Guidance for Collaborative, Unified Data and AI Development on AWS

Overview – SageMaker Unified Studio AIOps

This architecture diagram illustrates a multi-account AIOps framework with centralized governance, enabling systematic model progression across Dev, Test, and Prod environments through automated pipelines. This slide shows steps 1-4.

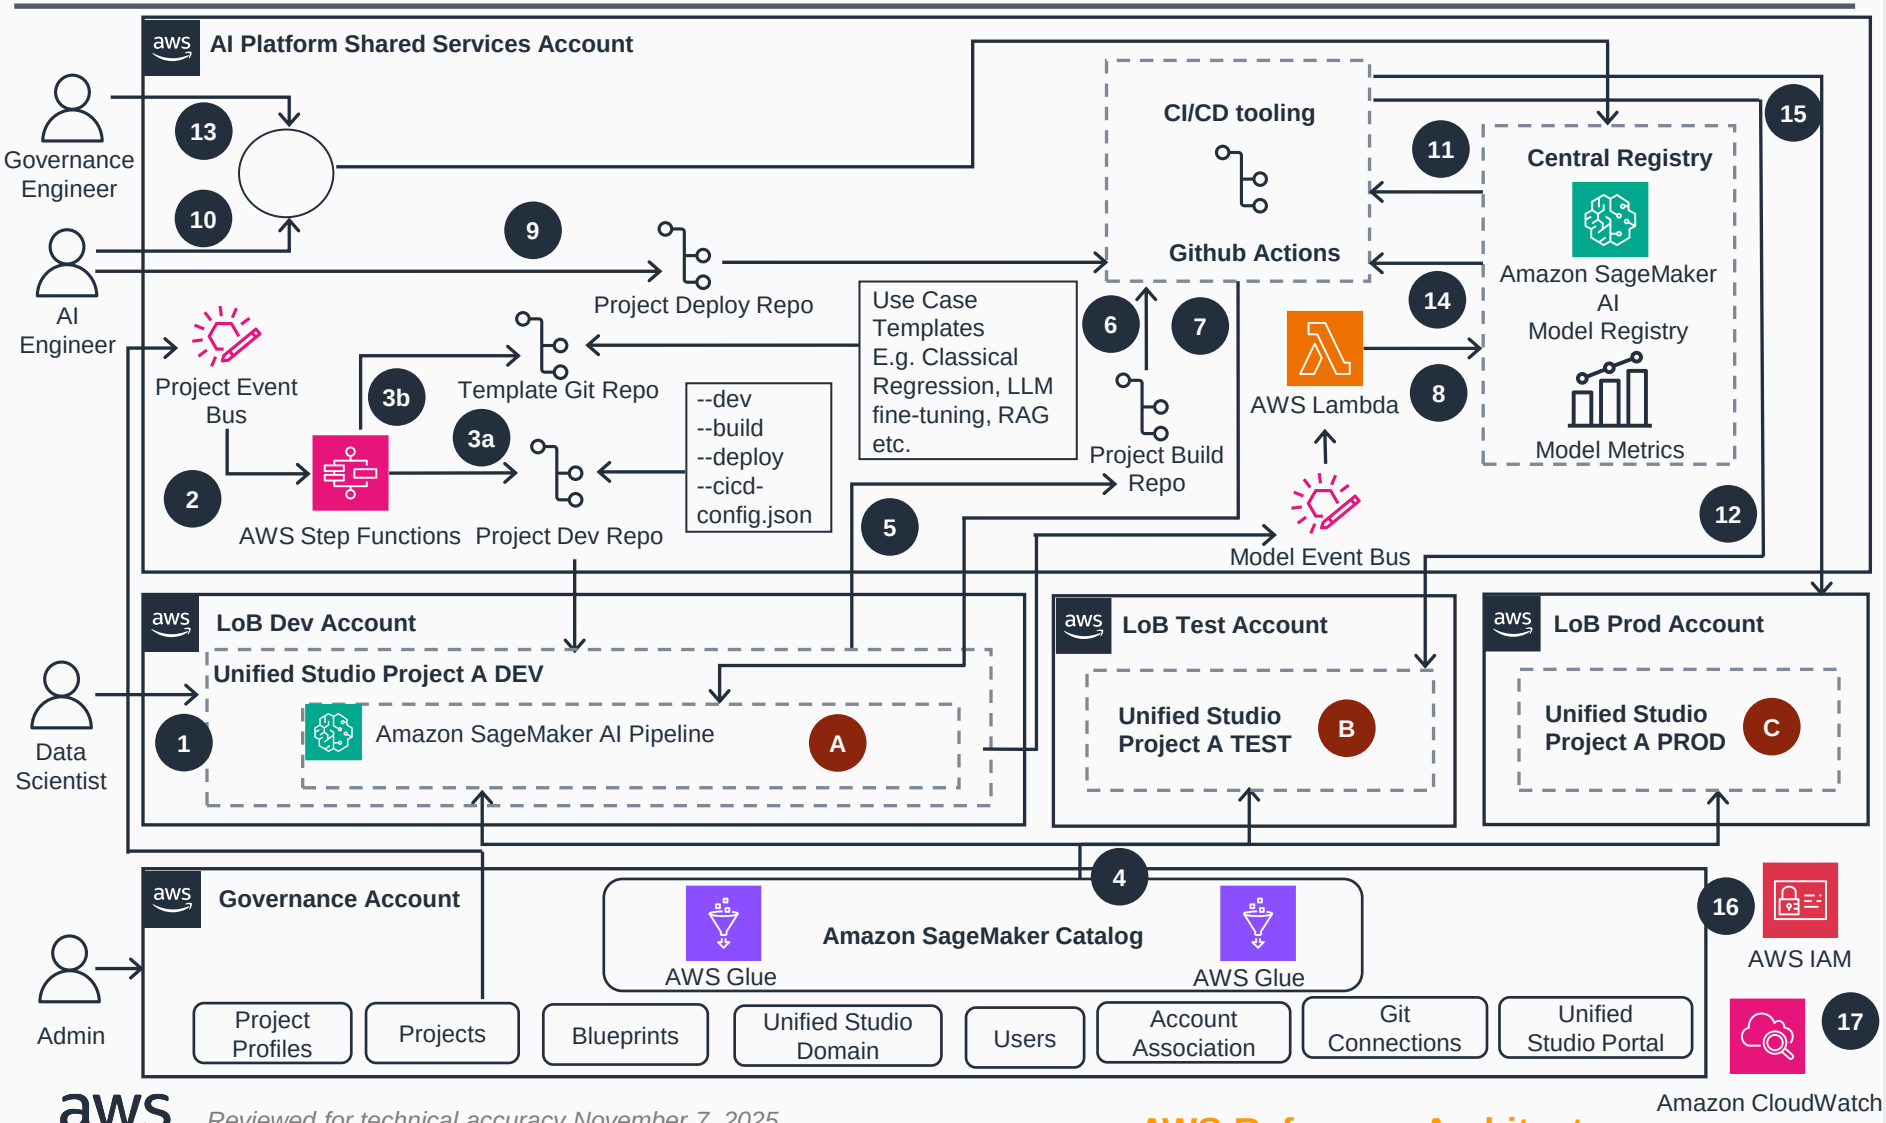

1. A Data Scientist creates Project A – DEV in **Amazon SageMaker** Unified Studio. This action publishes a CreateProject event captured by **AWS CloudTrail** and bridged to a custom **Amazon EventBridge** bus in the AI Shared Services account. Using the specified Git connection, a Git repository is created.

