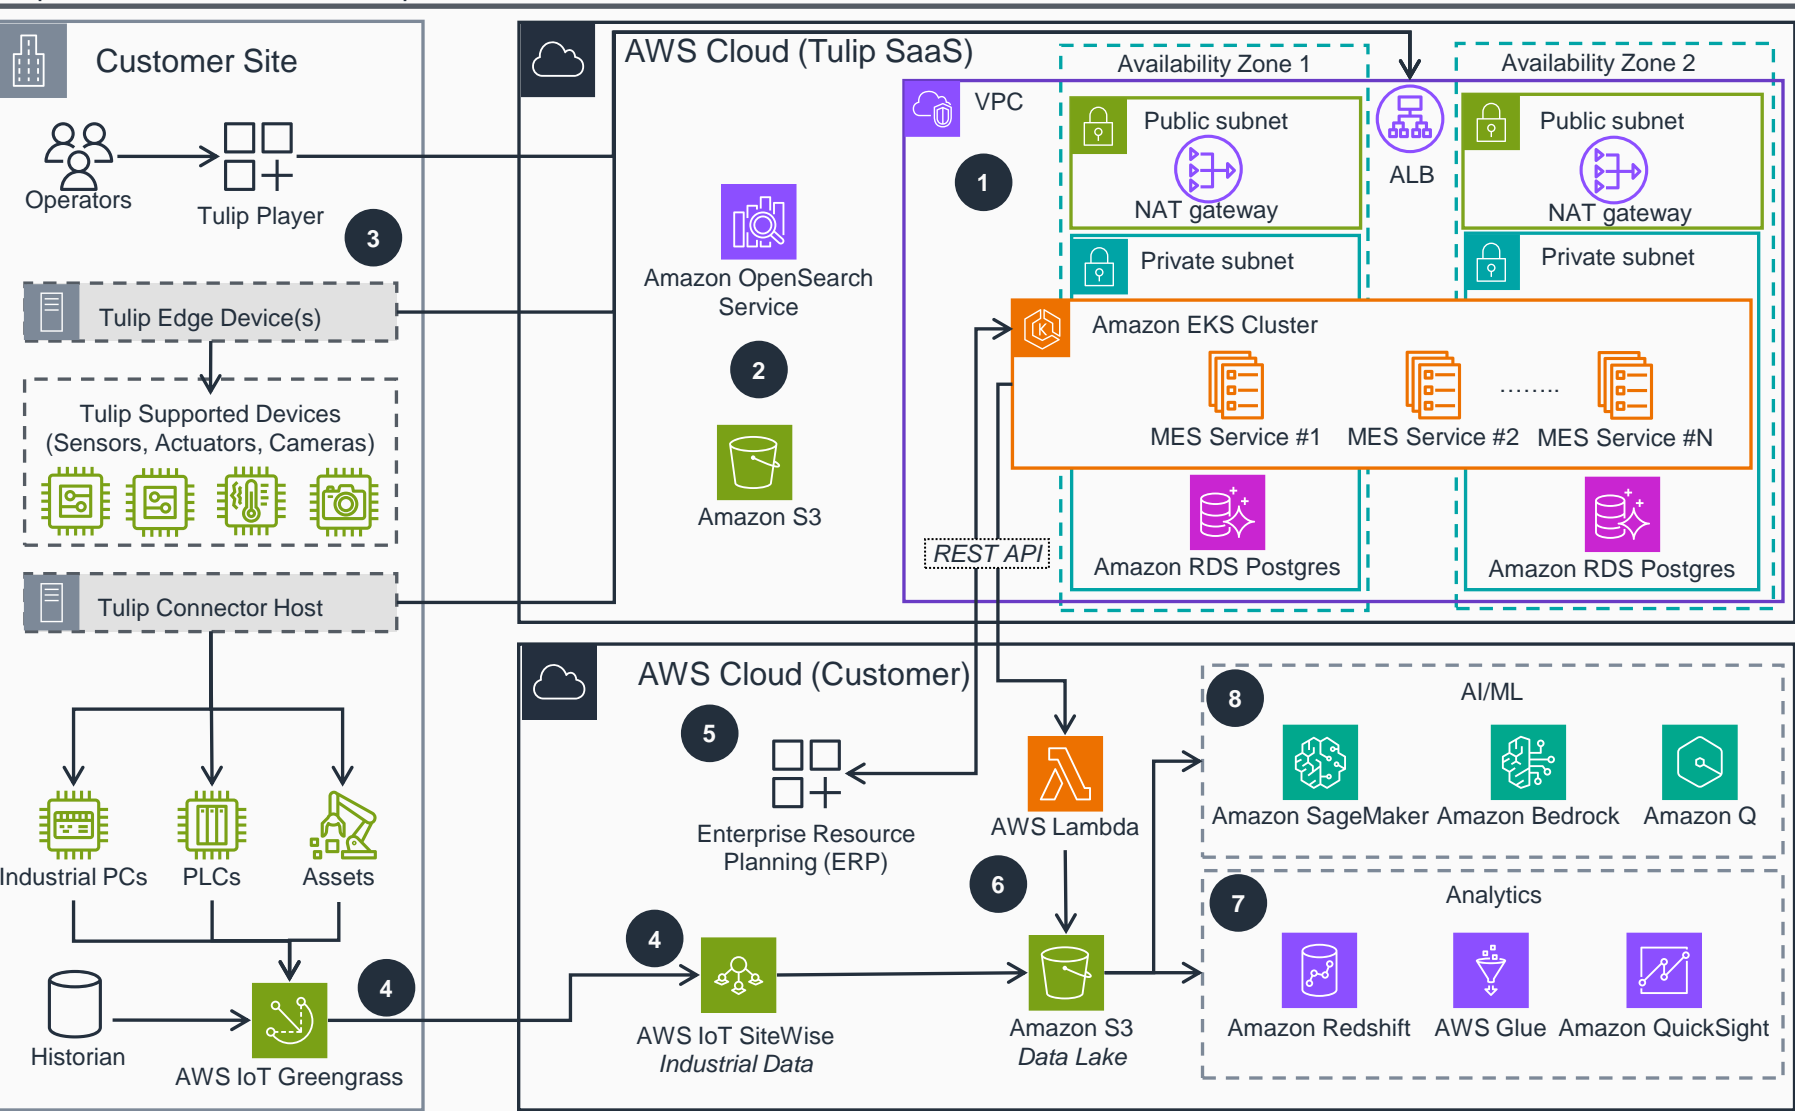


# Guidance for Tulip Manufacturing Execution System (MES) on AWS

This architecture diagram showcases comprehensive data flows between a Tulip Manufacturing Execution System (MES) and AWS. The Tulip MES is offered as a software as a service (SaaS) on AWS, enabling data transparency and generating actionable insights. Steps 1-5 are shown below; steps 6-8 are shown on the next slide.



1 Tulip MES on AWS is designed to be a highly scalable, multi-tenant SaaS solution. At the core is an **Amazon Elastic Kubernetes Service** (Amazon EKS) cluster, deployed across multiple Availability Zones for high availability. **Amazon RDS for PostgreSQL**, in a high availability configuration, is used to store application data. The **Application Load Balancer** (ALB) serves as an entry point for the application, forwarding the requests to the appropriate Kubernetes service with the right tenant context. A **NAT gateway** in the public subnet allows application and database servers to reach the internet.

2 **Amazon Simple Storage Service** (Amazon S3) is used to store customer specific files with each customer having their own set of **AWS Identity and Access Management** (IAM) credentials to access their isolated **Amazon S3** buckets. **Amazon OpenSearch Service** is used for log aggregation and analysis.

