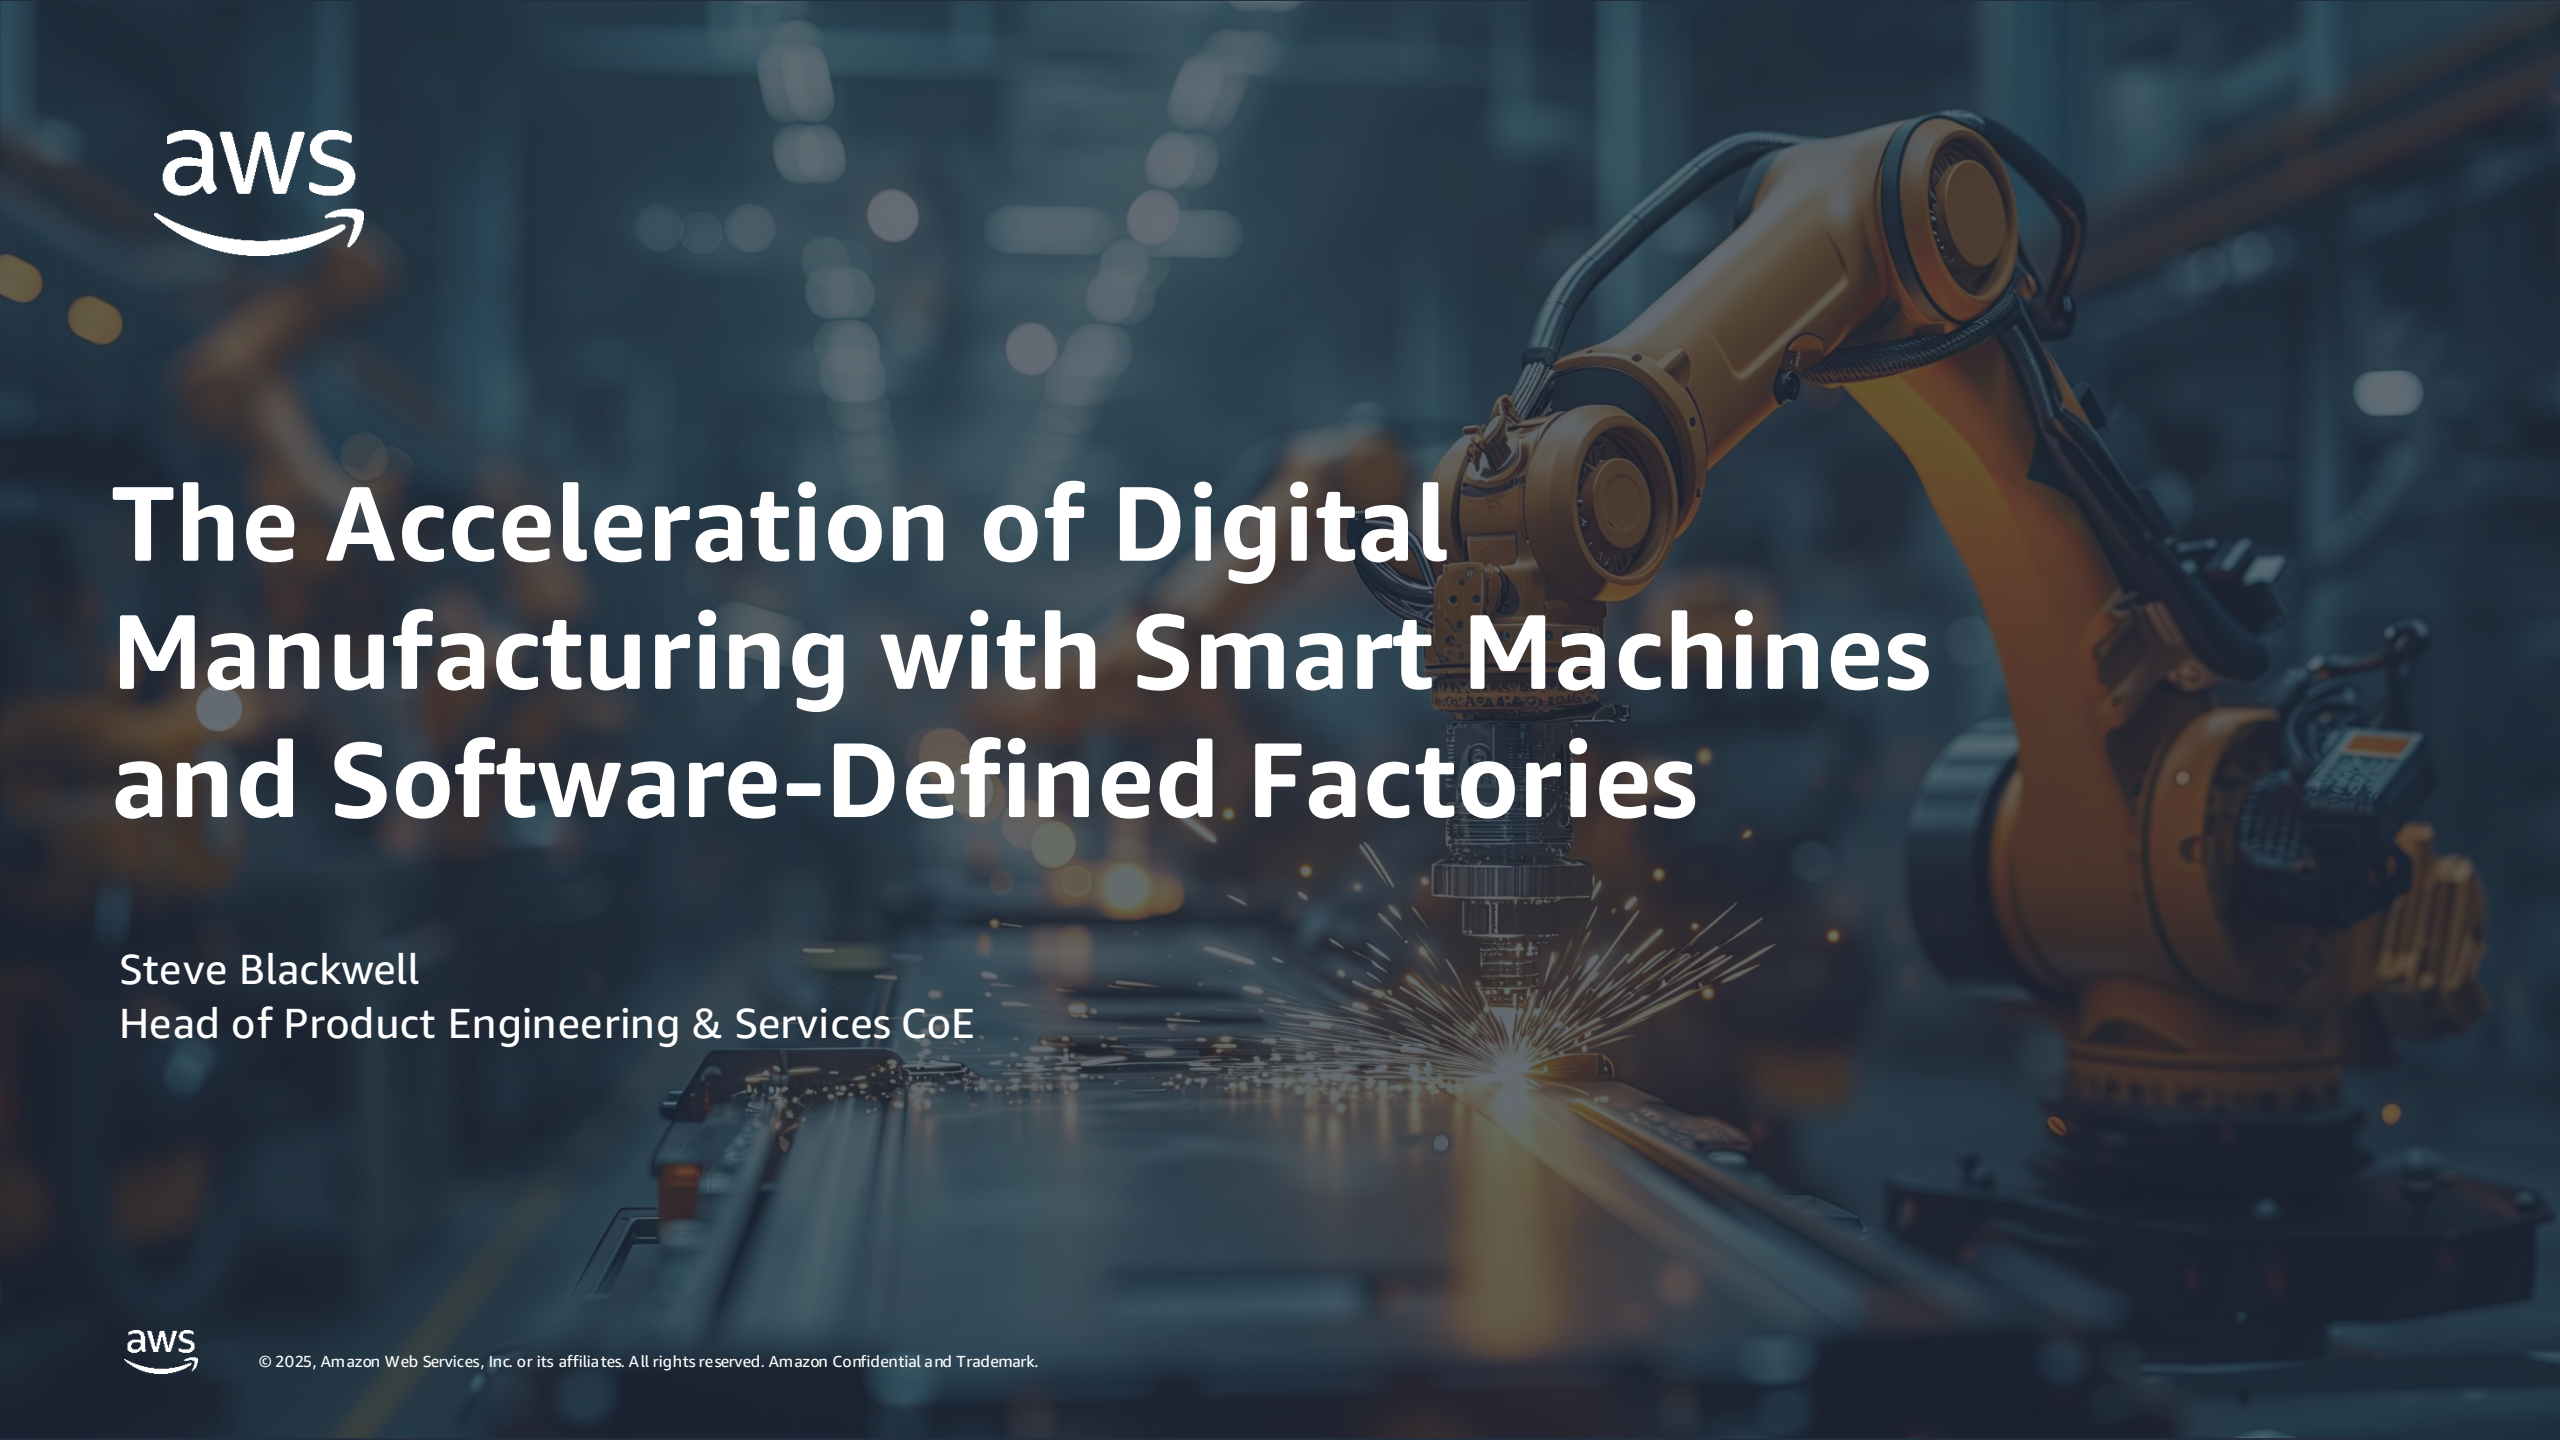

# The Acceleration of Digital Manufacturing with Smart Machines and Software-Defined Factories

Steve Blackwell
Head of Product Engineering & Services CoE

# What is Digital Manufacturing?

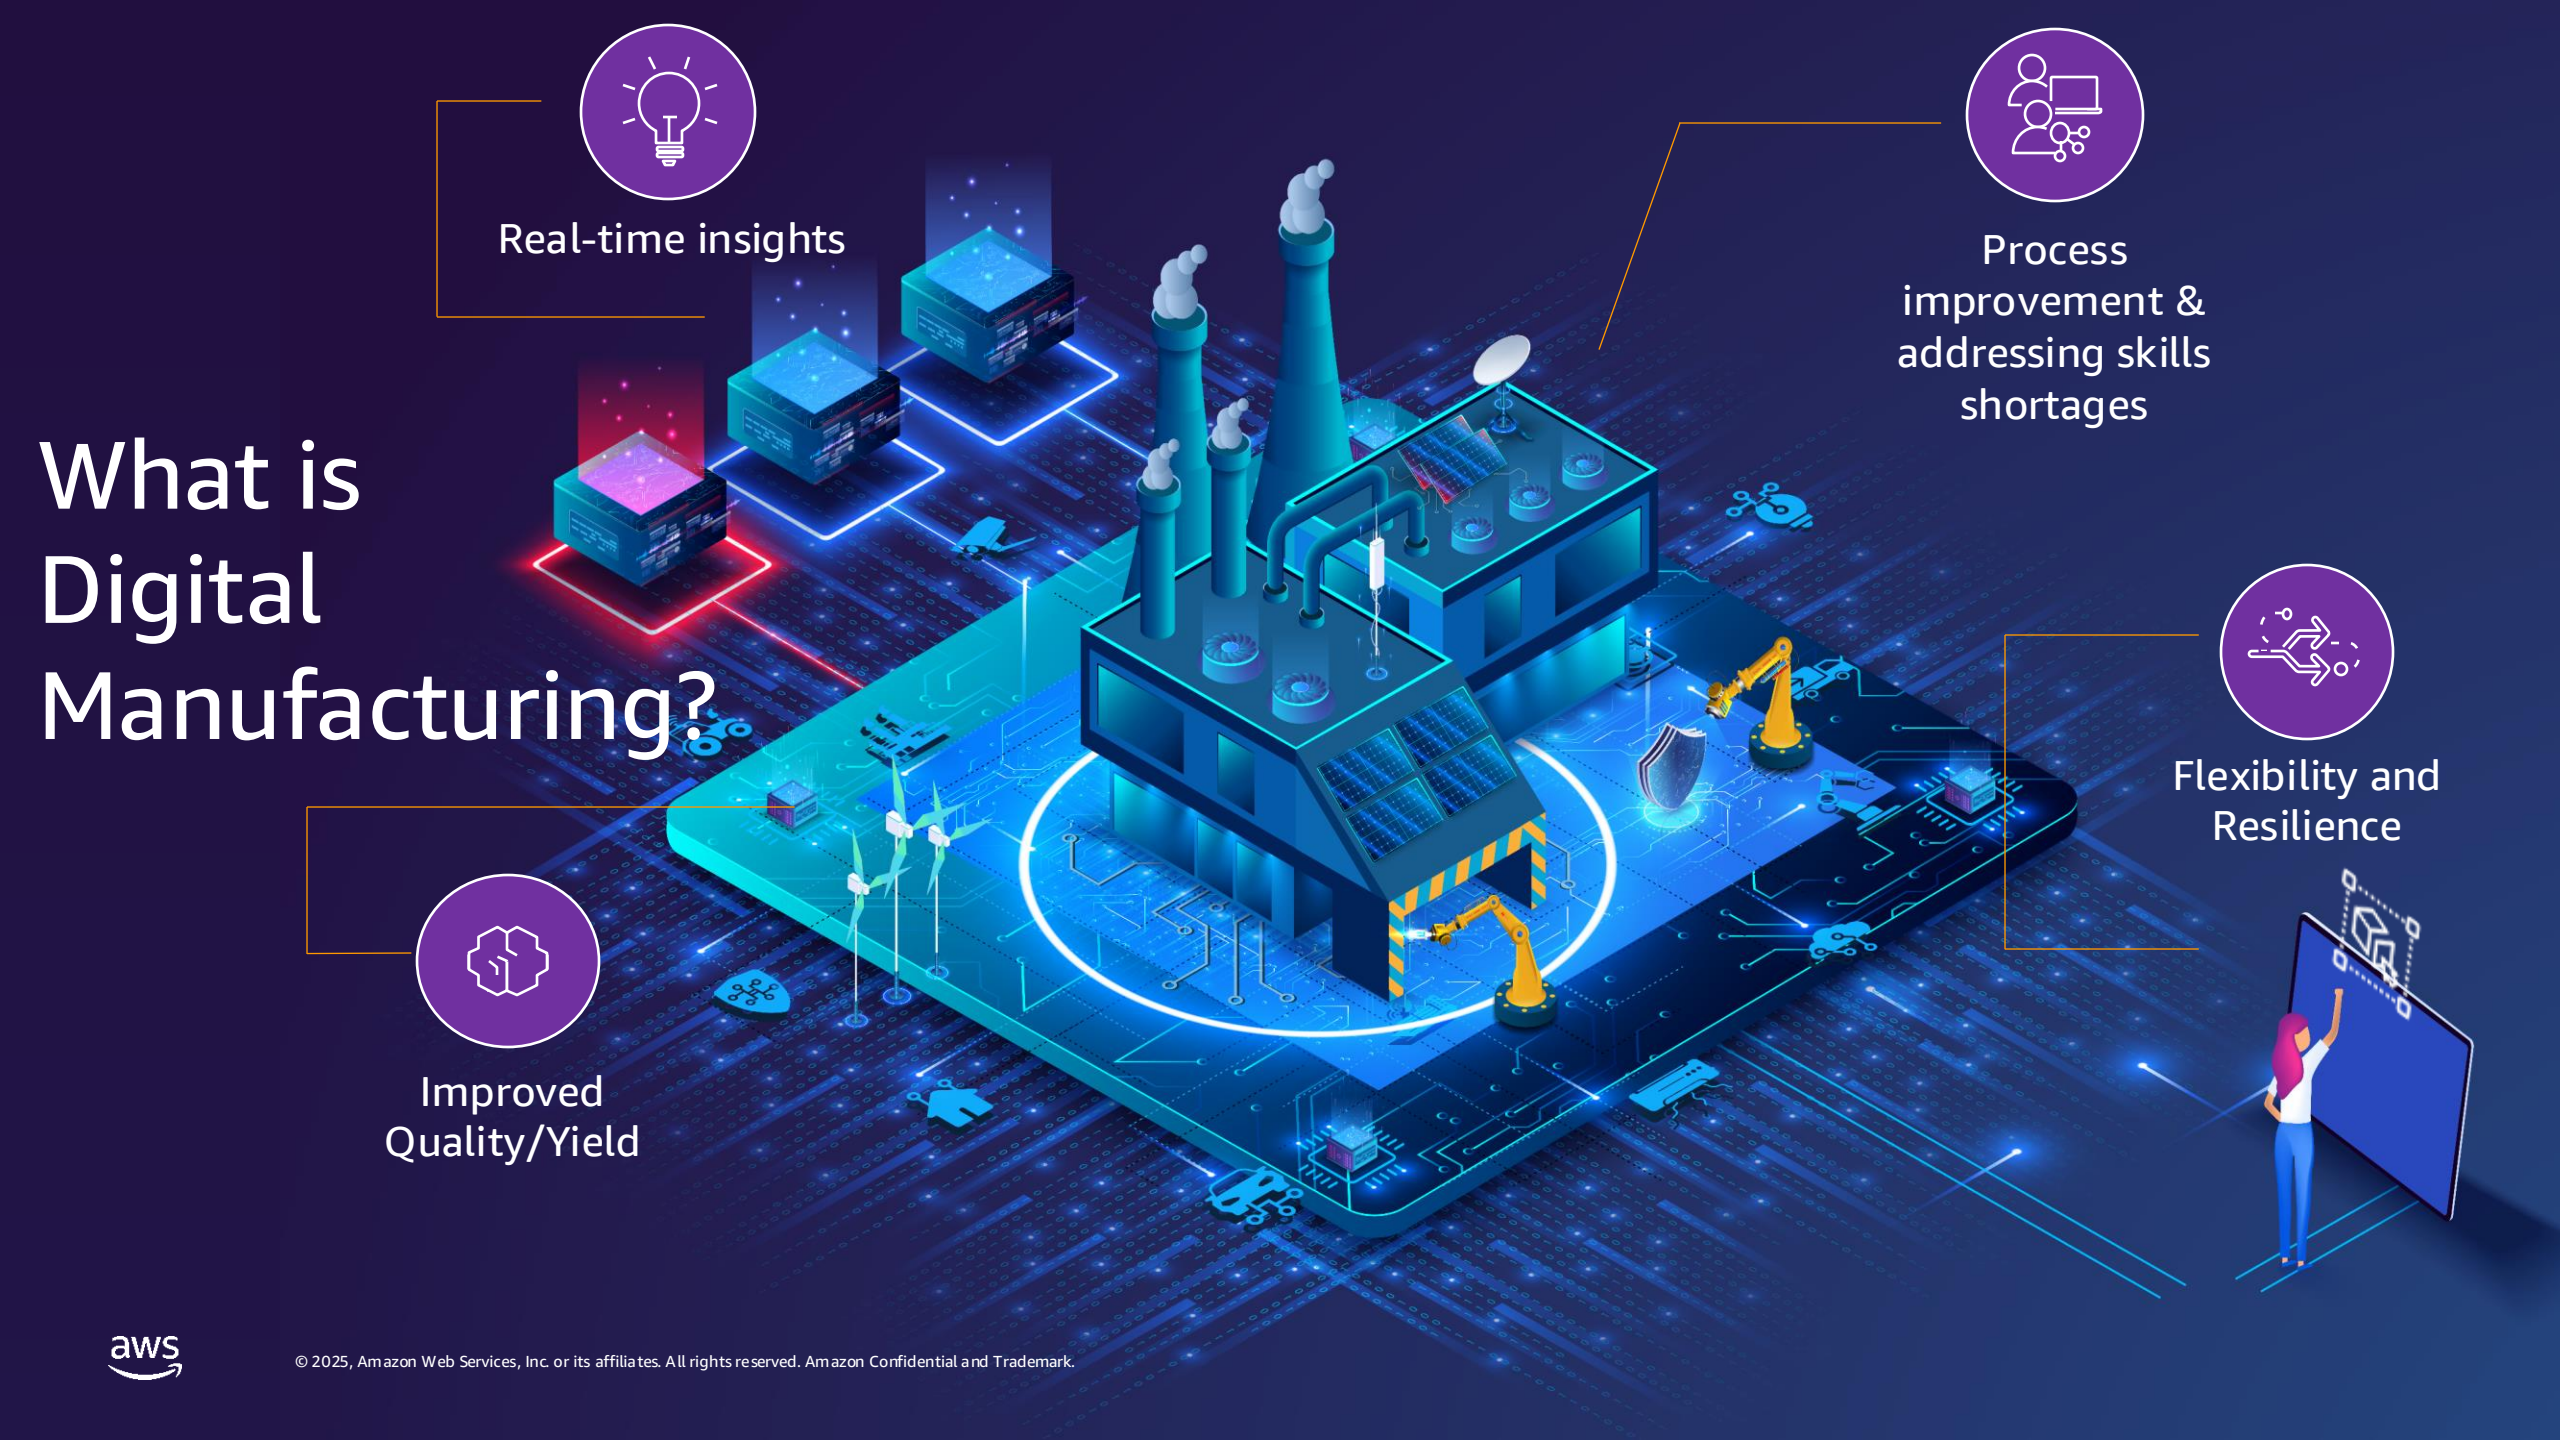

Real-time insights

Process improvement & addressing skills shortages

Flexibility and Resilience

Improved Quality/Yield

# Digital transformation leaders achieve outsized impact

Yield, energy, and throughput optimization

**5-10%**
energy reduction*

Optimized production

**16%**
Overall equipment effectiveness (OEE) increase†

Reduced water, air, gas, electricity, and steam use

**9%**
Greenhouse gas emission reduction*

Digitally enabled supply and demand planning

**5-10%**
improvement of on-time-in-full fulfillment*

Faster time to results, improved collaboration,

**22%**
increased speed to market †

## Modern Industrial Data Strategy

Democratizing data access to integrate operation and information technology networks



Source: *McKinsey, "The Digital Revolution is Brewing in the Industrials Sector," 2022
† The Hackett Group, **Business Impact of Cloud Adoption for Industrial Manufacturers, 2023

# Industrial customers face challenges

## Data access & integration

Integrate data from new and legacy equipment using different protocols

## Data context & management

Structure and organize large amounts of disparate machine data

## Real-time decision making

Relevant information at the right time

## Scale

Manage assets, device fleets and data across multiple machines, processes, and sites

**Security designed for the most sensitive industries**

# Automation & Robotics

# Levels of Automation in Digital Manufacturing

| Level 1 | Level 2 | Level 3 | Level 4 |
| --- | --- | --- | --- |
|  | 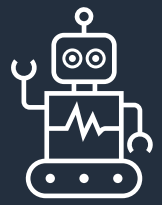 |  |  |
| **Programmed**<br>**Physical Automation** | **Adaptive**<br>**Automation + ML** | **Partially Autonomous**<br>**AI** | **Fully Autonomous**<br>**Physical AI** |
| Predefined tasks only<br>Controlled environments<br>Human-programmed actions | Flexible task sequencing<br>Real-time environmental adaptation<br>Dynamic action ordering | Intelligent planning & execution<br>Learning through demonstration<br>Limited human input required | Cross-domain operation<br>Minimal supervision needed<br>Fluid adaptation to new scenarios |

# What is Physical AI?

Physical AI describes systems that integrate understanding, reasoning, and learning to interact with the physical world. Physical AI is the next leap in advanced autonomy.