2. The CreateProject event is delivered to the Shared Services **Amazon EventBridge** bus, which serves as the integration hub for automation workflows across accounts.

3a. **Amazon EventBridge** triggers the Project Git Setup **AWS Lambda** function, which creates the dedicated project Git repository to host build and deployment assets.

3b. An **AWS Step Functions** project Git setup workflow provisions the corresponding use case templates (such as Classical Regression, LLM fine-tuning, RAG). These templates set up the repo's build and deploy folders with standardized pipeline code. Build and deploy can be kept in separate repositories, with mappings defined in cicd-config.json.

4. The **Amazon SageMaker** Catalog in the Governance Account serves as a centralized data registry, enabling secure discovery and access to approved data assets across accounts. It provides enterprise-wide management of data and model assets, with integrated controls for data access and sharing across different environments.

The Data Scientist subscribes to required datasets (e.g., **AWS Glue** tables) through **Amazon SageMaker** Catalog, which is an enterprise-wide business catalog within **Amazon SageMaker**. Once the data producer approves, the assets become available in the **Amazon SageMaker** Catalog for use in experiments. The **Amazon SageMaker** Catalog integrates with **AWS Glue** tables for structured data and supports **Amazon Simple Storage Service (Amazon S3)** Object collections for unstructured data.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture

Guidance for Collaborative, Unified Data and AI Development on AWS

Overview – SageMaker Unified Studio AIOps

This architecture diagram illustrates a multi-account AIOps framework with centralized governance, enabling systematic model progression across Dev, Test, and Prod environments through automated pipelines. This slide shows steps 5-12.

