


Guidance for Collaborative, Unified Data and AI Development on AWS

Overview

This architecture diagram shows how Amazon SageMaker provides a unified, collaborative experience for ML and data engineers, data stewards, and generative AI developers to accelerate data applications, from exploration to production.




- 1** **Amazon SageMaker** allows you to configure domains. Each domain provides a single and isolated portal for data teams to collaborate.
- 2** **SageMaker** offers a single sign-on (SSO) experience, allowing individual users or groups to access a particular domain. You can either use an identity provider (IdP) in AWS or connect your existing IdP using SAML.
- 3** Once you authenticate, **SageMaker** redirects to your assigned domain portal. This portal offers a common space for you to collaborate on projects. In this Guidance, you will use a sales forecasting project where a data engineer, ML engineer, and data analyst can work together.
- 4** In **SageMaker**, you can create projects backed by GitHub repositories, enabling user collaboration. The platform provides version control, a browser-based integrated development environment (IDE), and a low-code/no-code experience for solution development, with access to native generative AI services.
- 5** Data engineers transform datasets, such as for sales forecasting, before model development. The backend resources offer tools and infrastructure for data warehouses, engineering, pipelines, governance, ML, and generative AI. The data catalog, lineage, and pub/sub empower user collaboration on trusted, governed data.
- 6** The sales dataset is stored in persistent storage. In order to allow access to data, native services are used to implement a Lakehouse architecture for the project to use.
- 7** A studio (web-based IDE) experience is provided to allow both data and ML engineers to perform data transformations in addition to model training and validation operations.
- 8** **SageMaker** unifies data analytics tools, enabling you to build, deploy, and monitor data applications through a cohesive experience. The suite includes a Query Editor, **Amazon SageMaker** JumpStart, **SageMaker** Endpoints, **SageMaker** Workflows, and JupyterLab IDE.



Guidance for Collaborative, Unified Data and AI Development on AWS

Generative AI Lakehouse

This architecture diagram shows how Amazon SageMaker enables a collaborative data engineering and analytics experience for sales forecasting using a Lakehouse architecture, web-based studio with generative AI, and orchestration tools in a unified portal.




- 1** **AWS IAM Identity Center** manages user access and SSO to **Amazon SageMaker Unified Studio** for data engineers.
- 2** **SageMaker Unified Studio** allows data engineers and data analysts to collaborate on the sales forecasting project.
- 3** Store the sales dataset for the forecasting project in an **Amazon Simple Storage Service (Amazon S3)**-backed data Lakehouse architecture. Use **AWS Glue** for data cataloging and **Amazon Redshift Serverless** for fast data retrieval. Govern the data using **AWS Lake Formation** permissions and access it through the Iceberg API, enabling seamless integration between the **Amazon S3** and **Amazon Redshift Serverless** tiers.
- 4** **SageMaker Unified Studio** provides a web-based interface to the data engineer, allowing them to perform the necessary transformation to the sales dataset without needing to leave **SageMaker** or swap consoles. You can use **Amazon Q Developer** to provide in-place AI-generated coding recommendations.
- 5** The studio IDE interface in **SageMaker** allows you to leverage **Amazon Athena** and **Amazon Redshift** for data exploration and heavy data transformations, respectively. You can store the sales dataset and intermediate datasets resulting from data transformations in **Amazon S3**.
- 6** The Workflows tool automates the end-to-end process of data ingestion using **Amazon Managed Workflows for Apache Airflow (Amazon MWAA)** for orchestration and **AWS Glue** for task execution. The Query Editor in **SageMaker** provides a SQL notebook-style interface to write, run, and save queries against data sources in **Amazon Redshift** and the **AWS Glue Data Catalog**, allowing users to upload and view sample data.

Guidance for Collaborative, Unified Data and AI Development on AWS

Collaborative model deployment

This architecture diagram shows how Amazon SageMaker empowers ML engineers to collaboratively develop, evaluate, and deploy sales forecasting models using Amazon SageMaker, SageMaker JumpStart, and SageMaker Workflows within a unified portal.



- 1 **IAM Identity Center** manages user access and SSO to **SageMaker Unified Studio** for ML engineers.
- 2 **SageMaker Unified Studio** allows your ML engineer to collaborate on the sales forecasting project.
- 3 **SageMaker Unified Studio** provides a web-based interface to your ML engineer to train and validate the model using the dataset previously prepared and curated by the data engineer. You can use **Amazon Q Developer** to provide in-place AI-generated coding recommendations.
- 4 Use **SageMaker** compute to train and validate the model with the datasets prepared by your data engineers.
- 5 Use the **SageMaker** JumpStart tool to evaluate pre-trained models in **SageMaker** and the **SageMaker** Endpoints tool to deploy the final sales forecasting model for online predictions in **SageMaker**. Use the **SageMaker** Workflows tool to automate the end-to-end model training and deployment.
- 6 **SageMaker** stores the model artifacts in **Amazon S3**.