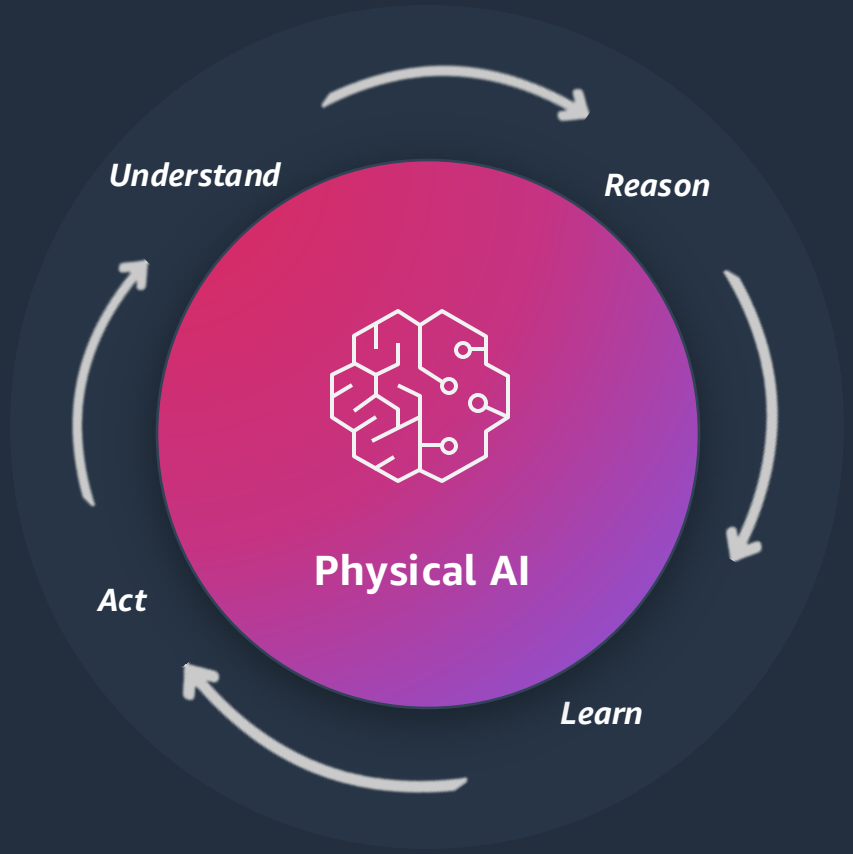


**Understand** — Integrating models and algorithms with sensors, real world and simulation data

**Reason** — Predicting real-time actions that will be realized in the physical world

**Learn** — Iteratively learning through feedback loops to improve overall performance of the system

**Act** — Applying predictions to perform real-world actions

# Amazon and Robotics, a decade of innovation

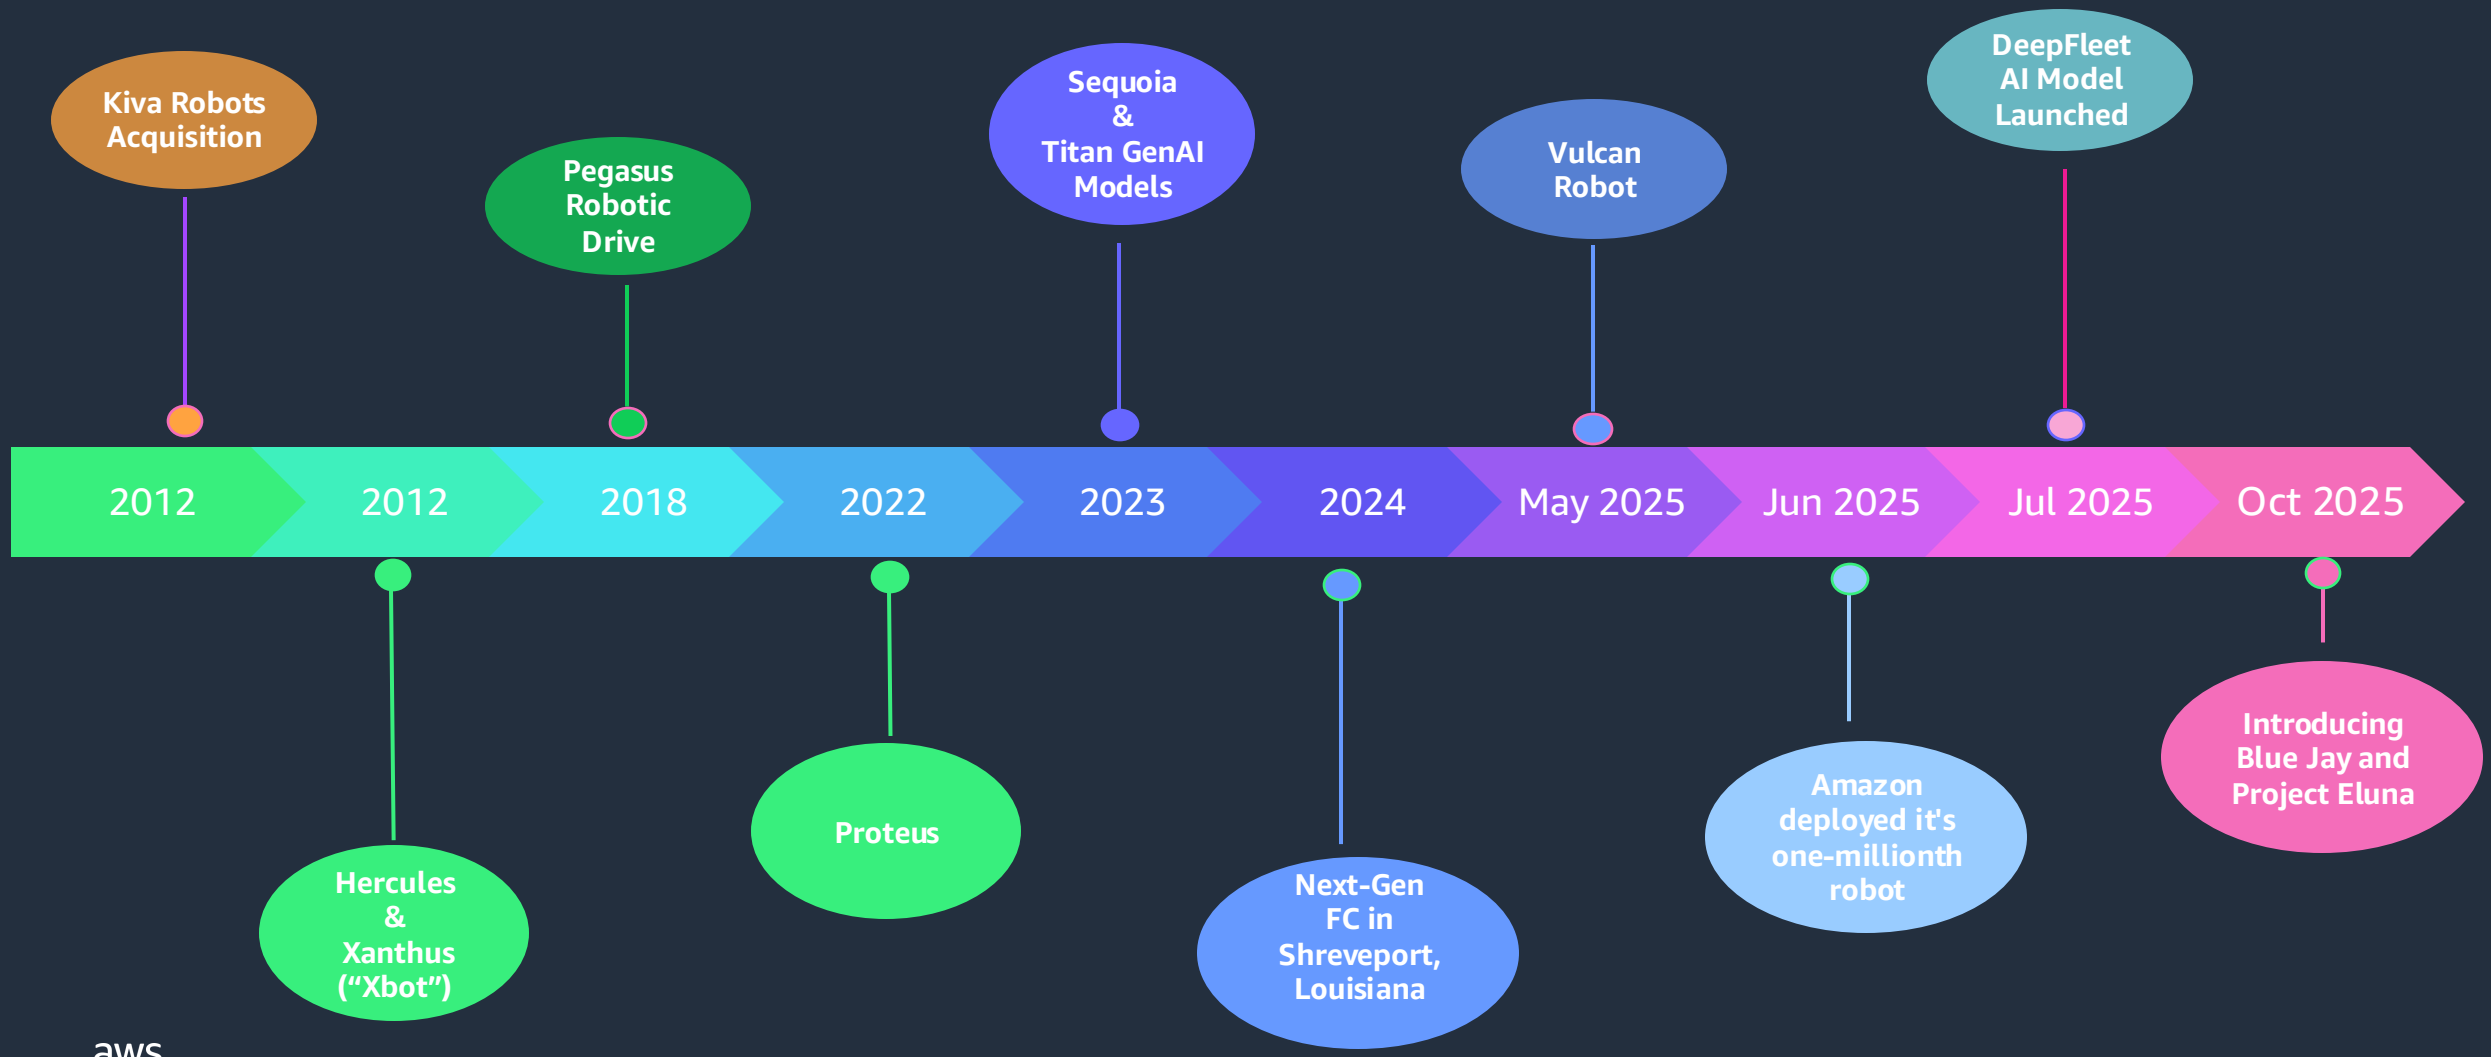

# Amazon Robotics Powered by AWS

## Sequoia

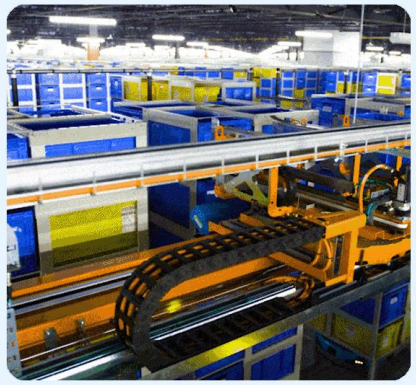

## Titan

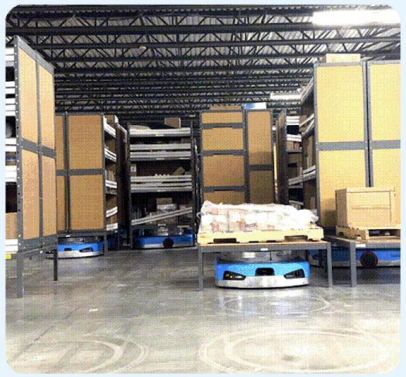

## Vulcan Pick

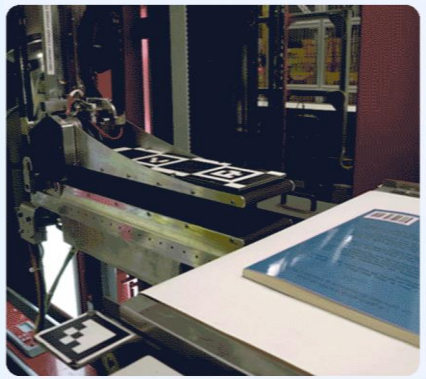

## Vulcan Stow

## Hercules

## Sparrow

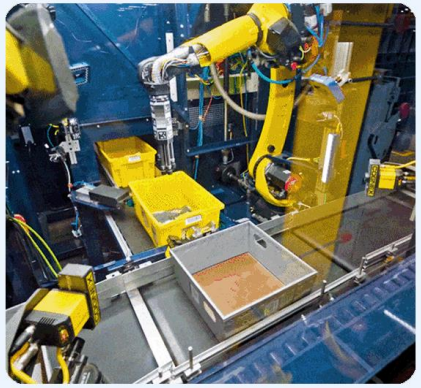

## Packaging Automation

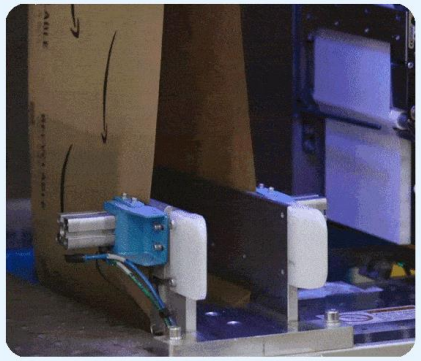

## Robin

## Cardinal

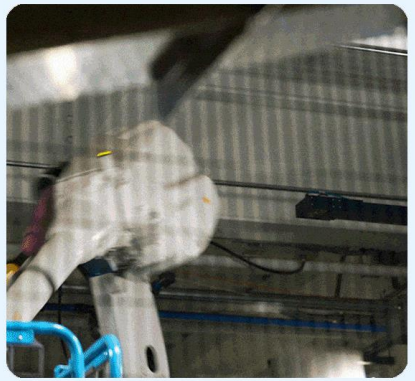

## Proteus

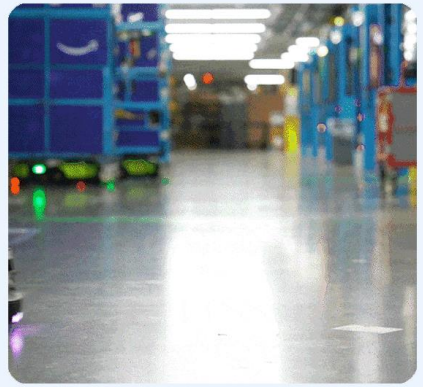

# Innovations that enable the development of Physical AI

## Improved Perception

Improvements in hardware (e.g., high-res cameras, LiDAR, etc.) and computer vision algorithms allow robots to perceive their environment better

## Advanced AI

Improving robot model approaches as well optimizing for edge hardware (e.g., simulation, RL, VLA) allowing robots to learn behaviors and adapt to their environments

## Ubiquitous Connectivity

Devices can be run on the edge and through the cloud

# Evolving manufacturing data landscape (ISA95)

## Today

**Monolithic pyramid**

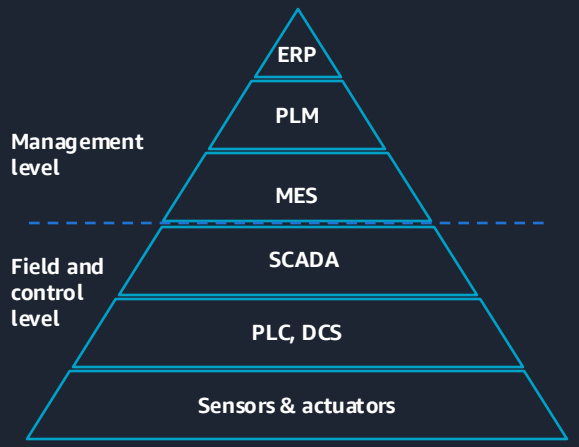


## Today/tomorrow

**Converging IT and OT**

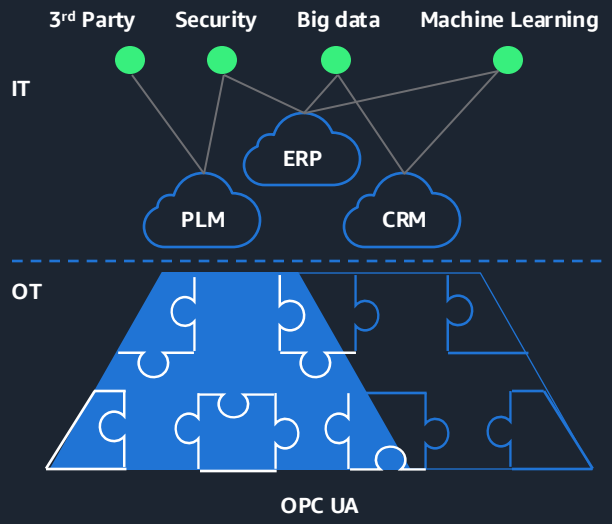


## Desired: Near future

**Smart factory**

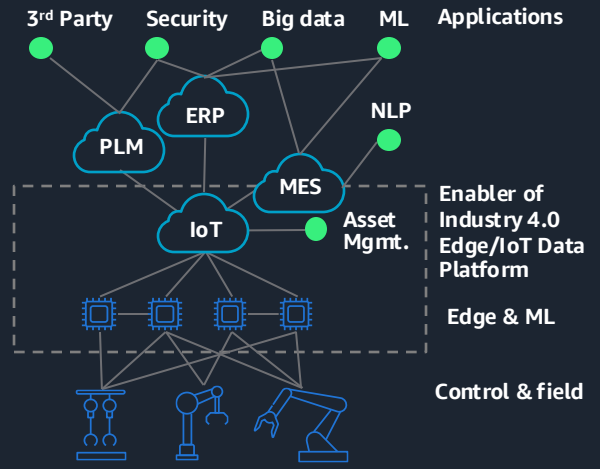

# Are Smart Machines the answer?

# Smart machines technical challenges

### Connect and manage machines at scale

Connect and manage smart machines at scale across customers and geographies

### Build industrial data foundation

Structure, store, and organize large amounts of disparate machine data

### Do more at the edge

Process information and use of AI/ML at the edge. In a secure and centralized way.

### Deploy applications powered by AI and digital twins

Innovate fast and build or improve applications for customers, engineers, and distributors

**Security and compliance challenges**

# Smart machines digital journey

01.

## Industrial connectivity

- KPI visualizations, diagnostics, remote control
- Provisioning, IT/OT interface, remote support

02.

## Machine monitoring

- Add functionality and security after shipment (OTA)
- Machine learning–based analytics
- Maintenance business transformation

03.

## Machine as a service

- New subscription or usage-based billing
- Enable horizontal solutions and value
- New service models with SLAs
- Zero downtime

INCREASING VALUE

# Smart Industrial Machines Benefits

## For Machine OEMs

Achieve Operational Excellence by reducing time to repair and on-site maintenance.