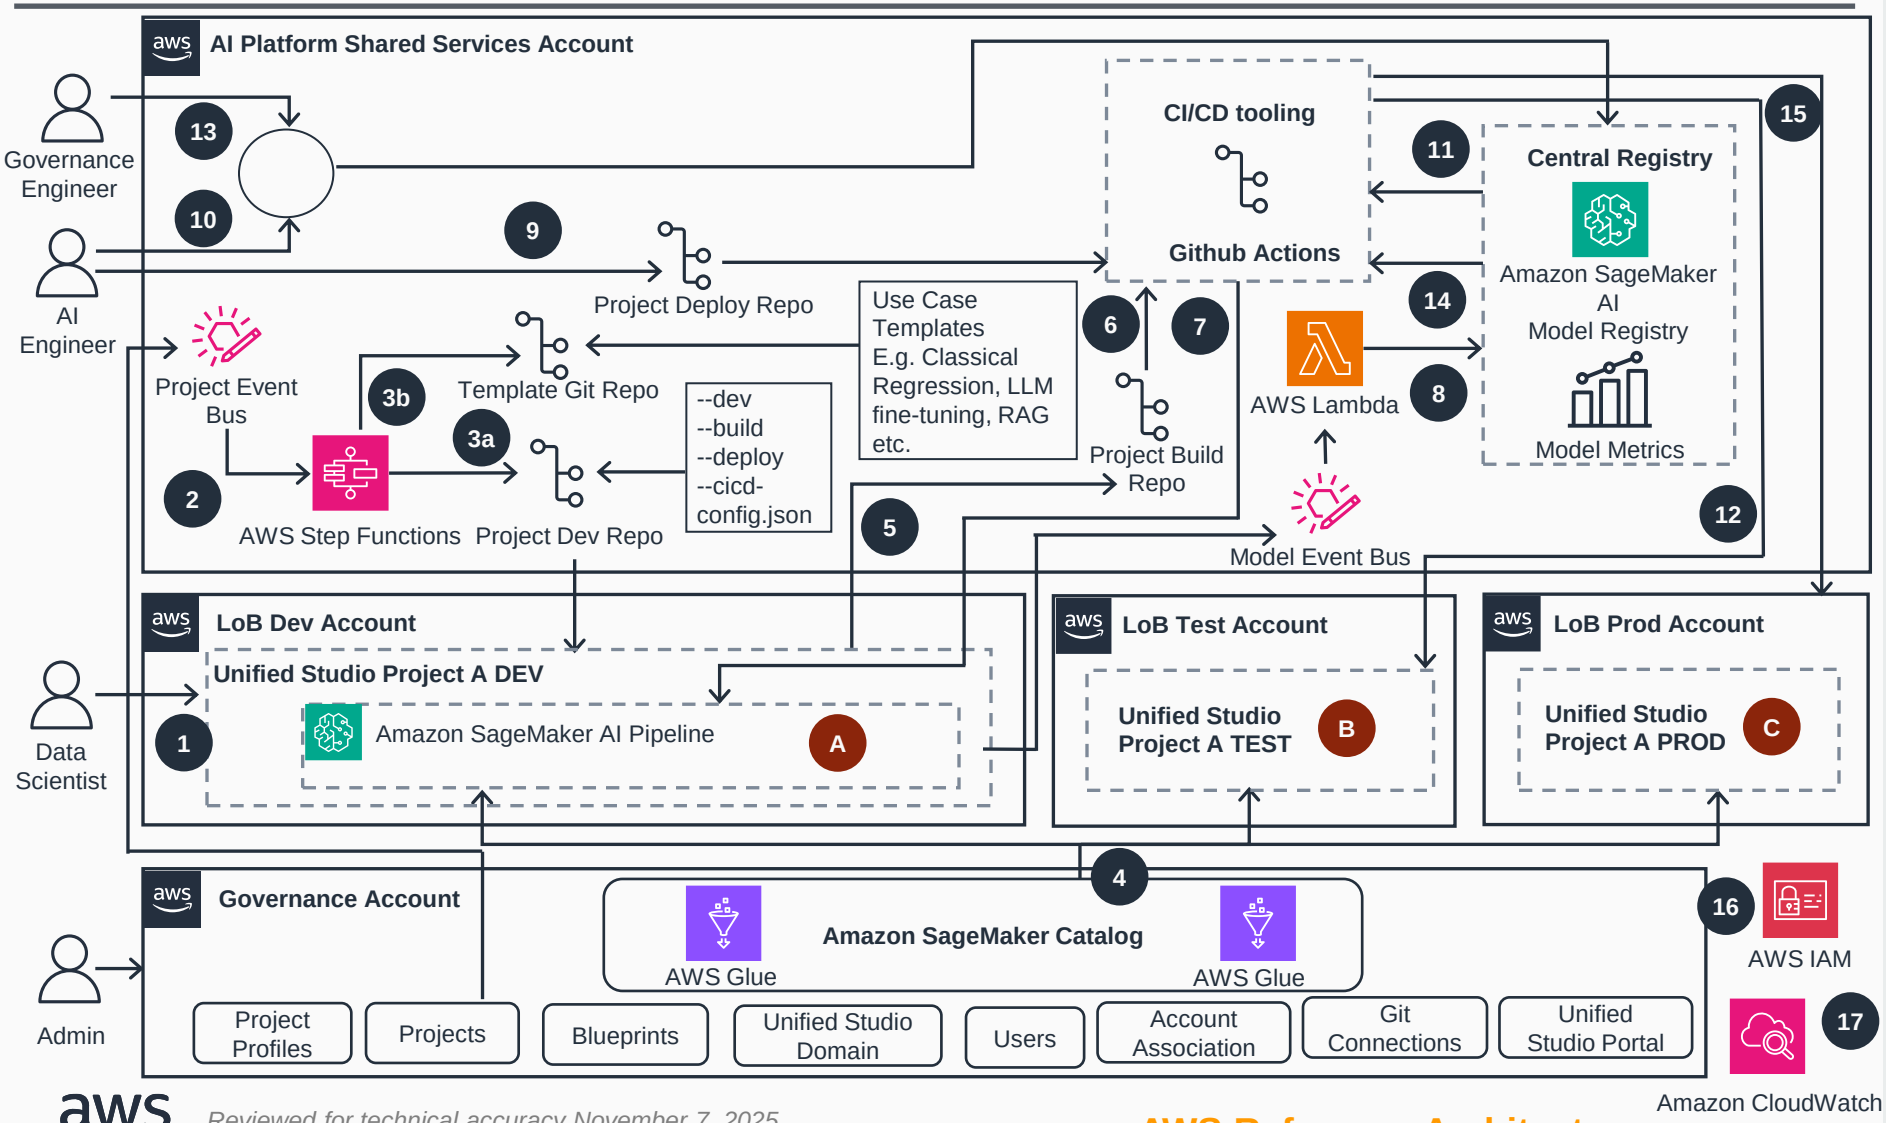

5 The Data Scientist customizes the project's build pipeline code such as adding MLflow experiment tracking and commits changes into the Git repository.

6 A CI/CD pipeline is automatically triggered based on path filter rules when changes are detected in the repository's build folder.

7 The CI/CD pipeline runs the **Amazon SageMaker AI** pipeline in the Project A DEV account to build, train, and evaluate the model. Metrics are tracked in MLflow, and upon successful evaluation, the model is automatically registered in the **Amazon SageMaker** Model Registry within DEV.

8 The Data Scientist approves the model in the Dev Model Registry (stage = Dev, status = Approved). This approval emits an **Amazon EventBridge** event in the Shared Services account, where an **AWS Lambda** function copies the model artifacts and metadata into the Central Model Registry for broader visibility and governance.

9 An AI Engineer validates and defines the deployment parameters (instance type, scaling, and endpoint configuration) by updating the project's deploy repo.

10 The AI Engineer promotes the model in the Central Model Registry to stage = Test with status = Approve, preparing it for test deployment.

11 This Test approval triggers the CI/CD pipeline for test deployment, initiating the automated build and deployment process.

12 The CI/CD pipeline deploys the model endpoint and optionally the **Amazon SageMaker AI** pipeline into the Project A TEST account. Comprehensive integration tests are performed, and the results are recorded back into the Central Model Registry.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture

Guidance for Collaborative, Unified Data and AI Development on AWS

Overview - SageMaker Unified Studio AIOps

This architecture diagram illustrates a multi-account AIOps framework with centralized governance, enabling systematic model progression across Dev, Test, and Prod environments through automated pipelines. This slide shows steps 13-17.