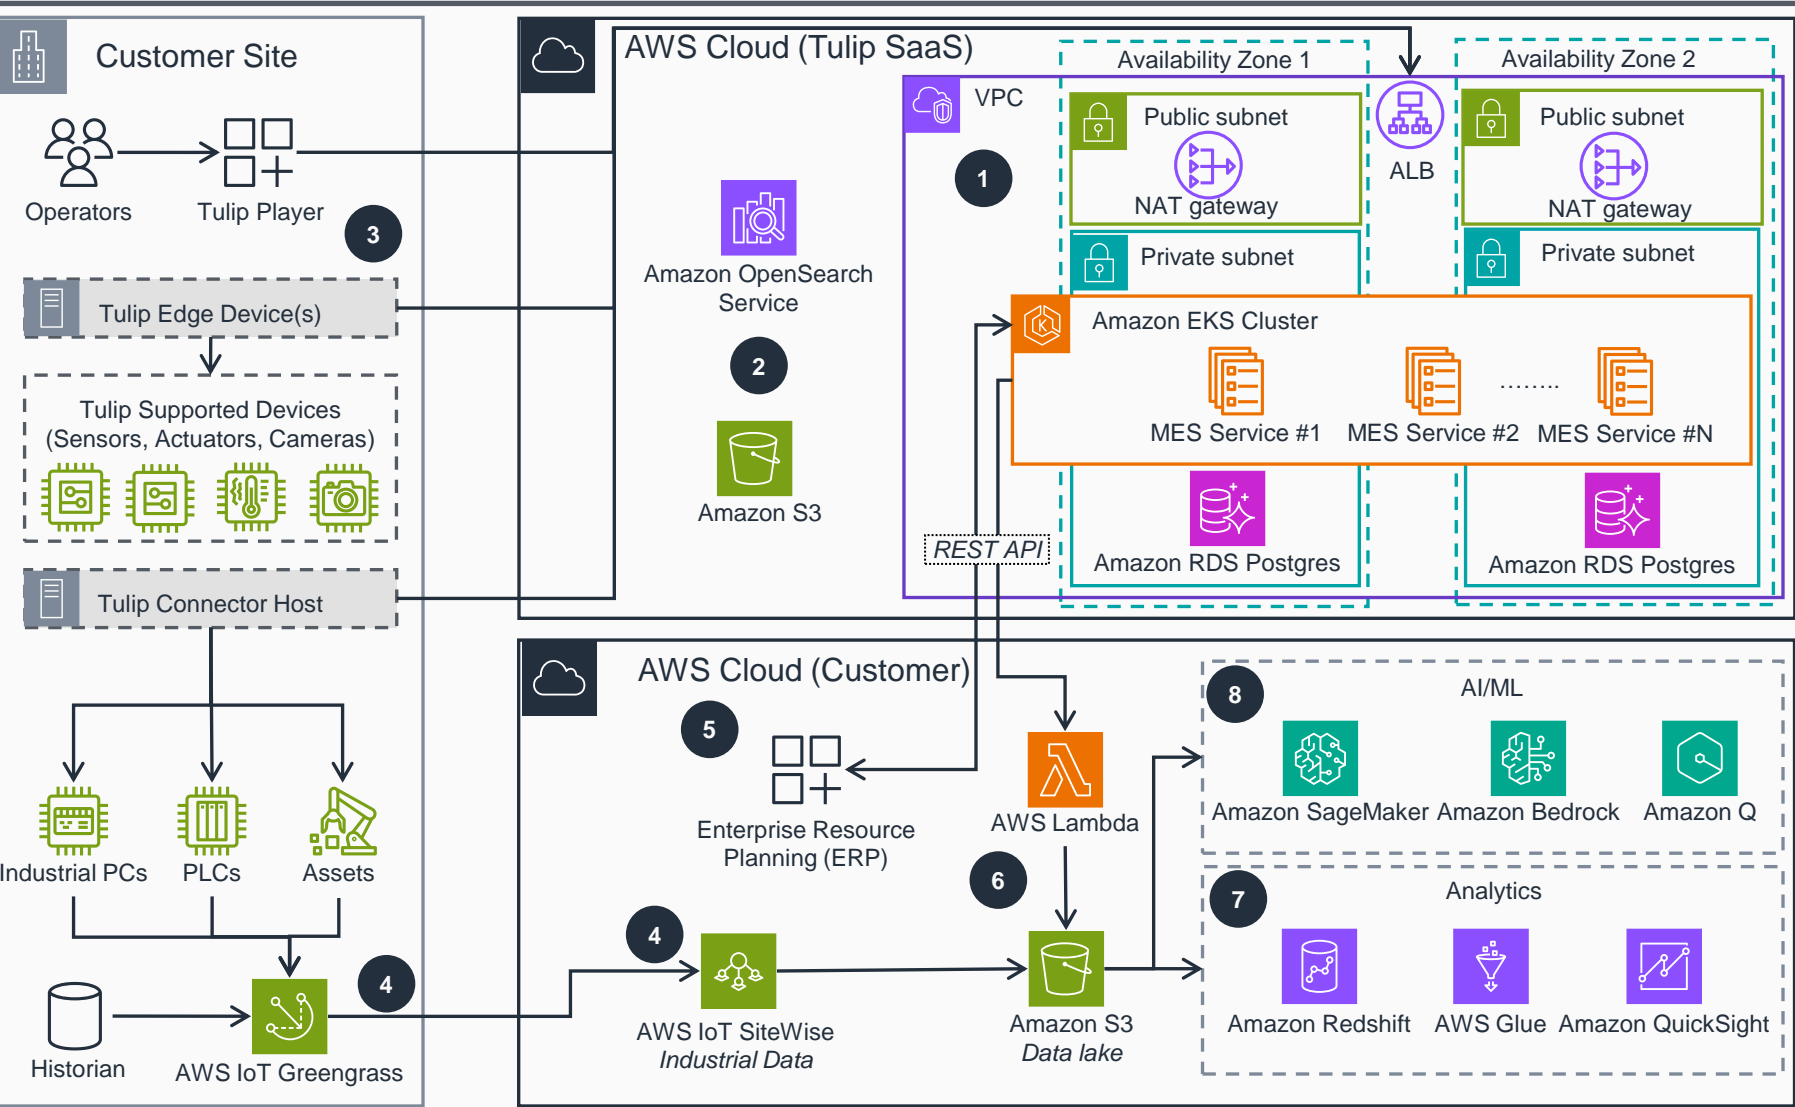
3 Operators at customer sites access Tulip through an application known as the Tulip Player. The Tulip platform also supports connectivity between select industrial devices, such as sensors, actuators, and cameras directly through Tulip edge device(s). Additionally, industrial assets, Programmable Logic Controllers (PLCs), and other external systems can be integrated with the Tulip platform using the Tulip Connector Host.