Increase customer satisfaction and loyalty

Device Management remotely and at scale

## For Manufacturers

Increase productivity

Reduce downtime and costs with Predictive Maintenance

Improve product quality

Reduce waste (energy, materials, etc)

# Is Software the answer for Manufacturing?

# Software-Defined Manufacturing

| | |
|---|---|
| L3 | Software-Defined Production |
| L2 | Software-Defined Automation |
| L1 | Software-Defined Control |

## Smart Machines

ARTIFICIAL INTELLIGENCE & DIGITAL TWINS

DATA FOUNDATION

SOFTWARE “DEFINED”

# Industrial Data Fabric framework

## Act

- Real-time OEE
- Process / Cell Optimization
- PdM
- Energy Management
- Quality/Yield

## Analyze

- AI, ML
- Digital twin
- Interactive query
- Real-time analytics
- Predictive analytics

## Contextualize

- Asset modeling
- Event processing
- Knowledge Graph DB

## Store

- Data warehousing and processing
- Industrial data lake
- Time series data

## Ingest

- Local decisions
- Third-party data
- Data store and forward
- Security and management

PLC's/DCS's | Historians | SCADA | Connected Product | MES | PLM | ERP | MRO | Logistics

# Software-Defined Manufacturing

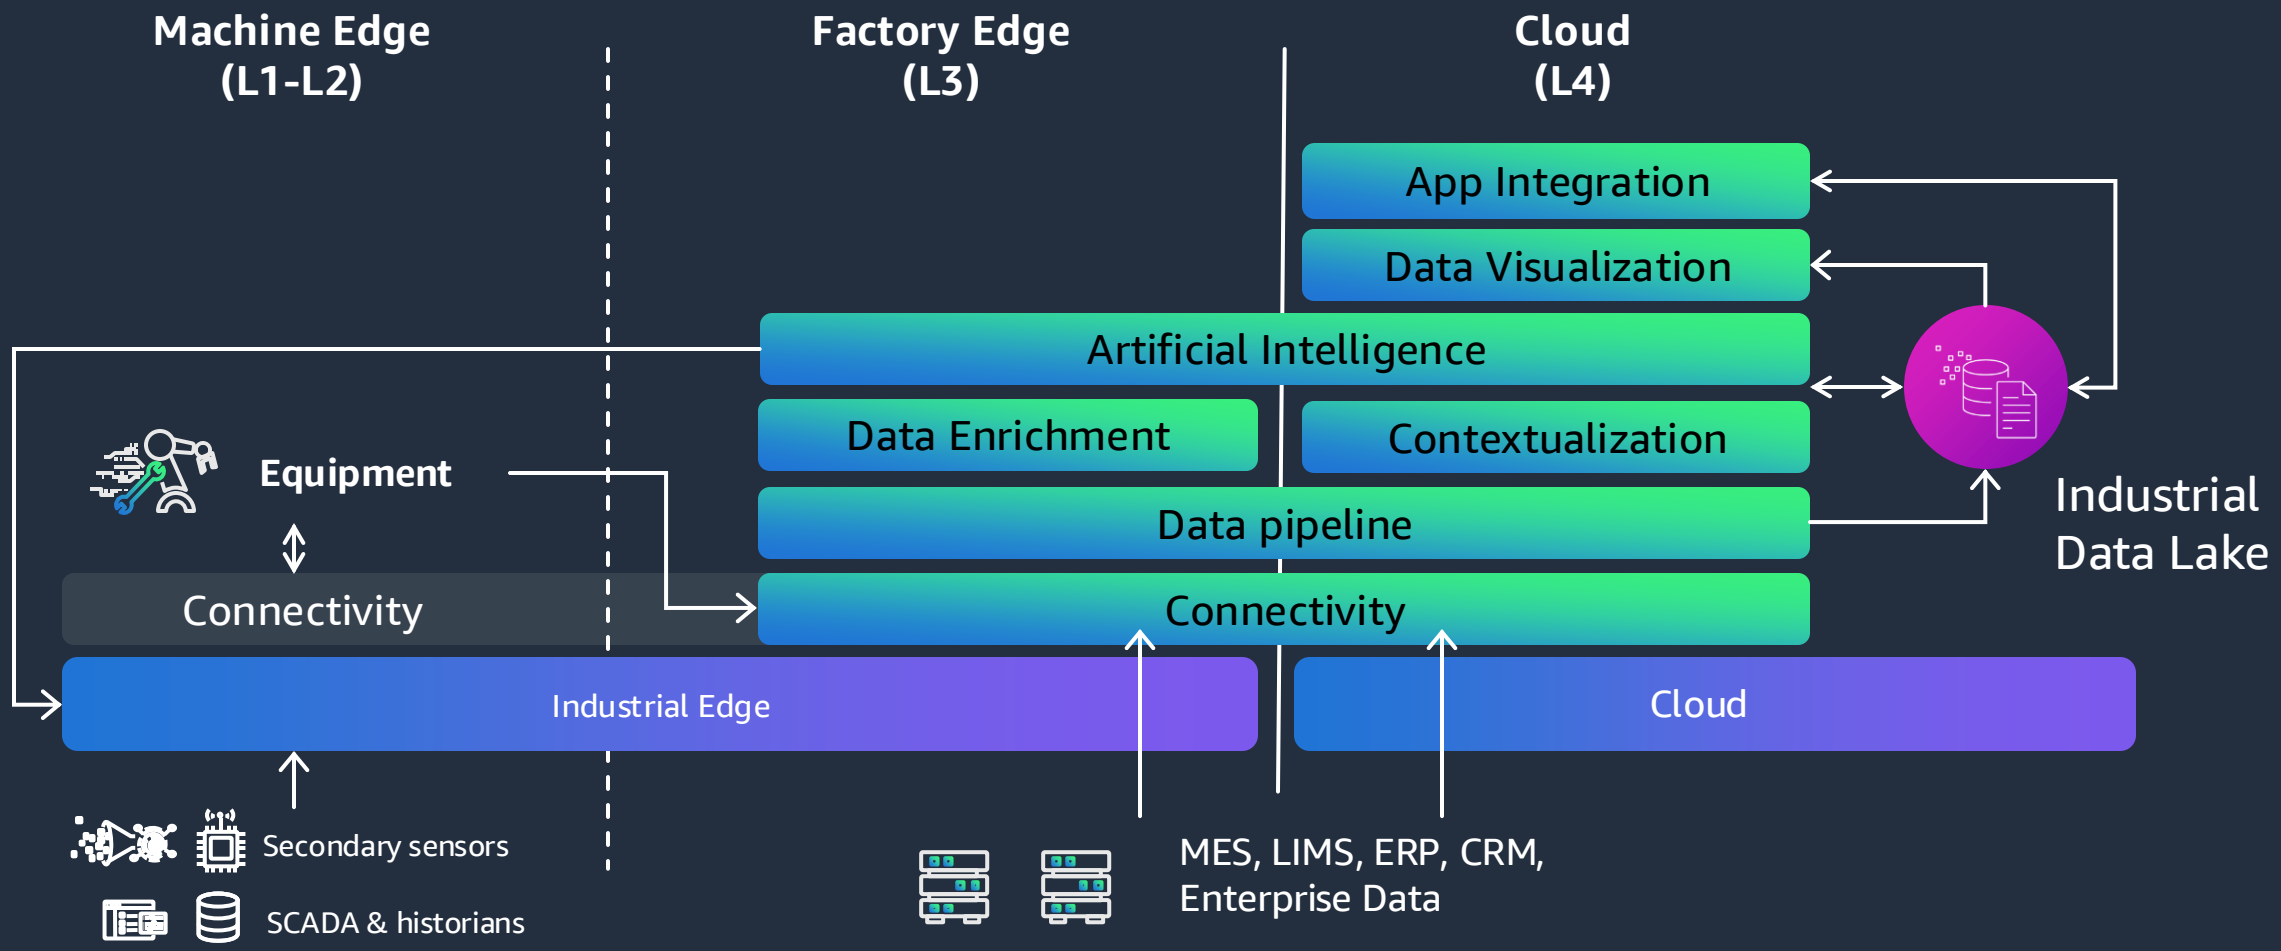

# Smart Machines Conceptual Architecture

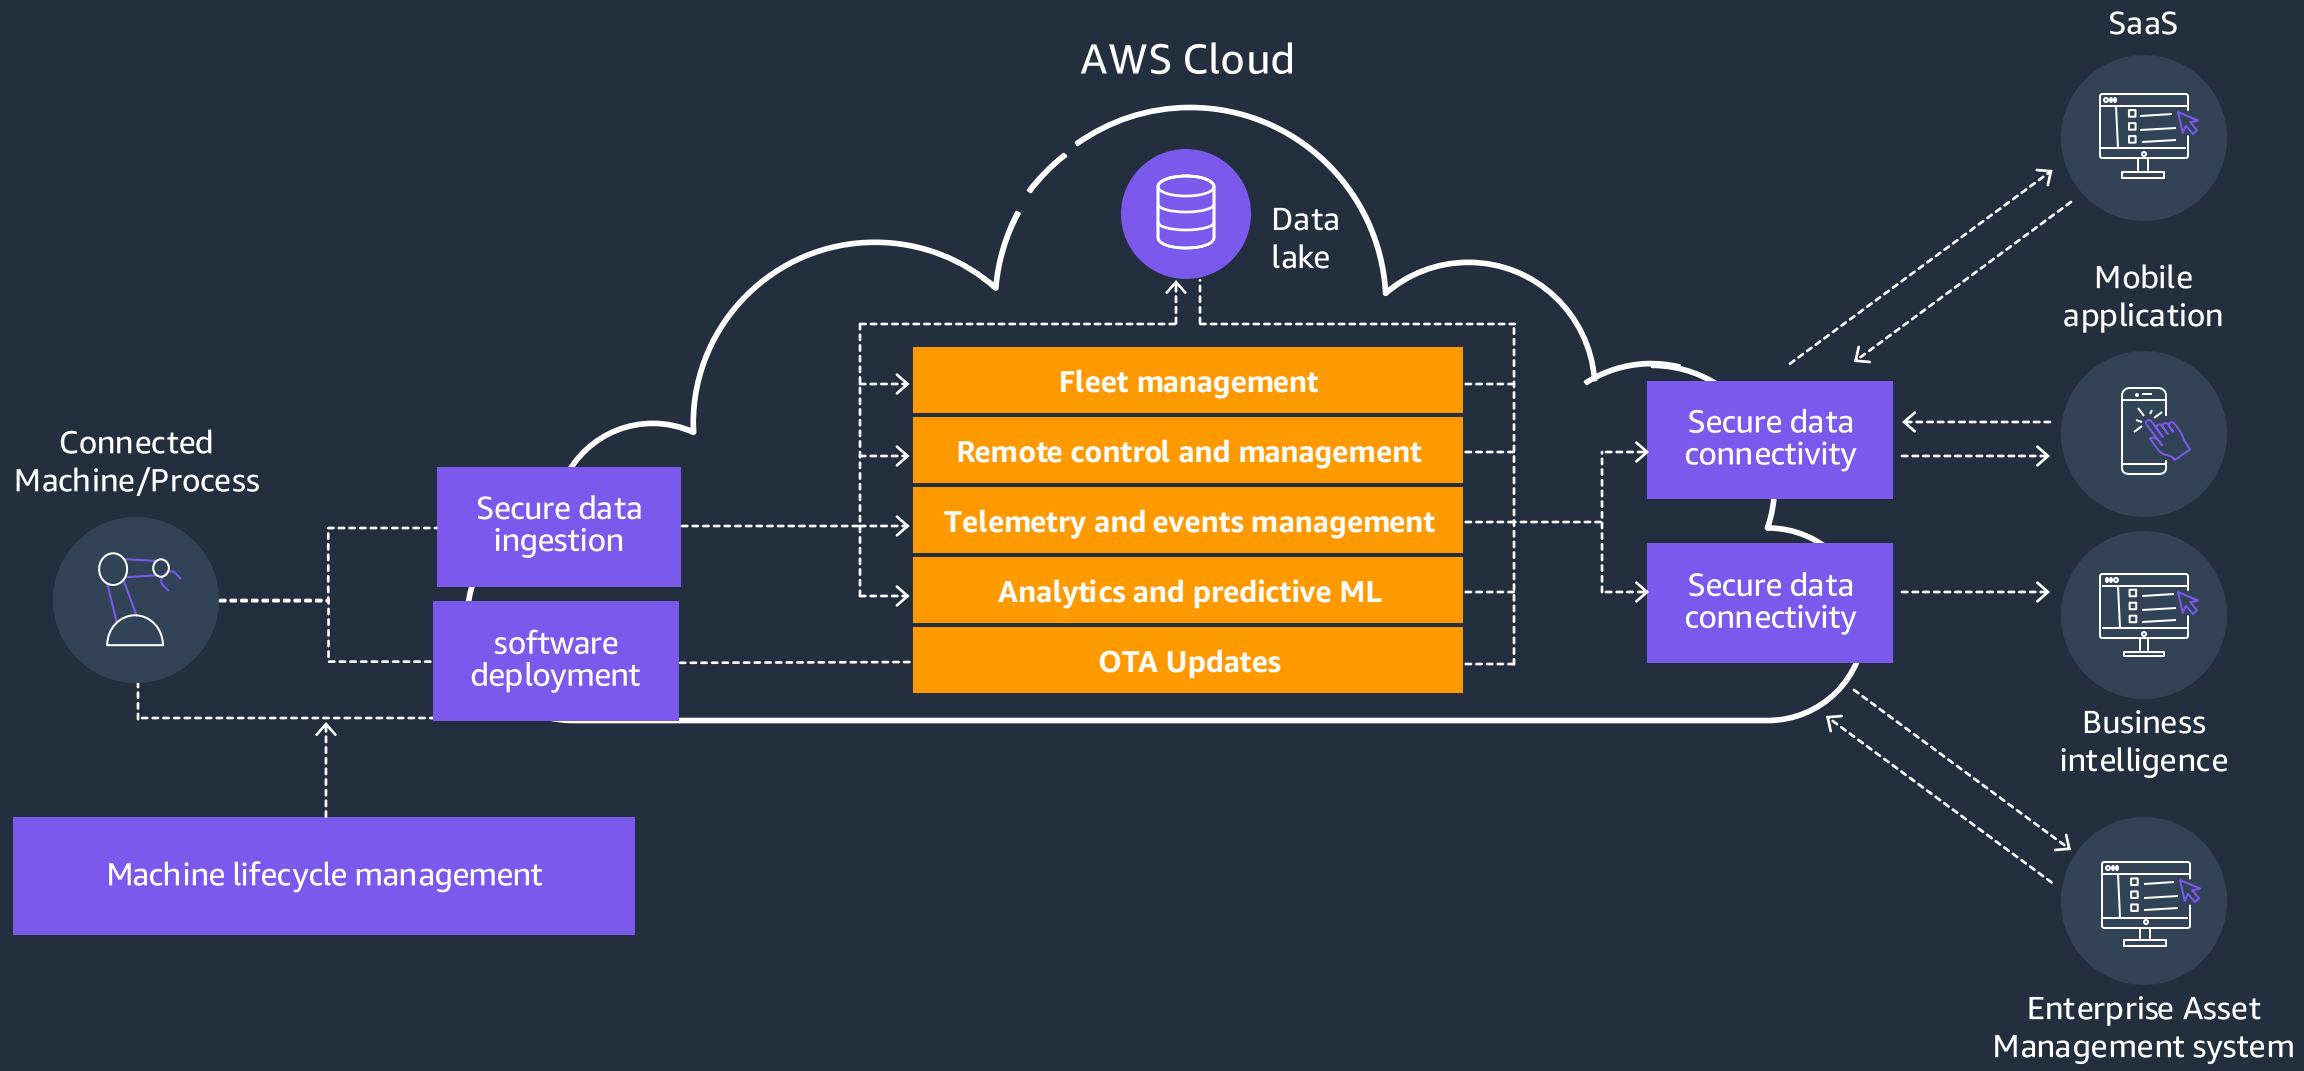

# Software-Defined Manufacturing with Smart Machines

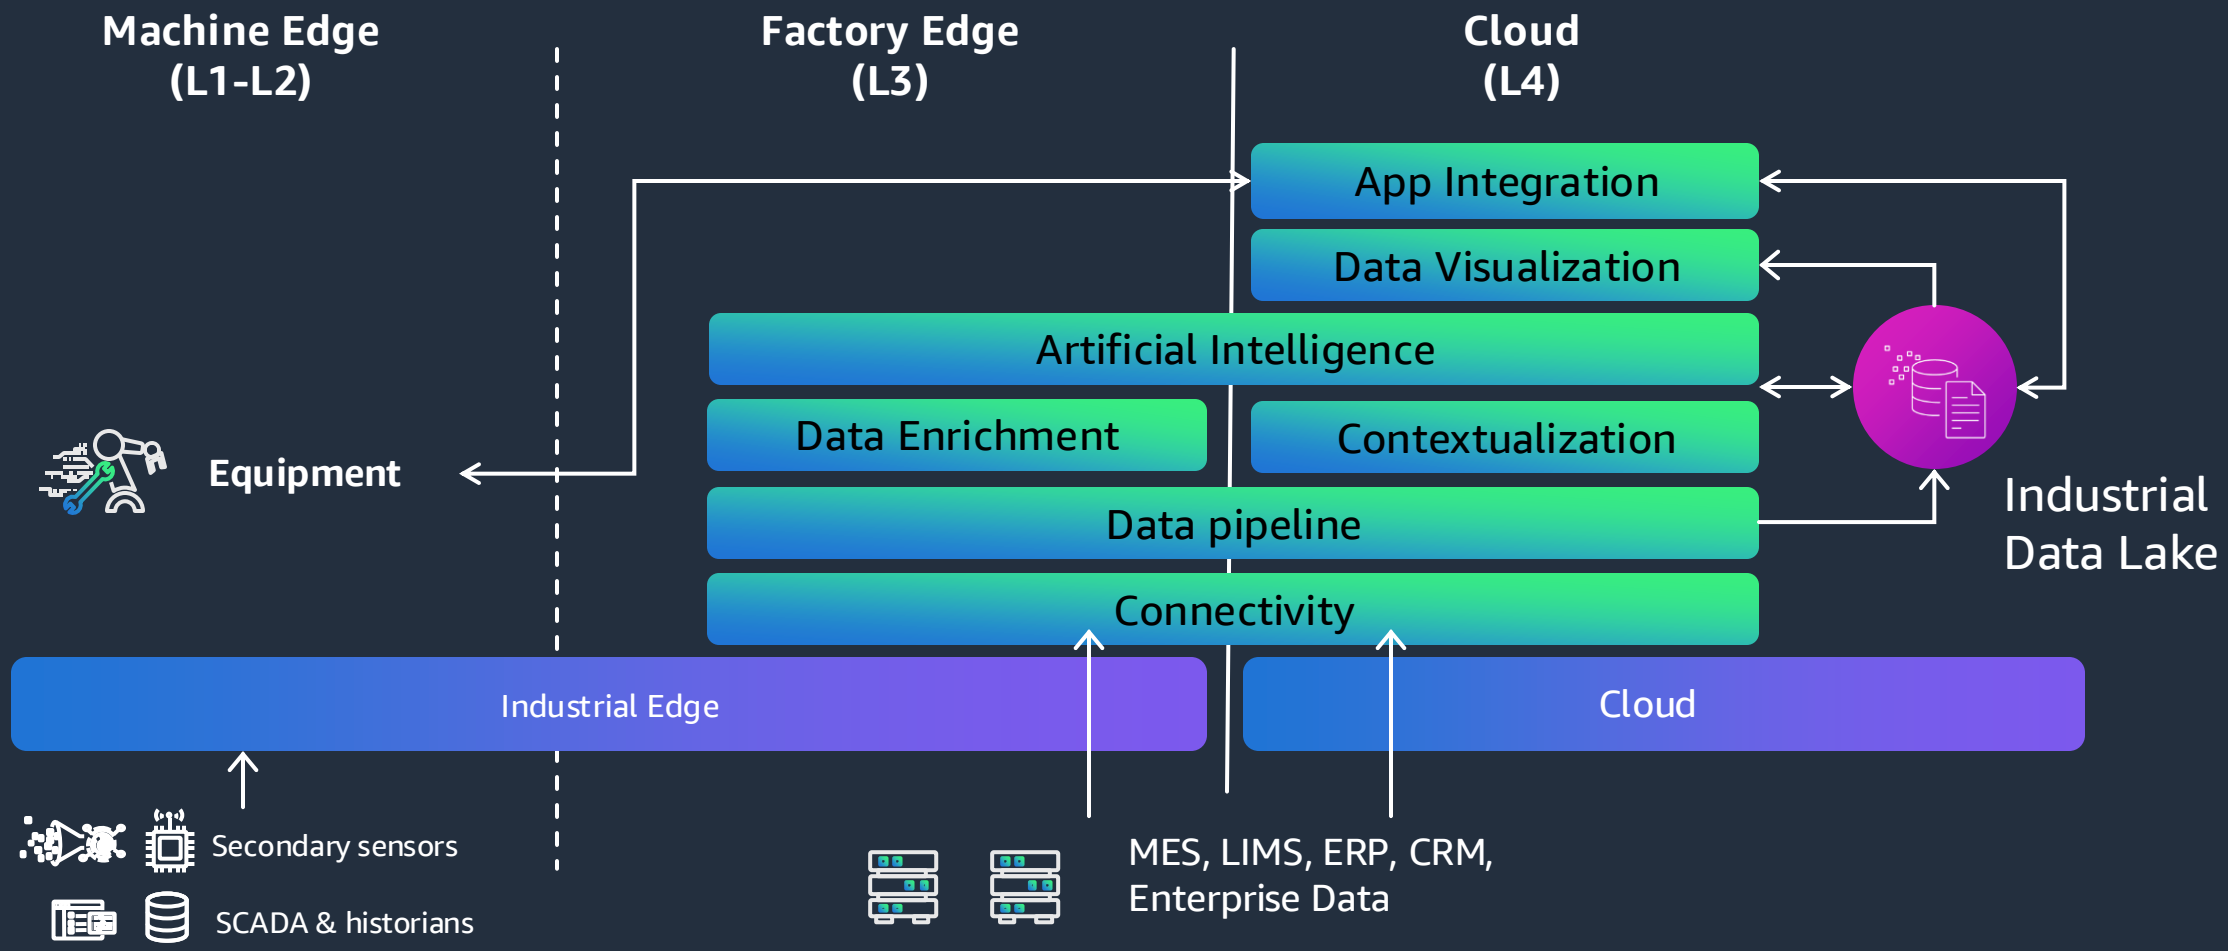

HEIDELBERG

## CHALLENGE

Heidelberg has been a reliable partner with great innovative strength for the global printing industry for many years.

Since the beginning of the digitalization, the printing industry is moving through a significant shift: In economic challenging environment, the printing industry is accelerating its digitization, which drives the automation in the industry helping to reduce manual work and make better use of resources and materials.

## SOLUTION

Heidelberg moved to the AWS Cloud and used AWS IoT Core and AWS IoT Device Management to support their Subscription Smart & Plus Service where Heidelberg provides equipment, workflow, matching consumables, training and consulting, and the service based on connected machinery data to power "pay-per-outcome" business model on print volume.

## IMPACT

Using AWS IoT Core and AWS IoT Device Management, Heidelberg helped their printing customers to boost productivity allowing them to remain competitive and accomplish print volume output at an increasing cost pressure level.

With more than 10,000 connected machines, a runtime of 2 seconds from the machine to the dashboard, Heidelberg saved up to 50% of operations costs of it's new data platform compared to the platform used before.

*"Key to your success lies in the coordinated interaction of people, processes, materials, and machines. We are a reliable and competent partner for all our customer's printing needs. AWS is supporting us to power our user-friendly systems, customized products, and innovative services to deliver maximum production quality, making controlling complex processes easier, and helping our customers to avoid errors."*

**Jochen Bender, Global Head of Contracts & Subscription, HEIDELBERG**

aws

# Accelerate time-to-value of machine digital services

## CHALLENGE

- Frontmatec, a leading food processing machine builder in the North American market, delivers digital services to its machines for improved performance and uptime.
- To accomplish this, Frontmatec needed to integrate and contextualize diverse data streams, for actionable insights, while at the same time, managing the solutions at the customer's facilities centrally, remotely and securely.

## SOLUTION

- To establish a centralized source for data-driven services across its machines, Frontmatec leveraged AWS IoT SiteWise Edge on Siemens Industrial Edge.