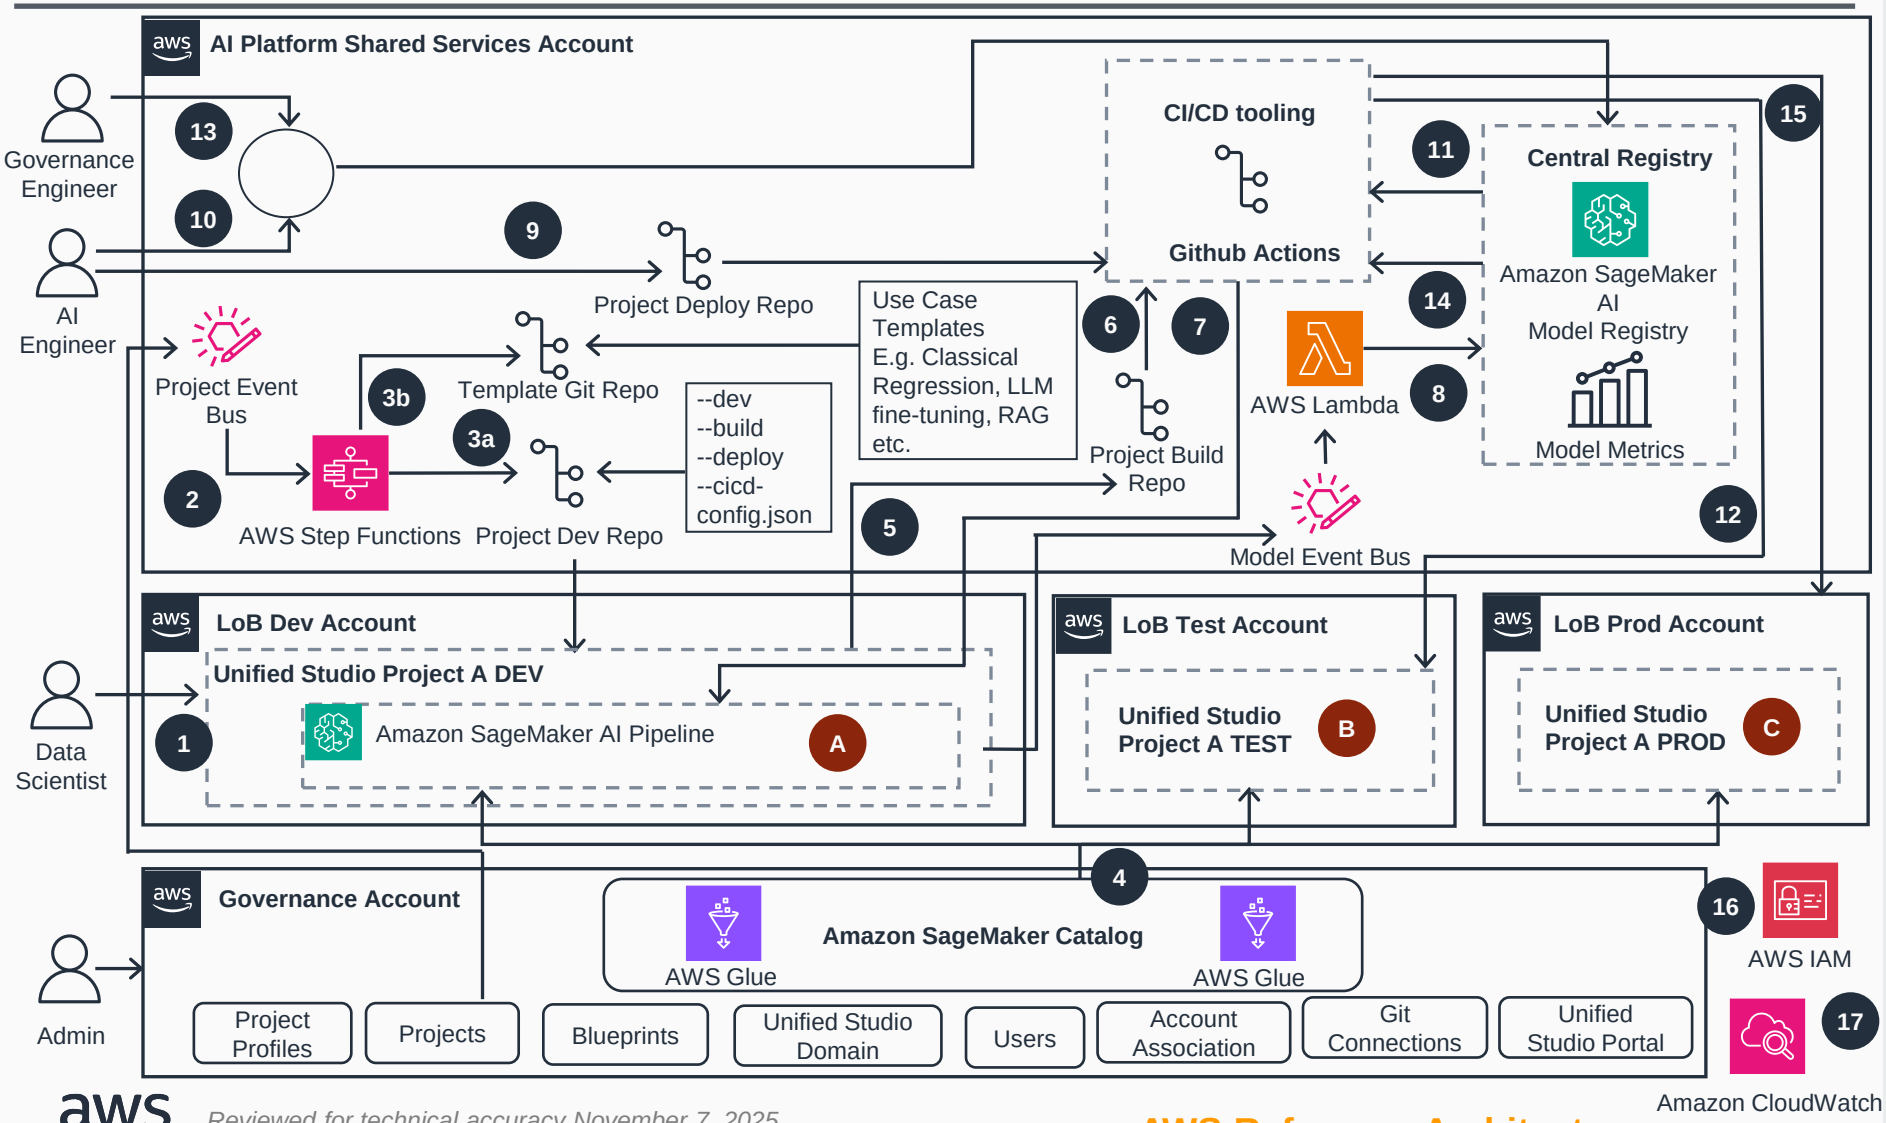

13 The Governance Officer reviews the integration test results. Upon meeting compliance and performance requirements, the model is approved in the Central Registry (stage = Prod, status = Approved) for production deployment.

14 The production approval triggers the Prod CI/CD pipeline, which builds the required artifacts and deploys the model into the Project A PROD account.

15 The model endpoint and optionally **Amazon SageMaker AI** pipeline is deployed in the production environment with A/B testing capability through **Amazon SageMaker AI** endpoints. Model performance is monitored, and results are logged into the Central **Amazon SageMaker AI** Model Registry, enabling continuous monitoring and governance through **Amazon SageMaker AI**'s monitoring capabilities.

Security Control Flow:

16 OIDC JWT provides secure, short-lived credentials for GitHub Actions workflows without storing secrets. Cross-account access is managed through **AWS Identity and Access Management (AWS IAM)** roles with least-privilege permissions. The **Amazon SageMaker** Catalog implements fine-grained access controls to ensure data scientists can only access authorized datasets. **Amazon EventBridge** rules in each account are configured to forward specific events to the central Project Event Bus using resource-based policies, enabling the automated workflow across organizational boundaries.

17 The architecture implements continuous monitoring and governance across the multi-account setup through integrated services. Model performance is monitored through **Amazon SageMaker** endpoints, with results continuously logged into the Central Model Registry. The solution leverages **Amazon CloudWatch** through **AWS CloudFormation** stack for comprehensive logging across Dev, Test, and Prod environments, enabling end-to-end visibility of model performance throughout the automated deployment pipeline.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture

Guidance for Collaborative, Unified Data and AI Development on AWS

ML Model Development and Registration Pipeline (Part A)

This architecture diagram illustrates the Dev pipeline for a multi-account AIOps framework.

