies using Retrieval Augmented Generation (RAG), through managed **Amazon Bedrock** Knowledge Bases. Connect to Model Context Protocol (MCP) servers to allow agents to connect to systems where data lives.
- 3 Integrate with cloud-hosted manufacturing solutions, including MCP servers. Exchange data with enterprise platforms using Iceberg REST Catalog. Build ETL pipelines with **Amazon SageMaker** Zero-ETL for **Amazon S3**, **Amazon S3 Tables**, and **Amazon Redshift**. Query data using **Amazon Athena**.