- They centrally deployed and managed Siemens Industrial Edge machine connectivity applications. They acquired machine data from various shopfloor assets, processed and compressed it at the edge, and securely transmitted it to the cloud via AWS IoT SiteWise Edge.
- Frontmatec employed asset modeling in AWS IoT SiteWise for structuring data across their machines, processing it to calculate metrics.

## IMPACT

- Accelerated deployment time from hours to 15 minutes
- Optimized compute costs per machine through the flexible edge setup
- Achieved scalability with repeatable configuration to expand from four to 400 machines across the US
- Empowered automation engineers with easy-to-deploy solution
- Delivered improved Overall Equipment Effectiveness (OEE) for end customers

SIEMENS

### AWS PARTNER: SIEMENS

Siemens Digital Industries (DI) is an innovation leader in automation and digitalization. Frontmatec used Siemens Industrial Edge, an open, ready-to-use edge computing platform with a central management system for a wide variety of industrial Edge Apps and Devices from 3rd party vendors and Siemens.

TOYOTA

# Predictive Maintenance Solution to Improve Assets Availability with AWS

## CHALLENGE

With more than 200 CNC machines per site running 24/7, preventative maintenance is time consuming and costly. The Maintenance Team had difficulty identifying high-priority tasks and work plans. The amount of data to analyze and required skills to implement predictive maintenance were too high.

## SOLUTION

AWS and Toyota co-developed a Predictive Maintenance solution capable of predicting failures days in advance, reducing unplanned downtime. The solution uses AWS IoT SiteWise and Amazon Lookout for Equipment to reduce data processing infrastructure and the need for data science skills.

## IMPACT

- Prevented 16 incidents and 20 hours of downtime (~$80k cost savings)
- Operational Availability in focus line improved 10% vs. previous 12-month avg.
- Solution scaled to 300+ assets across TMMK Powertrain; TMNA selected the solution for the region

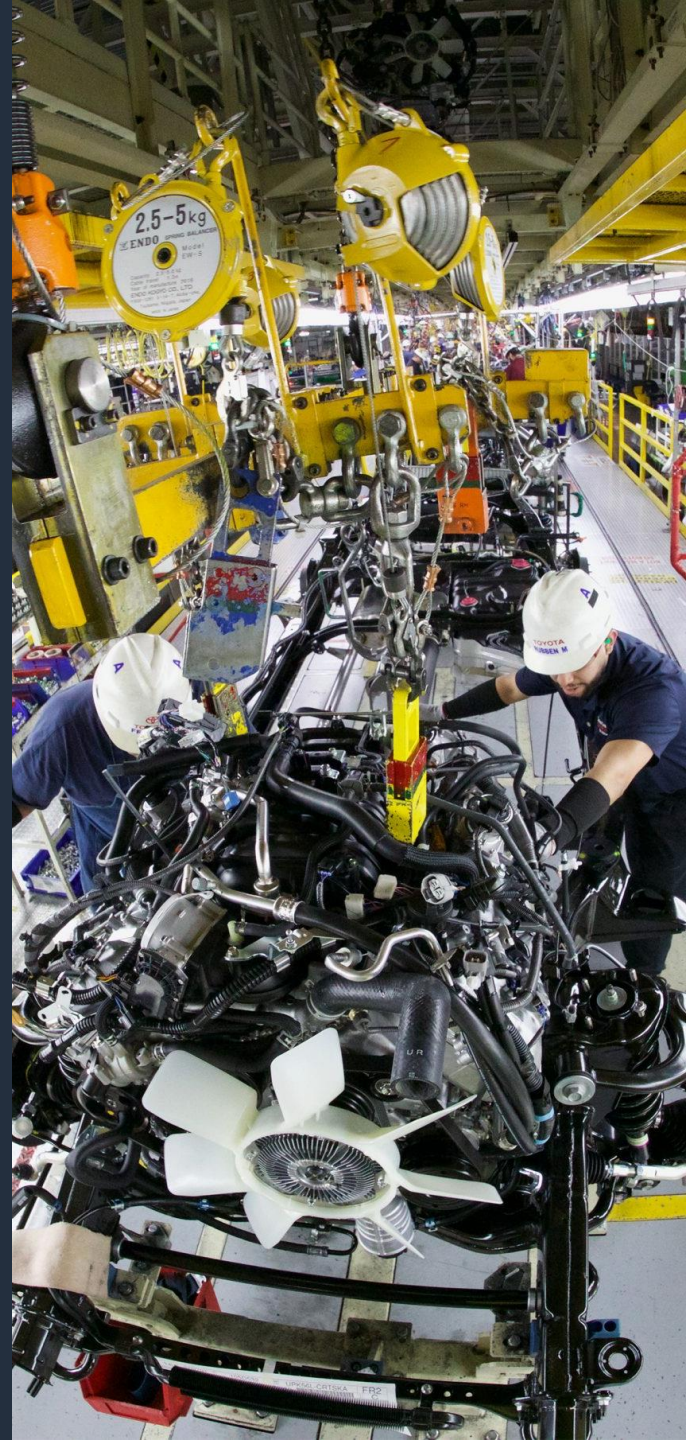


"The Operation Availability of our focus line was between 78-82%, incurring around 40 hours of downtime each month. With the help of AWS, we have found many problems in our machines, if left unnoticed would lead to critical failure. Now our OA is 92% and the downtime is around 20 hours!" - **Braden Burford, Sr. Maintenance Engineer, Toyota**

aws

## CHALLENGE

Founded in 1905 to solve the emerging famine in Europe, Yara has established a unique position as the industry's only global crop nutrition company.

To increase uptime, improve performance, and reduce waste in the value chain, Yara wanted to build a new Digital Production Platform (DPP) to digitize its production system, covering 24 sites, 122 production units and one mine.

## SOLUTION

AWS Professional Services worked side by side with Yara, to implement a new Digital Production Platform based on AWS IoT SiteWise to model production sites, collect, and store production data.

The Industrial Data Lake enables data processing, cataloging, access, and governance to support data-driven products, analytics, and applications.

## IMPACT

Within 9 months time, Yara, assisted by AWS Professional Services, was able to develop and deploy the new DPP to the first 6 production sites containing 50 production units and exposing hundreds of thousands datapoints on the platform.

This enables Yara to predict product quality and composition, improve balancing of the site utilities, and detect when machines need repair or maintenance to keep production at optimal efficiency levels.


"It is quite amazing what we have been able to do in just 9 months - this is what Yara has been talking about for years."