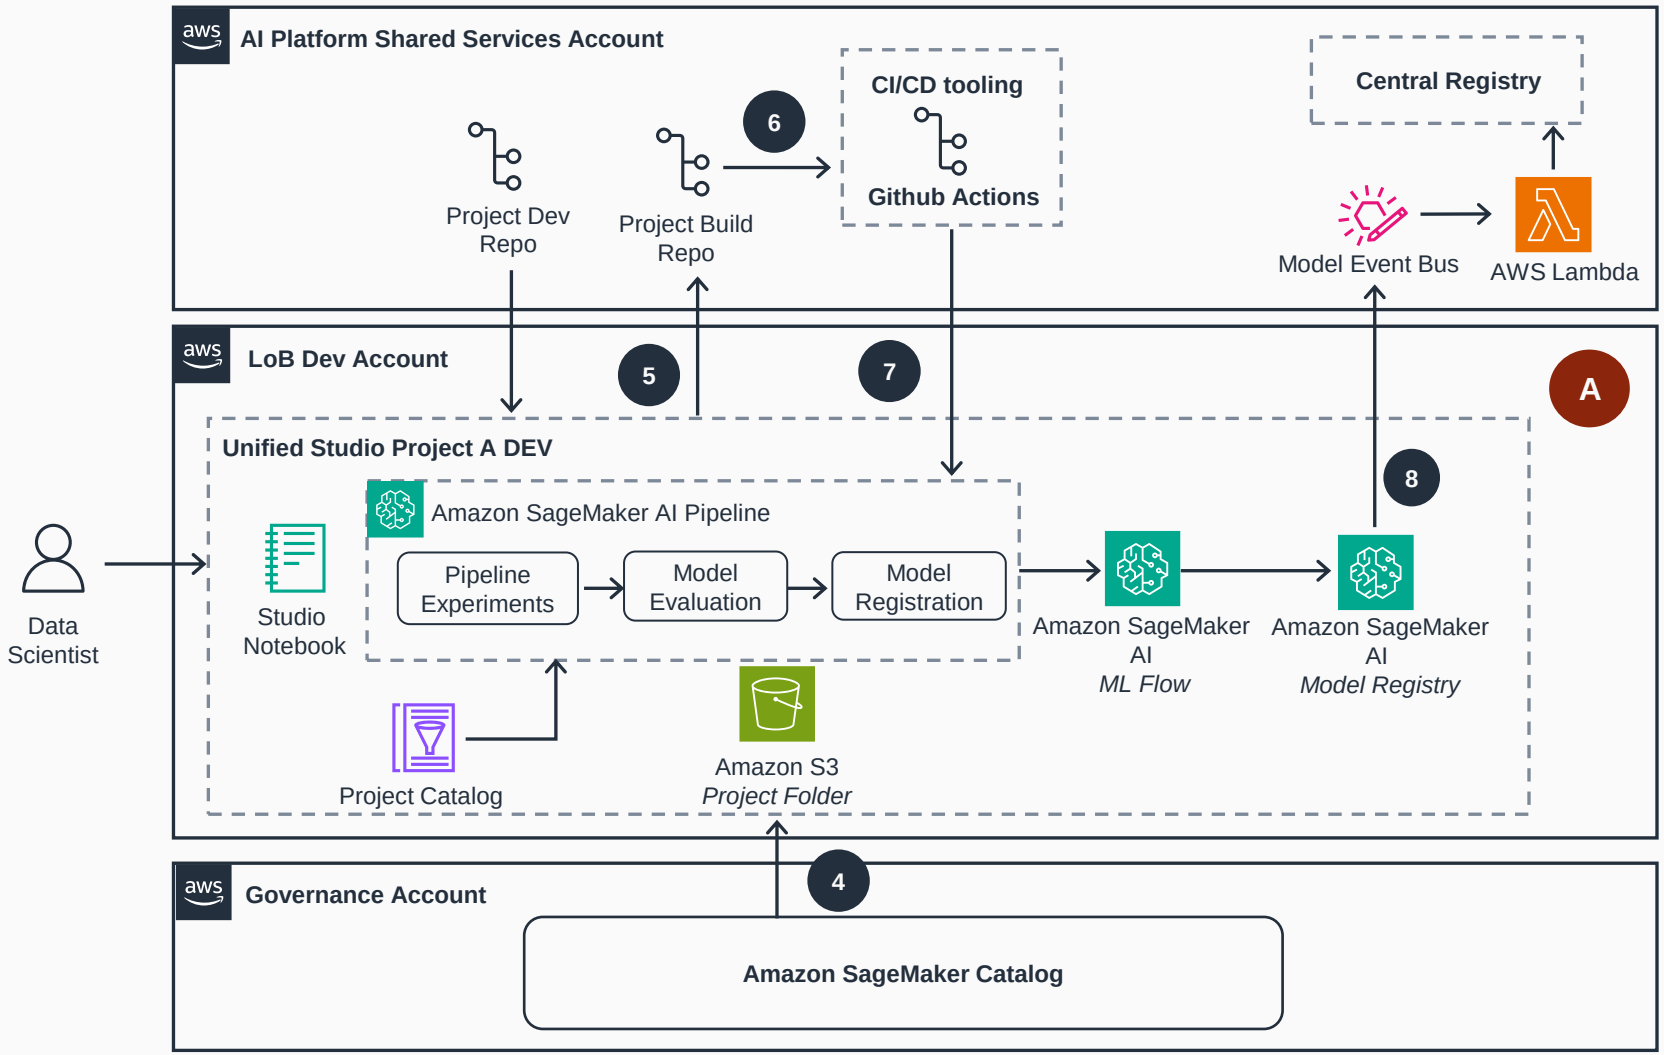

4 Data Scientist subscribes to required datasets (for example, **AWS Glue** tables) through **Amazon SageMaker** Catalog. Once the data producer approves, the assets become available in the project catalog for use in experiments.

5 Data Scientist customizes the project's build pipeline code such as adding MLflow experiment tracking and commits changes into the Git repository.

6 A CI/CD pipeline is automatically triggered based on path filter rules when changes are detected in the repository's build folder.

7 The CI/CD pipeline runs the **Amazon SageMaker AI** pipeline in the Project A DEV account to build, train, and evaluate the model. Metrics are tracked in MLflow, and upon successful evaluation, the model is automatically registered in the **Amazon SageMaker AI** Model Registry within DEV.

8 Data Scientist approves the model in the Dev Model Registry (stage = Dev, status = Approved). This approval emits an **Amazon EventBridge** event in the Shared Services account, where an **AWS Lambda** function copies the model artifacts and metadata into the Central Model Registry for broader visibility and governance.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture

Guidance for Collaborative, Unified Data and AI Development on AWS

ML Model Deployment Pipeline (Part B and C)

This architecture diagram illustrates the Test and Prod pipelines for a multi-account AIOps framework.