4 **AWS IoT Greengrass** can also connect in parallel to industrial assets, Programmable Logic Controllers (PLCs), and historian systems to collect and store time series industrial data to **AWS IoT SiteWise**.

5 AWS customers can integrate their Enterprise Resource Planning (ERP) application to Tulip using REST API.

# Guidance for Tulip Manufacturing Execution System (MES) on AWS

Steps 6-8



- 6 Customers can also create an **Amazon S3** data lake in their own account for data analytics. Data from the Tulip application can be extracted and transferred to the **Amazon S3** data lake using **AWS Lambda** to call Tulip's REST API. Time series data from **AWS IoT SiteWise** can also be exported to the **Amazon S3** data lake.
- 7 Customers can extract value from industrial data stored in the **Amazon S3** data lake using various AWS analytics services. For data warehousing, they can use **Amazon Redshift**. To manage the data catalog, they can use **AWS Glue**. And for business intelligence capabilities, they can take advantage of **Amazon QuickSight**.
- 8 Customers can use the data stored in the **Amazon S3** data lake to build, train, and deploy machine learning models using **Amazon SageMaker**. They can also use **Amazon Bedrock** to develop generative artificial intelligence (AI) applications. **Amazon Q Business** can be configured to obtain relevant answers to pressing questions, solve problems, generate content, and take action on the user's behalf.