**Kyrre Tømmerberg, Director of Technology Digital Production, Yara International**

# Guidance for Smart Manufacturing on AWS

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

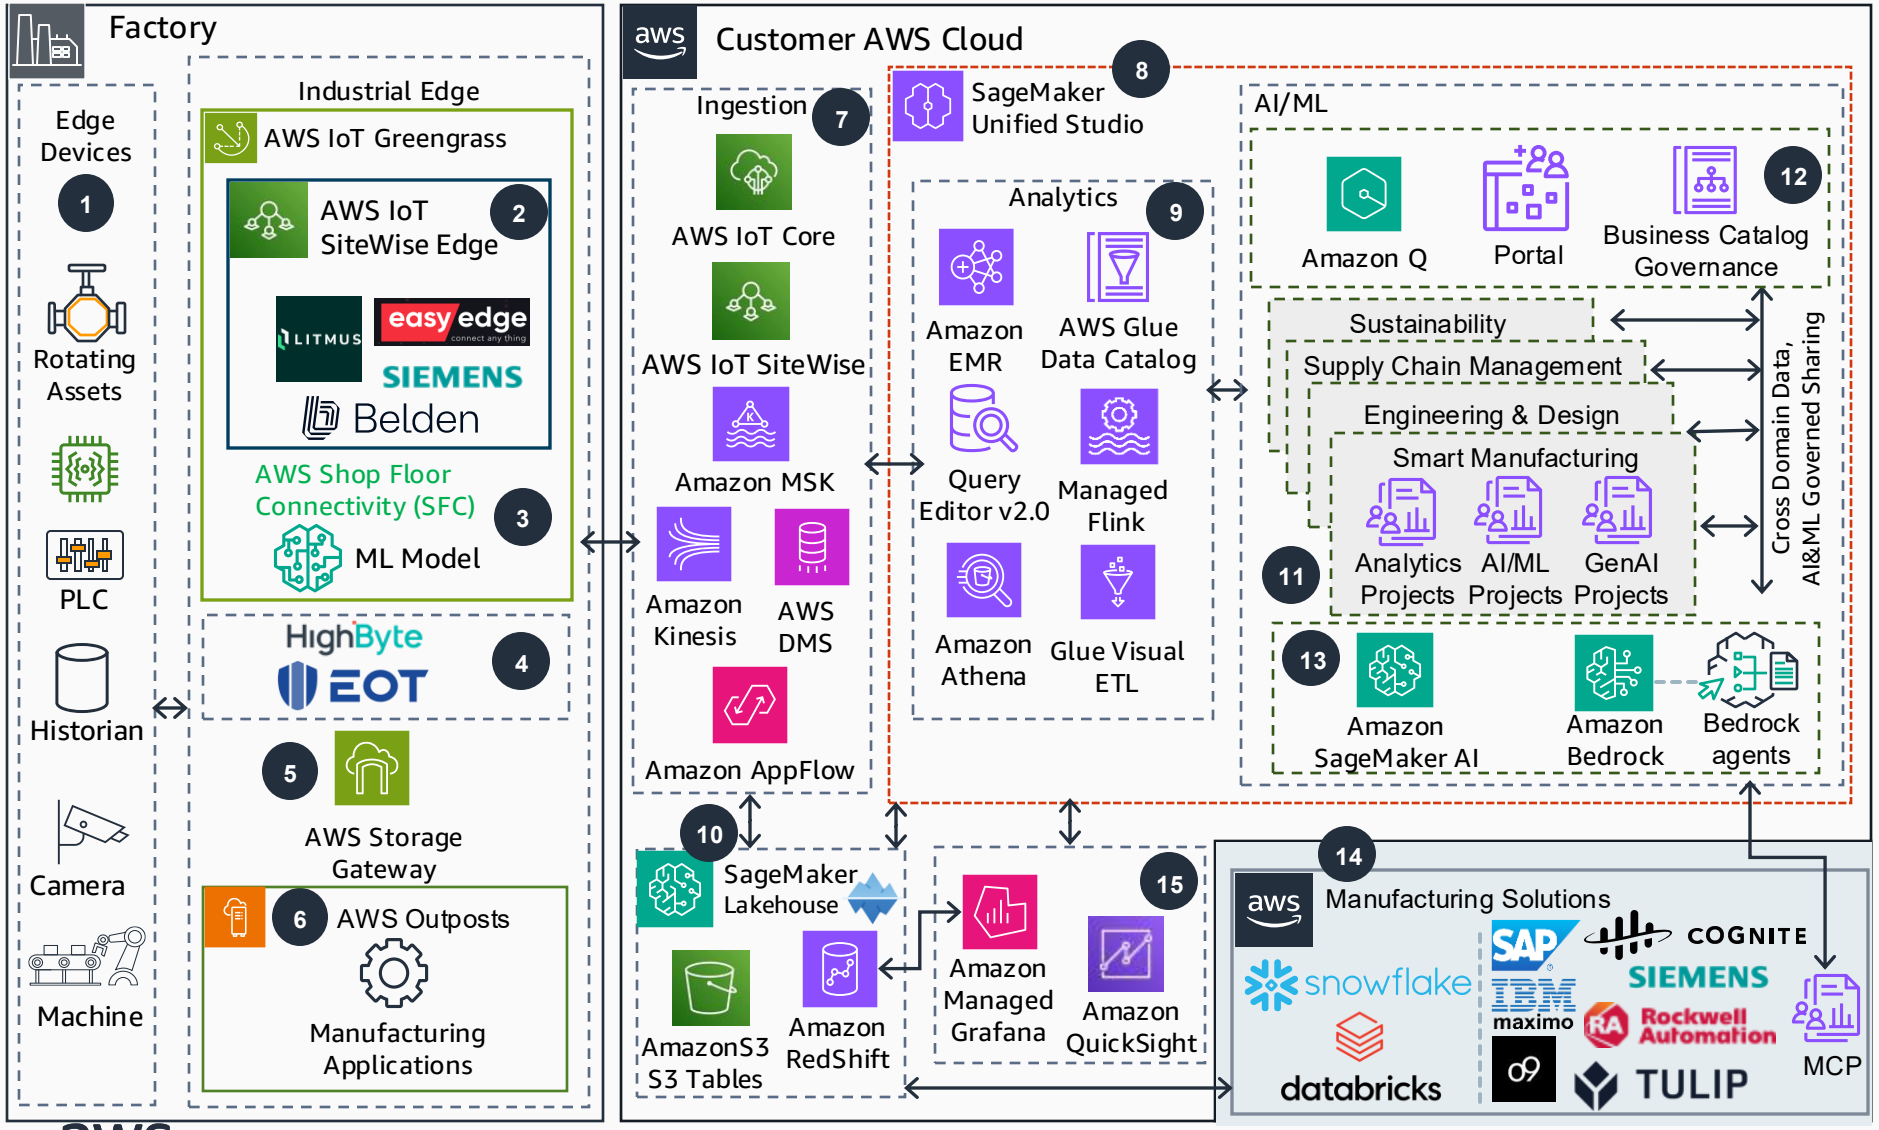


1. Identify information related to industrial activities from on-premise equipment.
2. Collect real-time data from edge devices and transmit data streams securely to **AWS IoT SiteWise** in the cloud. Leveraging partners like Litmus, Domatica's EasyEdge, Siemens Industrial Edge and Belden accelerate your integration with **AWS IoT SiteWise Edge**.
3. Connect to your edge devices through AWS Shop Floor Connectivity (SFC) Framework, an open-source solution from AWS, using multiple industrial protocols to securely stream data to AWS cloud services. Deploy and run your cloud-developed machine learning models at the edge through **AWS IoT Greengrass** for defect detection and anomaly inference.
4. Connect to edge devices, industrial protocol translators, and IT systems using AWS partner solutions like HighByte or Embassy of Things (EoT). These solutions contextualize your data and securely send it to AWS services such as **AWS IoT SiteWise**, **Amazon S3**, **AWS IoT Core**, **Amazon Kinesis**, and **Amazon MSK**.
5. Connect your on-premises applications to **Amazon S3** through **AWS Storage Gateway** using NFS and SMB file shares.
6. Extend AWS infrastructure and services to your premises with **AWS Outposts**, a fully managed service. Run your manufacturing-specialized applications on AWS services locally at your plant and integrate with AWS cloud infrastructure.

*Reviewed for technical accuracy October 21, 2024*



**AWS Reference Architecture**

# Guidance for Smart Manufacturing on AWS

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

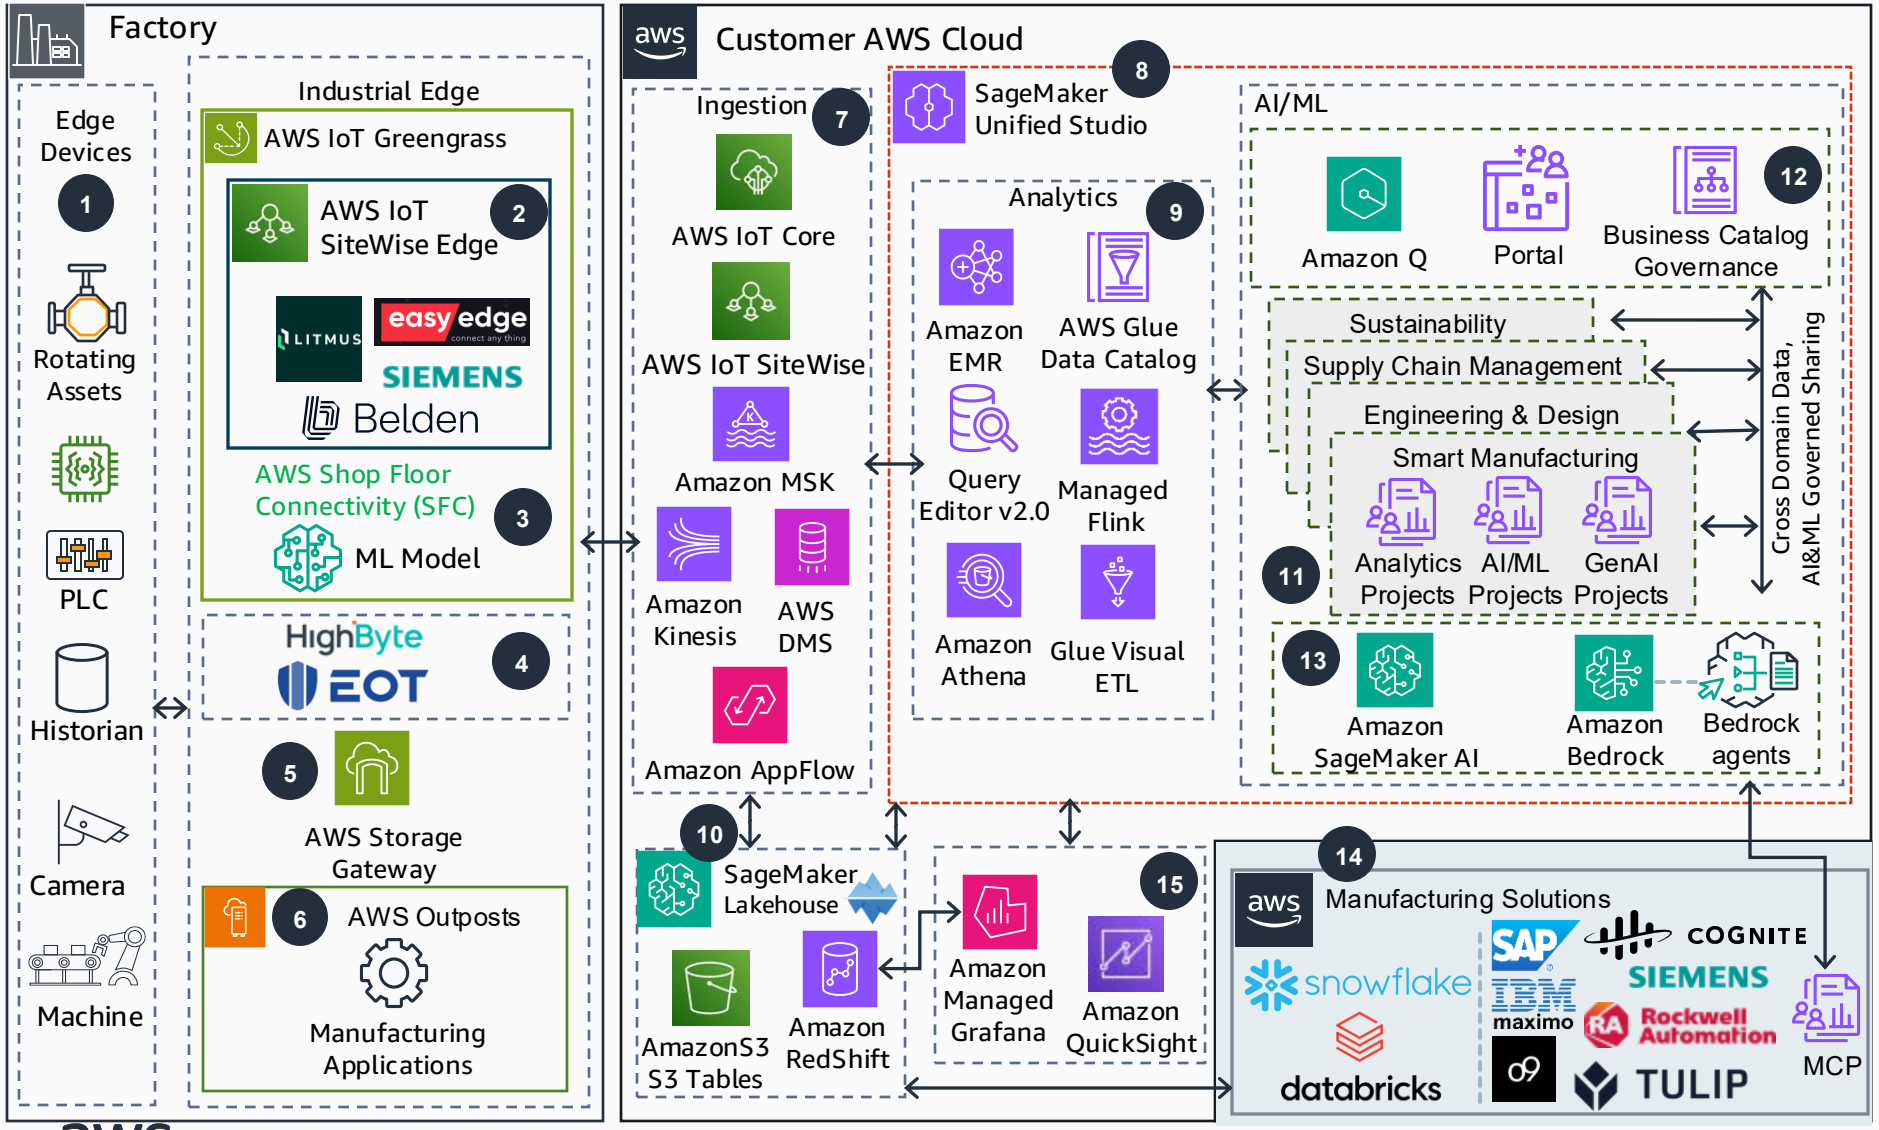


7 Process diverse data types through the data ingestion layer using AWS services. Stream real-time data through **AWS IoT SiteWise**, **AWS IoT Core**, **Amazon Kinesis**, and **Amazon MSK**. Transfer structured data from legacy on-premises platforms and data warehouses using **AWS DMS** and **AWS Glue**. Store and process unstructured and semi-structured data with **Amazon S3** and **Amazon AppFlow**. You can use **AppFlow** to extract, create and update data with ERPs, like SAP.

8 Access functionality and tools from AWS Analytics and AI/ML services through **Amazon SageMaker Unified Studio**'s single development environment. Find, access, and query data and AI assets across your organization. Collaborate on projects to build and share analytics and AI artifacts, including data, models, and generative AI applications, in a secure environment.

9 Process, model, and analyze your combined OT and IT data through AWS analytics services accessed via **Amazon SageMaker Unified Studio**'s Portal. Catalog your data using **AWS Glue Catalog**, transform it with Glue Visual ETL, and run SQL analytics through **Amazon Athena**. Process streaming data with **Managed Flink** and perform distributed data processing using **Amazon EMR** to generate business insights.