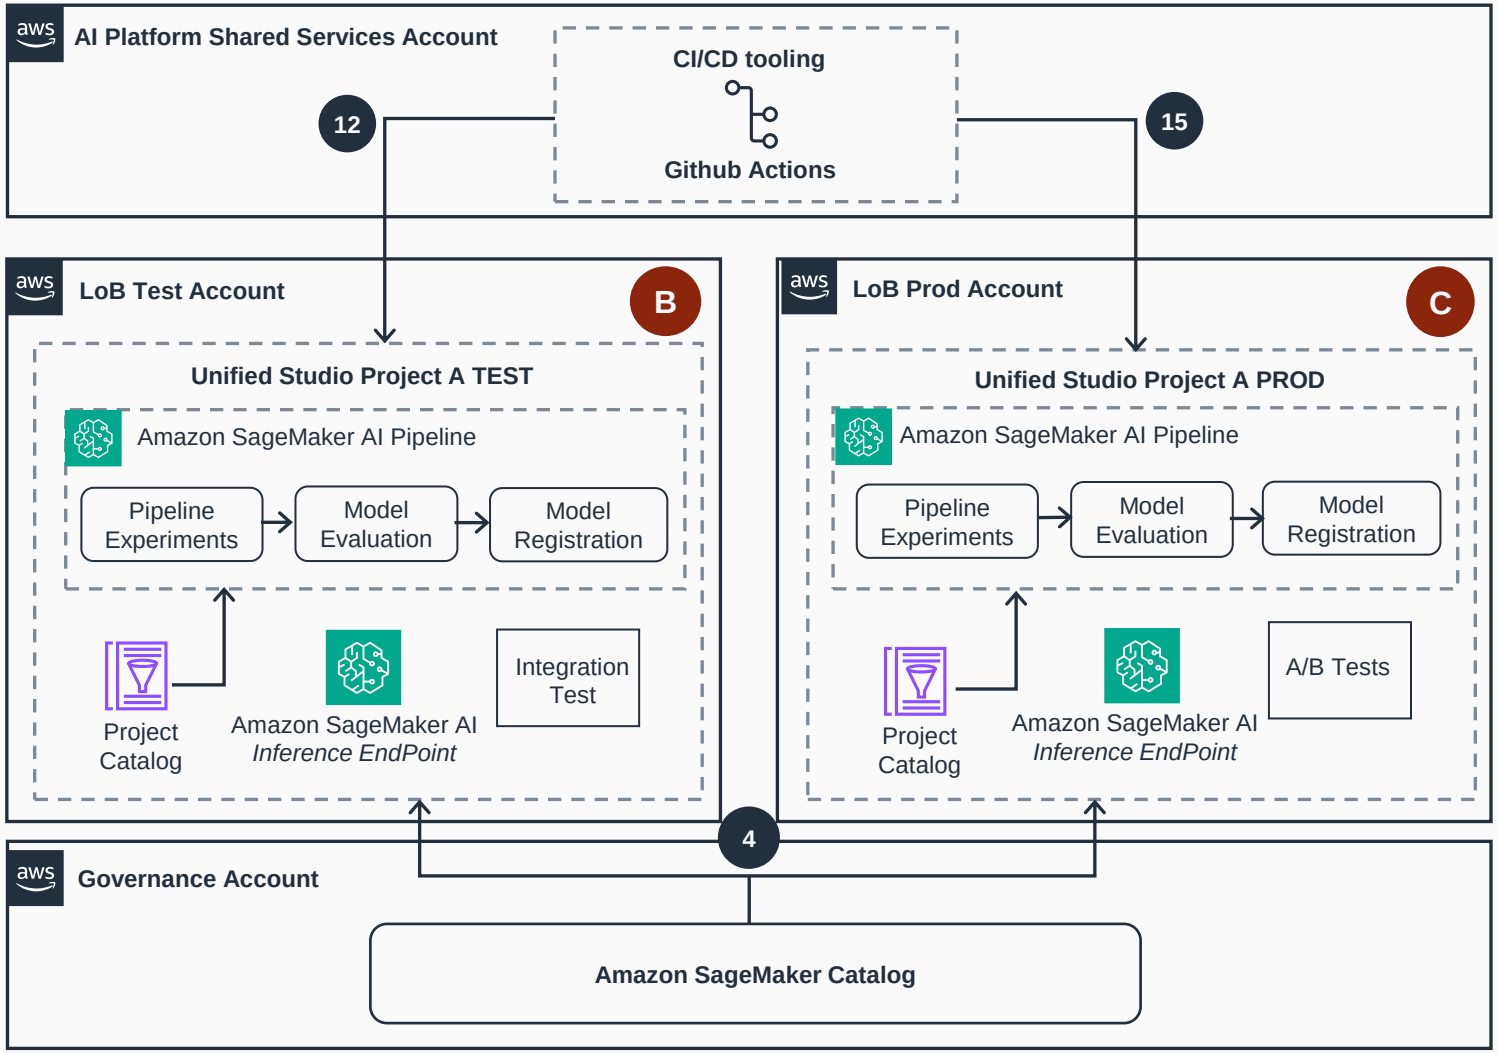

4 The **Amazon SageMaker** Catalog in the Governance Account connects to both the LoB Test Account and LoB Prod Account. This integration enables access to required datasets through the Project Catalog in both test and production Unified Studio Project A environments.

12 The CI/CD pipeline deploys the model endpoint and optionally the **Amazon SageMaker AI** pipeline into the Project A TEST account. Comprehensive integration tests are performed, and the results are recorded back into the Central Model Registry. The **Amazon SageMaker AI** Pipeline in the TEST environment executes pipeline experiments, model evaluation, and model registration, utilizing the **Amazon SageMaker AI** Inference EndPoint.

15 The model endpoint and optionally SageMaker pipeline is deployed in the production environment with A/B testing capability. Model performance is monitored and results are logged into the Central Model Registry, enabling continuous monitoring and governance. The **Amazon SageMaker AI** Pipeline in the PROD environment manages pipeline experiments, model evaluation, and model registration, interacting with the **Amazon SageMaker AI** Inference EndPoint.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture

Guidance for Collaborative, Unified Data and AI Development on AWS

Automated Project Provisioning (Technical Implementation)

This architecture demonstrates an automated AIOps workflow within Amazon SageMaker Unified Studio, orchestrating the AI lifecycle from project initiation to model deployment through integrated CI/CD pipelines. This slide shows steps 1-6.