10 Unify your 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.

*Reviewed for technical accuracy October 21, 2024*



**AWS Reference Architecture**

# Guidance for Smart Manufacturing on AWS

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

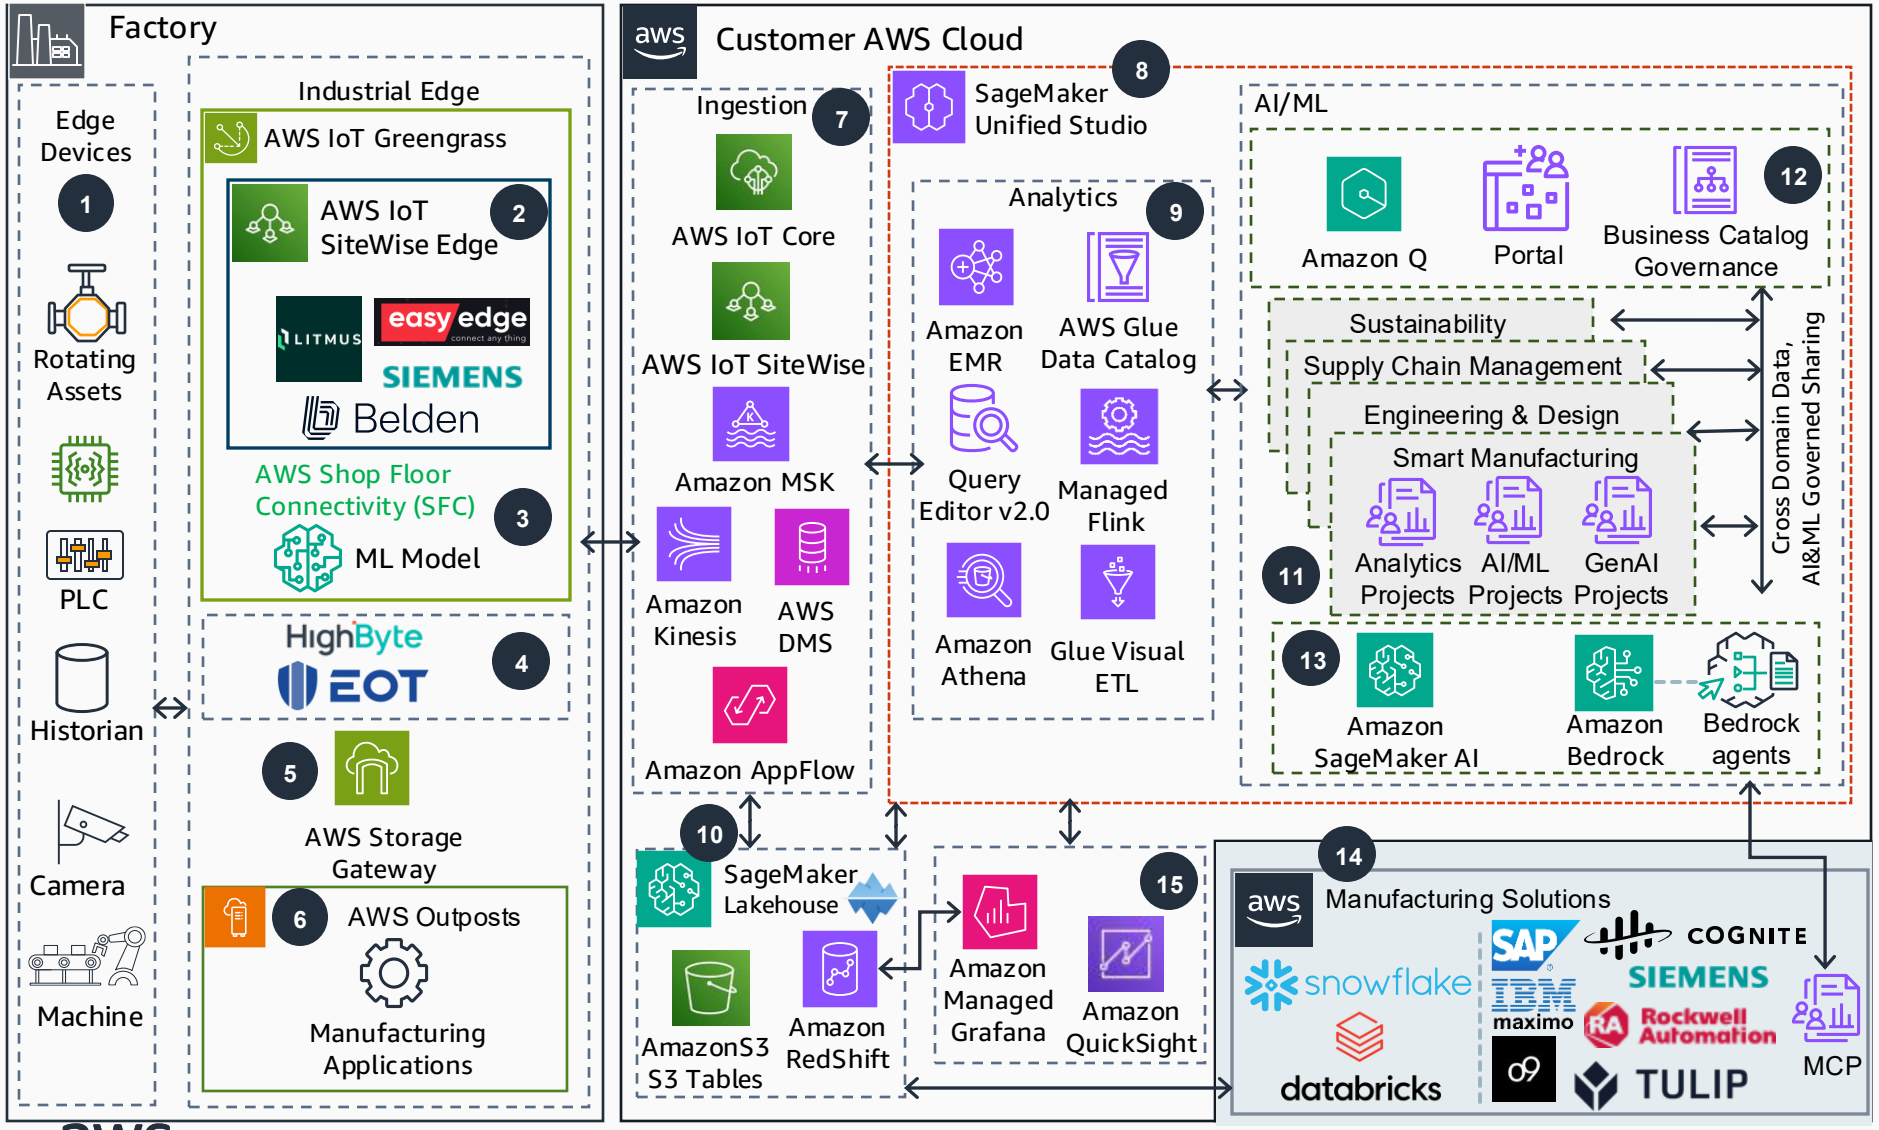


11 Organize your assets, users, and projects within **Amazon SageMaker Unified Studio** domains. Create single or multiple domains to match your enterprise structure. Collaborate in projects to manage data assets, analyze data, develop ML models, and build generative AI applications for specific business needs.

12 Enrich metadata from your technical catalogs with business context using **Amazon SageMaker Catalog**. Discover and access approved data and models through semantic search powered by generative AI. Monitor data quality, track lineage, and enforce access policies centrally in **Amazon SageMaker Unified Studio**.

13 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 & Design, all impacting your Sustainability.

14 Integrate with cloud-hosted manufacturing solutions (ERP, Supply Chain, Maintenance, WMS/TMS), including **MCP** servers for industrial knowledge sources. Exchange data with enterprise data platforms like Snowflake and Databricks to enhance your manufacturing analytics capabilities.

15 Visualize your data with **Amazon Managed Grafana** natively from **Amazon RedShift** or from **S3** via **Amazon Athena**. Build dashboards with **Amazon QuickSight** via **Amazon Athena**.

*Reviewed for technical accuracy October 21, 2024*



**AWS Reference Architecture**

# Guidance for Deploying Smart Machines on AWS

## Connect and Manage Machines

This architecture diagram shows the process of connecting smart machines, remotely managing them, and constructing an industrial data management layer. The following slides show further details on building a data foundation and managing the device lifecycle.

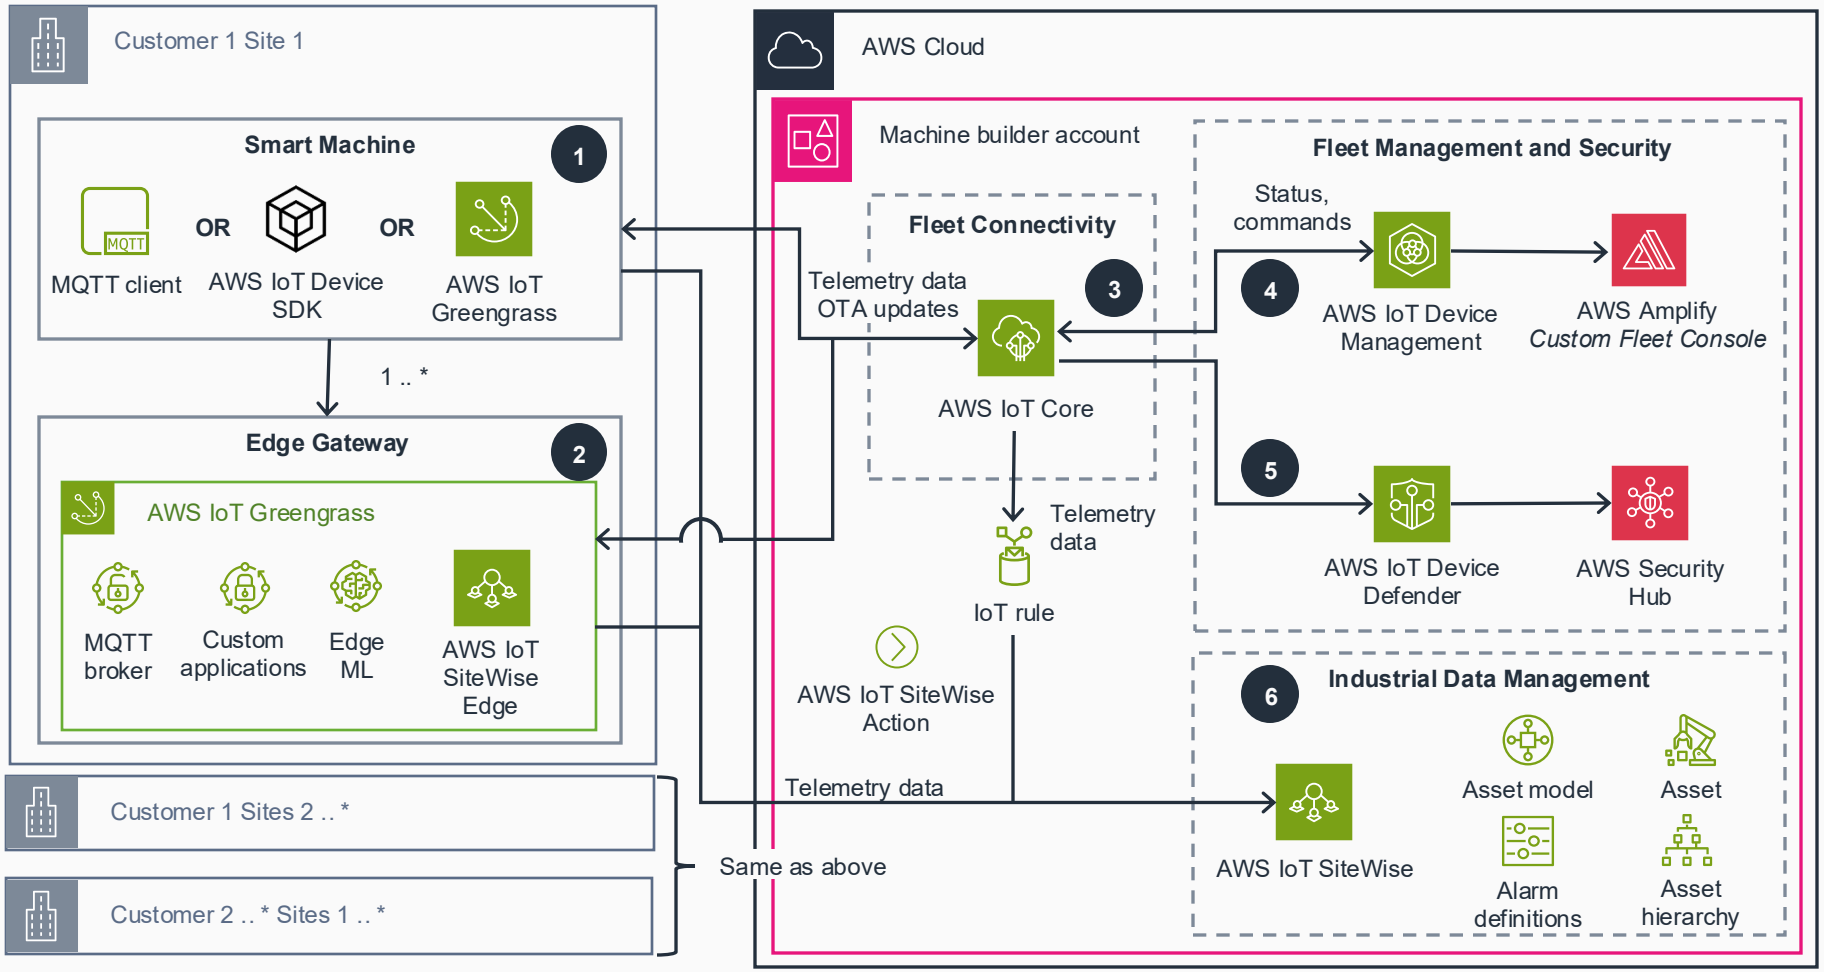


1. A smart machine connects to **AWS IoT Core** using a Message Queuing Telemetry Transport (MQTT) client, AWS IoT Device SDK, or the edge runtime provided by **AWS IoT Greengrass**. The telemetry data is then ingested into **AWS IoT SiteWise** directly through **AWS IoT SiteWise Edge** or through **AWS IoT Core**.

2. If the machine lacks direct internet connectivity, use an edge gateway as a cloud connectivity layer. The edge gateway collects data from the machines, data historians, applications, then processes, stores, and forwards it to the AWS Cloud. Run custom applications and ML inferences at the edge.

3. Facilitate scalable two-way communication between machines or edge gateways and the AWS Cloud, without the need to manage infrastructure, using **AWS IoT Core**.

4. Remotely provision, monitor, update, and troubleshoot machines or edge gateways by leveraging **AWS IoT Device Management**. Build a custom fleet management console using **AWS Amplify** to visualize your fleet, and search across it to view machine state and health data.

5. Audit your fleet for compliance with security best practices and continuously monitor it using **AWS IoT Device Defender**. Any security findings are sent to **AWS Security Hub** for a centralized view of all security issues from various AWS services.