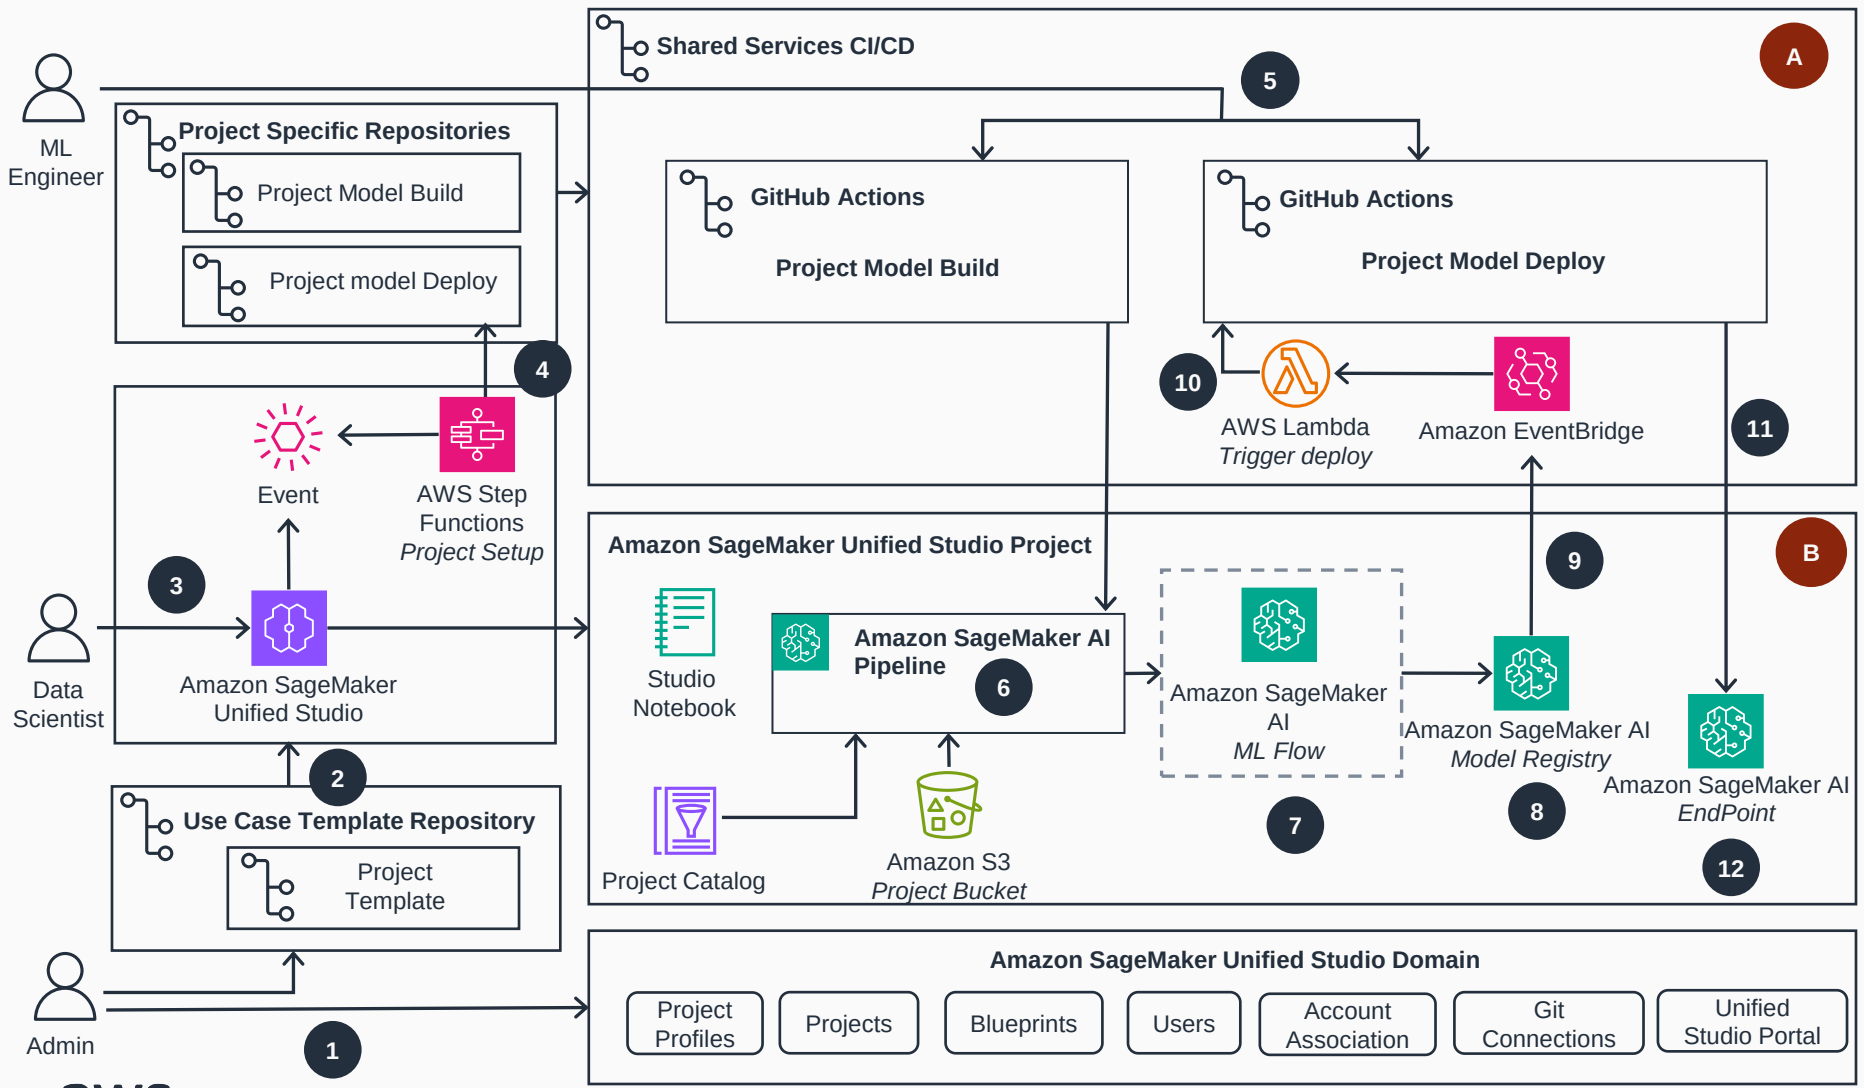

1. The Administrator configures the **Amazon SageMaker** Unified Studio environment by setting up domains, AWS infrastructure, authentication, GitHub connections, and project template repositories. This provides the foundation for consistent, governed ML project creation.
2. The Use Case Template Repository is prepared with standard templates that define configurations for model build and deployment, ensuring every new project begins with an approved baseline.
3. When a Data Scientist creates a new project in **Amazon SageMaker** Unified Studio, a Create Project event is emitted. This event is captured by **Amazon EventBridge** and triggers an **AWS Lambda** function to automate setup.
4. The **AWS Step Functions** Project Setup provisions dedicated repositories for model build and model deploy. These repositories are configured and prepopulated with seed code and CI/CD workflows, including GitHub Action secrets.
5. The project repositories are linked to the Shared Services CI/CD system, enabling automated build and deployment pipelines to be centrally managed and consistently applied.
6. During development, the **Amazon SageMaker AI** Pipeline runs within the project, orchestrating steps for data preprocessing, feature engineering, training, evaluation, and registration. Outputs are stored in the project's **Amazon S3** bucket.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture

Guidance for Collaborative, Unified Data and AI Development on AWS

Model Approval and Deployment Pipeline(Technical Implementation)

This architecture demonstrates an automated AIOps workflow within Amazon SageMaker Unified Studio, orchestrating the AI lifecycle from project initiation to model deployment through integrated CI/CD pipelines. This slide shows steps 7-12.