6. Ingest and contextualize operational data from your machines using **AWS IoT SiteWise** data streams and modeling capabilities. Additionally, compute performance metrics, store timeseries data, create alarm definitions, and provide flexible data access to external applications.

*Reviewed for technical accuracy October 21, 2024*



**AWS Reference Architecture**

# Guidance for Deploying Smart Machines on AWS

## Build an Industrial Data Foundation

This architecture diagram demonstrates how the industrial data foundation can enable operations monitoring, alarm notifications, AI/ML models, business intelligence dashboards and reports, AI assistants, APIs, lifecycle management—empowering contact center agents with contextual machine information.

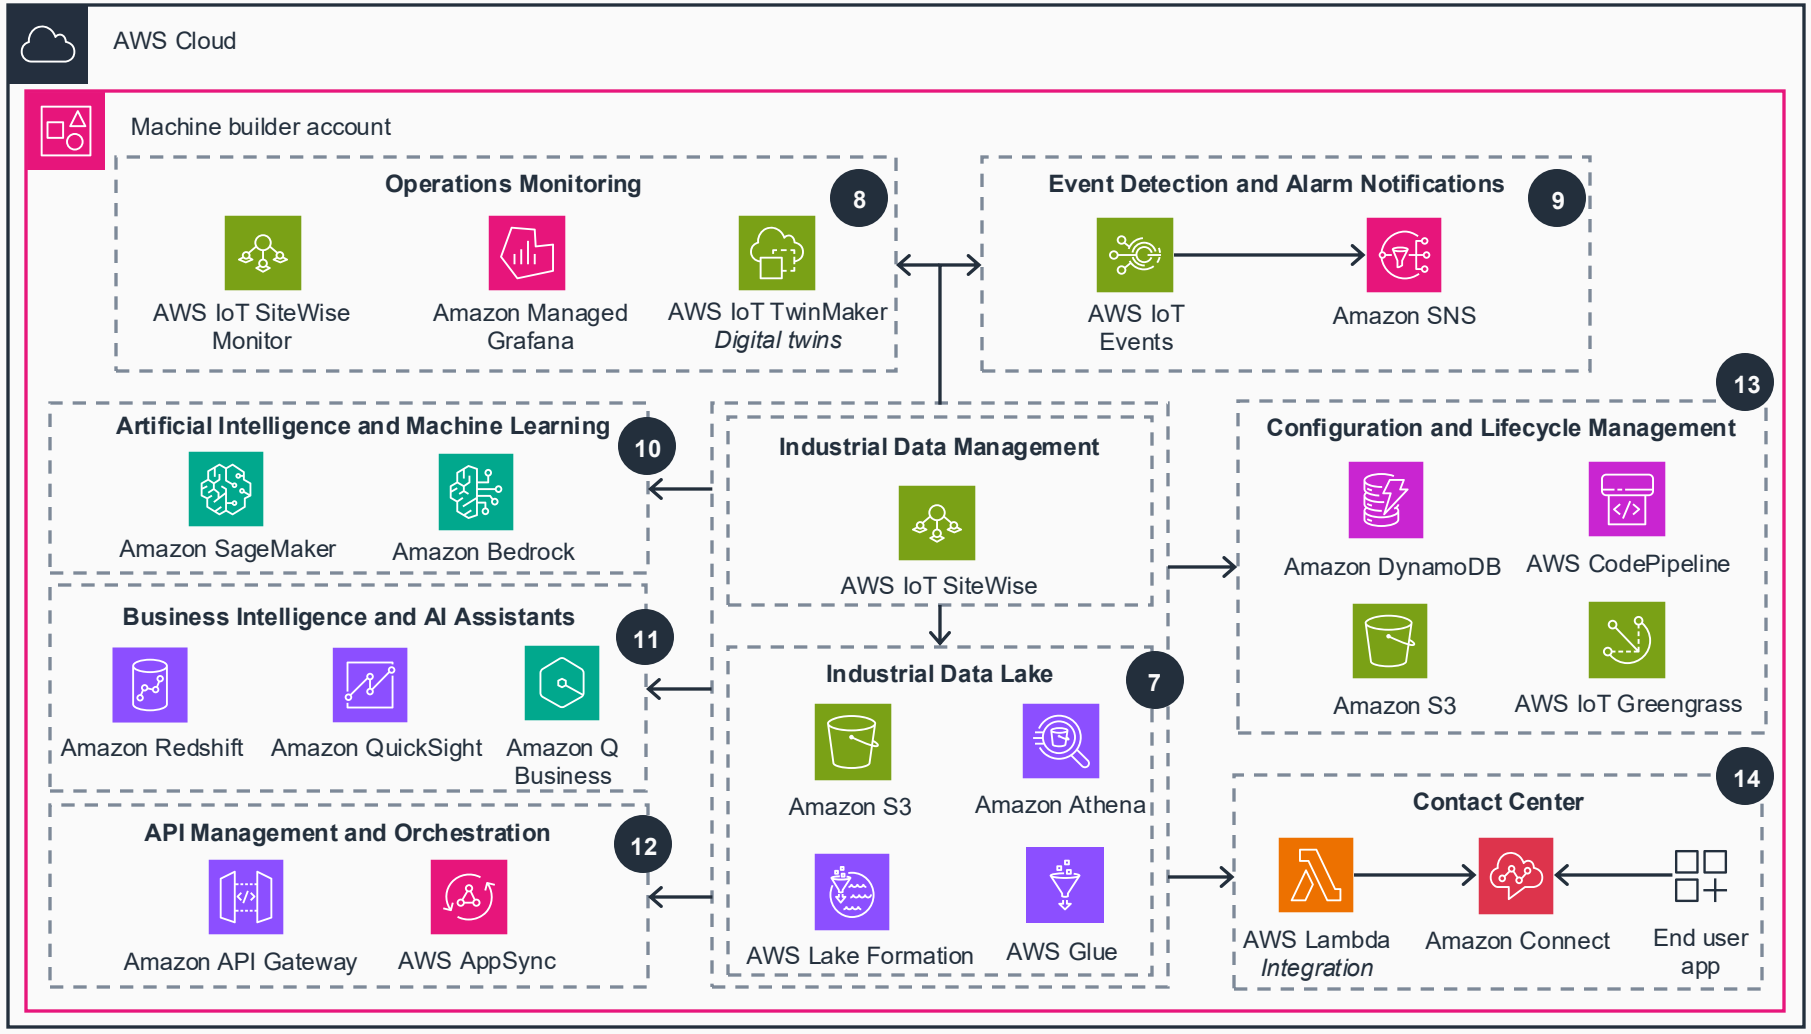




**AWS Reference Architecture**

7 Build an industrial data lake using the contextual data from **AWS IoT SiteWise**. Govern, secure, and share data using **AWS Lake Formation** for advanced analytics. Catalog and analyze data with services like **AWS Glue** and **Amazon Athena**.

8 Remotely monitor machines using **AWS IoT SiteWise** Monitor or with **Amazon Managed Grafana** for rich, contextual dashboards. Build digital twins powered by **AWS IoT TwinMaker** to improve equipment performance.

9 Notify operational personnel about the health of machines using **AWS IoT Events** and **Amazon Simple Notification Service** **(Amazon SNS)**. Create state machines and event monitoring applications with **AWS IoT Events**.

10 Develop AI/ML solutions for predictive maintenance with **Amazon SageMaker** and build generative AI solutions using **Amazon Bedrock**.

11 **Amazon QuickSight** enables data-driven decisions. With the **Amazon Q** add-on, business users can ask natural language queries for quick insights. Empower employees with enterprise information using **Amazon Q Business**.

12 Provide historical and real-time machine data to customers by building serverless APIs using **Amazon API Gateway** and **AWS AppSync** that can scale to millions of users.

13 Use **Amazon DynamoDB** for storing machine configuration, **AWS CodePipeline** for automating continuous integration and continuous delivery (CI/CD), **Amazon Simple Storage Service** **(Amazon S3)** for storing artifacts, and **AWS IoT Greengrass** for managing edge devices.

14 Leverage **Amazon Connect** to meet customer service needs and empower agents with contextual machine information.

# Guidance for Deploying Smart Machines on AWS

## DevOps Lifecycle Management

This architecture diagram illustrates the process of enhancing machine capabilities and resolving issues through over-the-air (OTA) updates, leveraging an automated CI/CD pipeline that involves various stages of development, including build, test and deployment. This DevOps lifecycle helps close the loop to quickly respond to customer needs in the market.

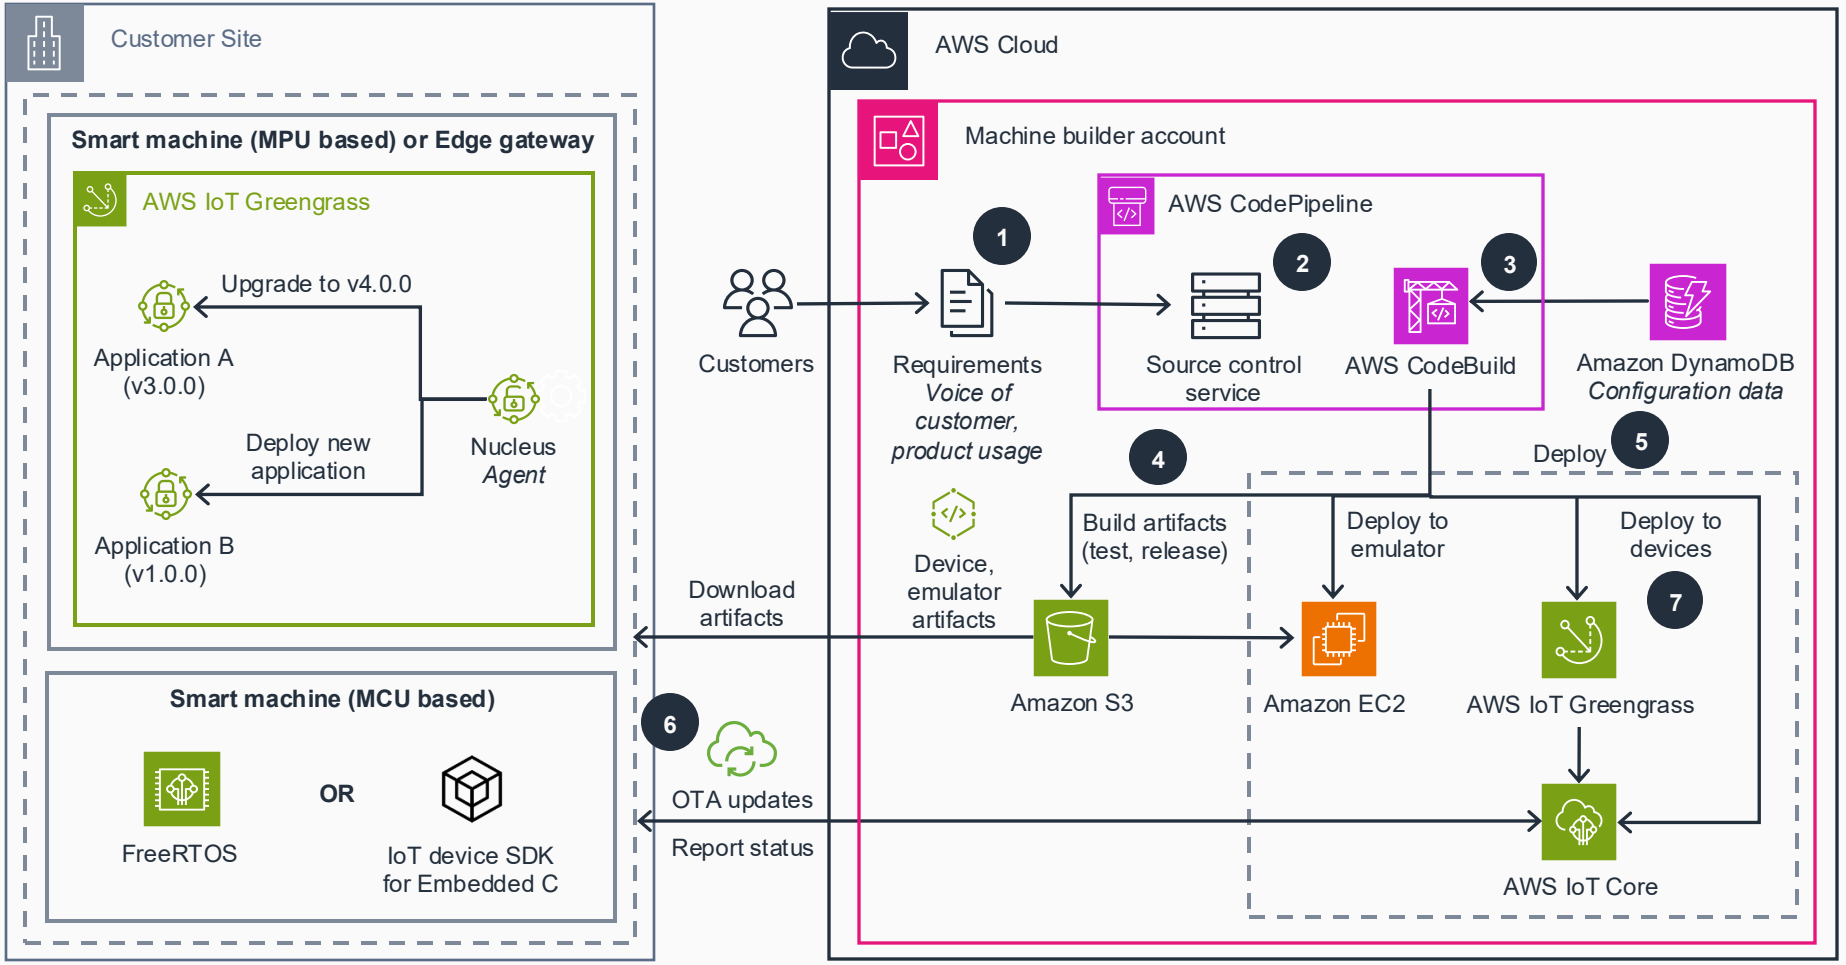


1. The machine builder gathers requirements through Voice of Customer feedback and product usage analysis in an effort to enhance machine capabilities or resolve ongoing issues.
2. Software developers and embedded developers make changes to the source code hosted by source control services such as GitHub, GitLab, and Bitbucket.
3. Leverage **AWS CodeBuild** with cross-build tools to create artifacts for devices and emulators. **DynamoDB** provides the necessary machine-specific configuration. **CodePipeline** automates the CI/CD process by orchestrating various stages of development.
4. Store the artifacts meant for testing and production release securely in **Amazon S3**.
5. Test the artifacts by deploying them to emulated environment and a test group of physical devices. Emulated environments can created using emulators such as Quick Emulator (QEMU) and Arm Virtual Hardware (AVH) on **Amazon Elastic Compute Cloud (Amazon EC2)**. Use thing groups from **AWS IoT Core** to organize the test devices for testing.
6. Devices receive over-the-air (OTA) updates from **AWS IoT Core** and securely download the necessary artifacts from **Amazon S3** using pre-signed URLs or MQTT file streams. They then update the firmware or software and report the status back to **AWS IoT Core**. The machine builder verifies the update for improved security, usability, reliability, and functionality and then approves it.
7. Deploy approved artifacts to all devices with configurable rollout rates and schedules, and monitor continuously during and after deployment.

*Reviewed for technical accuracy October 21, 2024*



**AWS Reference Architecture**

# Thank you!

Steve Blackwell
/in/smblackwell