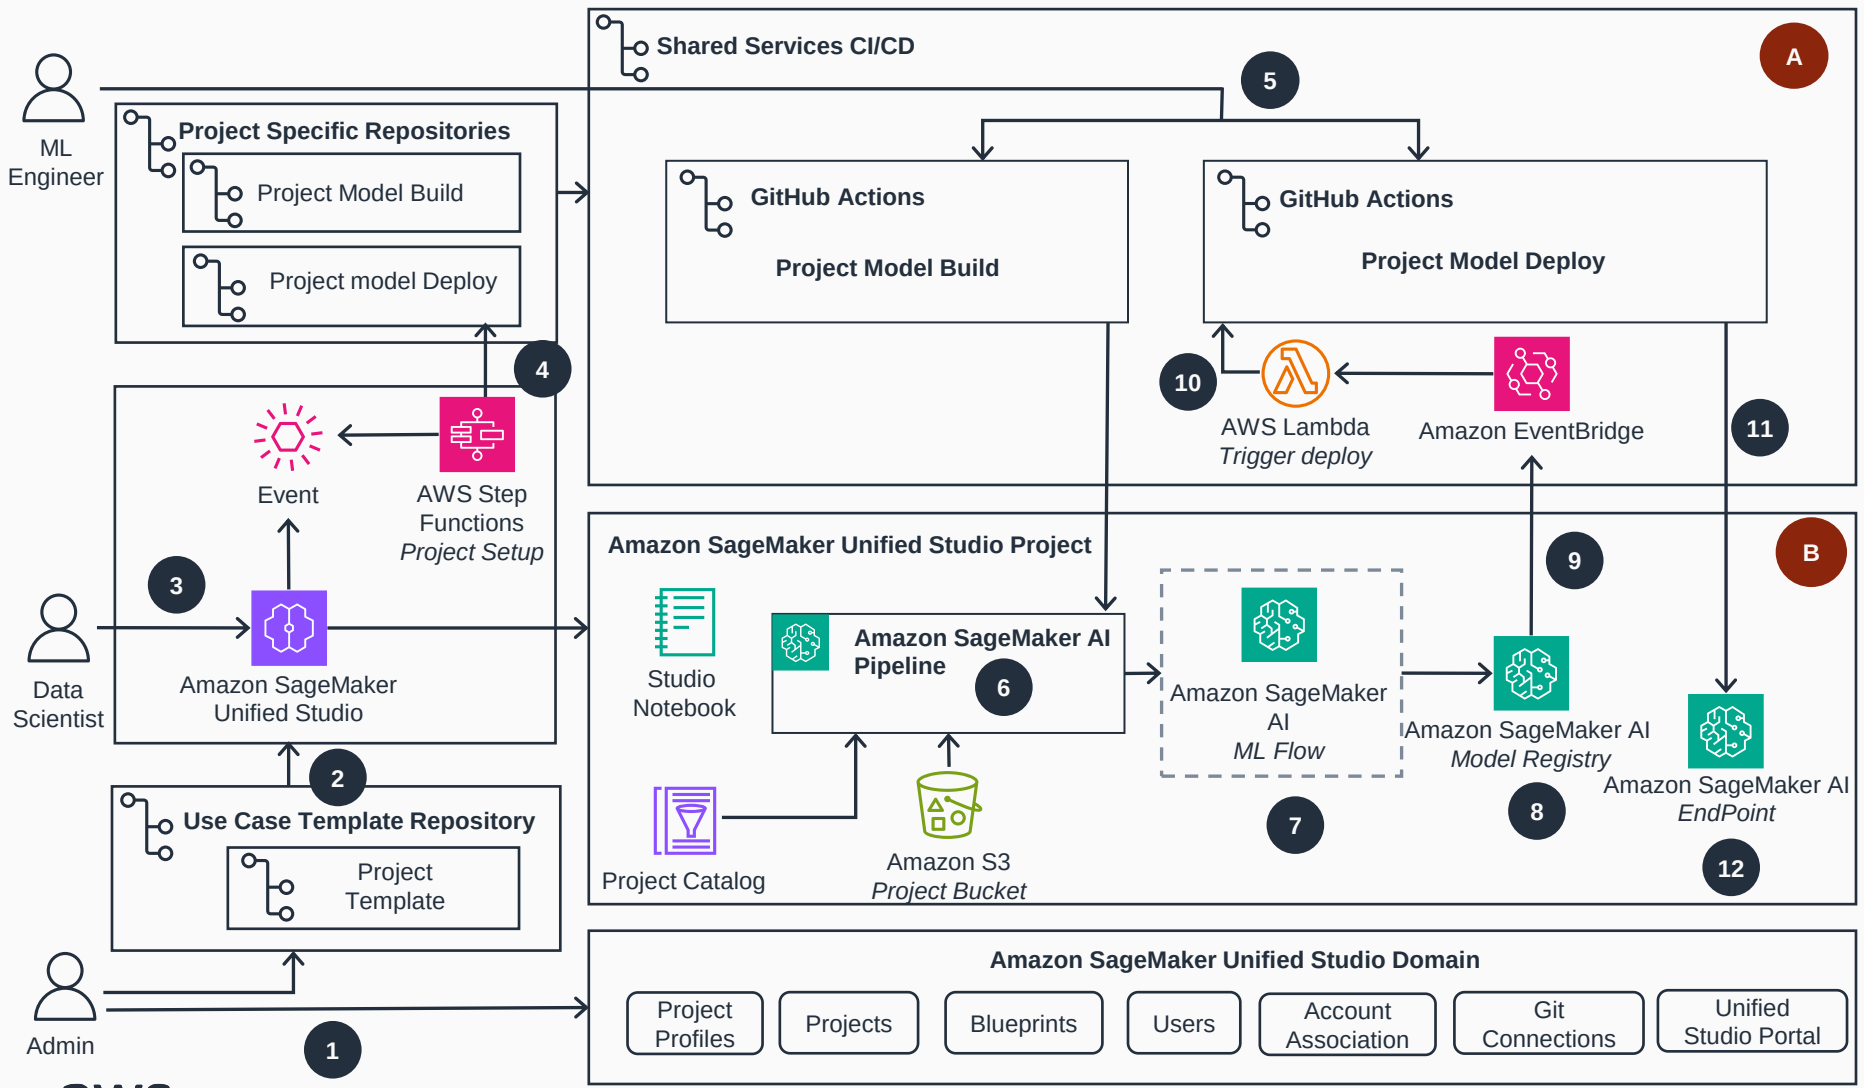

7 All experiments are tracked using **Amazon SageMaker AI** MLflow integration, which logs metrics, artifacts, and experiment details for full traceability.

8 When training and evaluation succeed, the model is registered in the **Amazon SageMaker AI** Model Registry, awaiting further validation and approval.

9 The registered model undergoes review. Once approved, it is marked ready for deployment, ensuring only validated models progress to the next phase.

10 An **Amazon EventBridge event** is emitted upon approval, which invokes a Deploy Lambda function to initiate the deployment process.

11 The Project Model Deploy repository runs its GitHub Actions workflow to fetch the approved model, validate configurations, and provision/update the **Amazon SageMaker AI** Endpoint.

12 The model is deployed as an **Amazon SageMaker AI** Endpoint, completing the automated journey from development to production. The endpoint is live and ready to serve inference requests.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture

Guidance for Collaborative, Unified Data and AI Development on AWS

ML Pipeline Execution and Model Deployment (Technical Implementation)

This architecture demonstrates an automated AIOps workflow within Amazon SageMaker Unified Studio, orchestrating the AI lifecycle from project initiation to model deployment through integrated CI/CD pipelines.

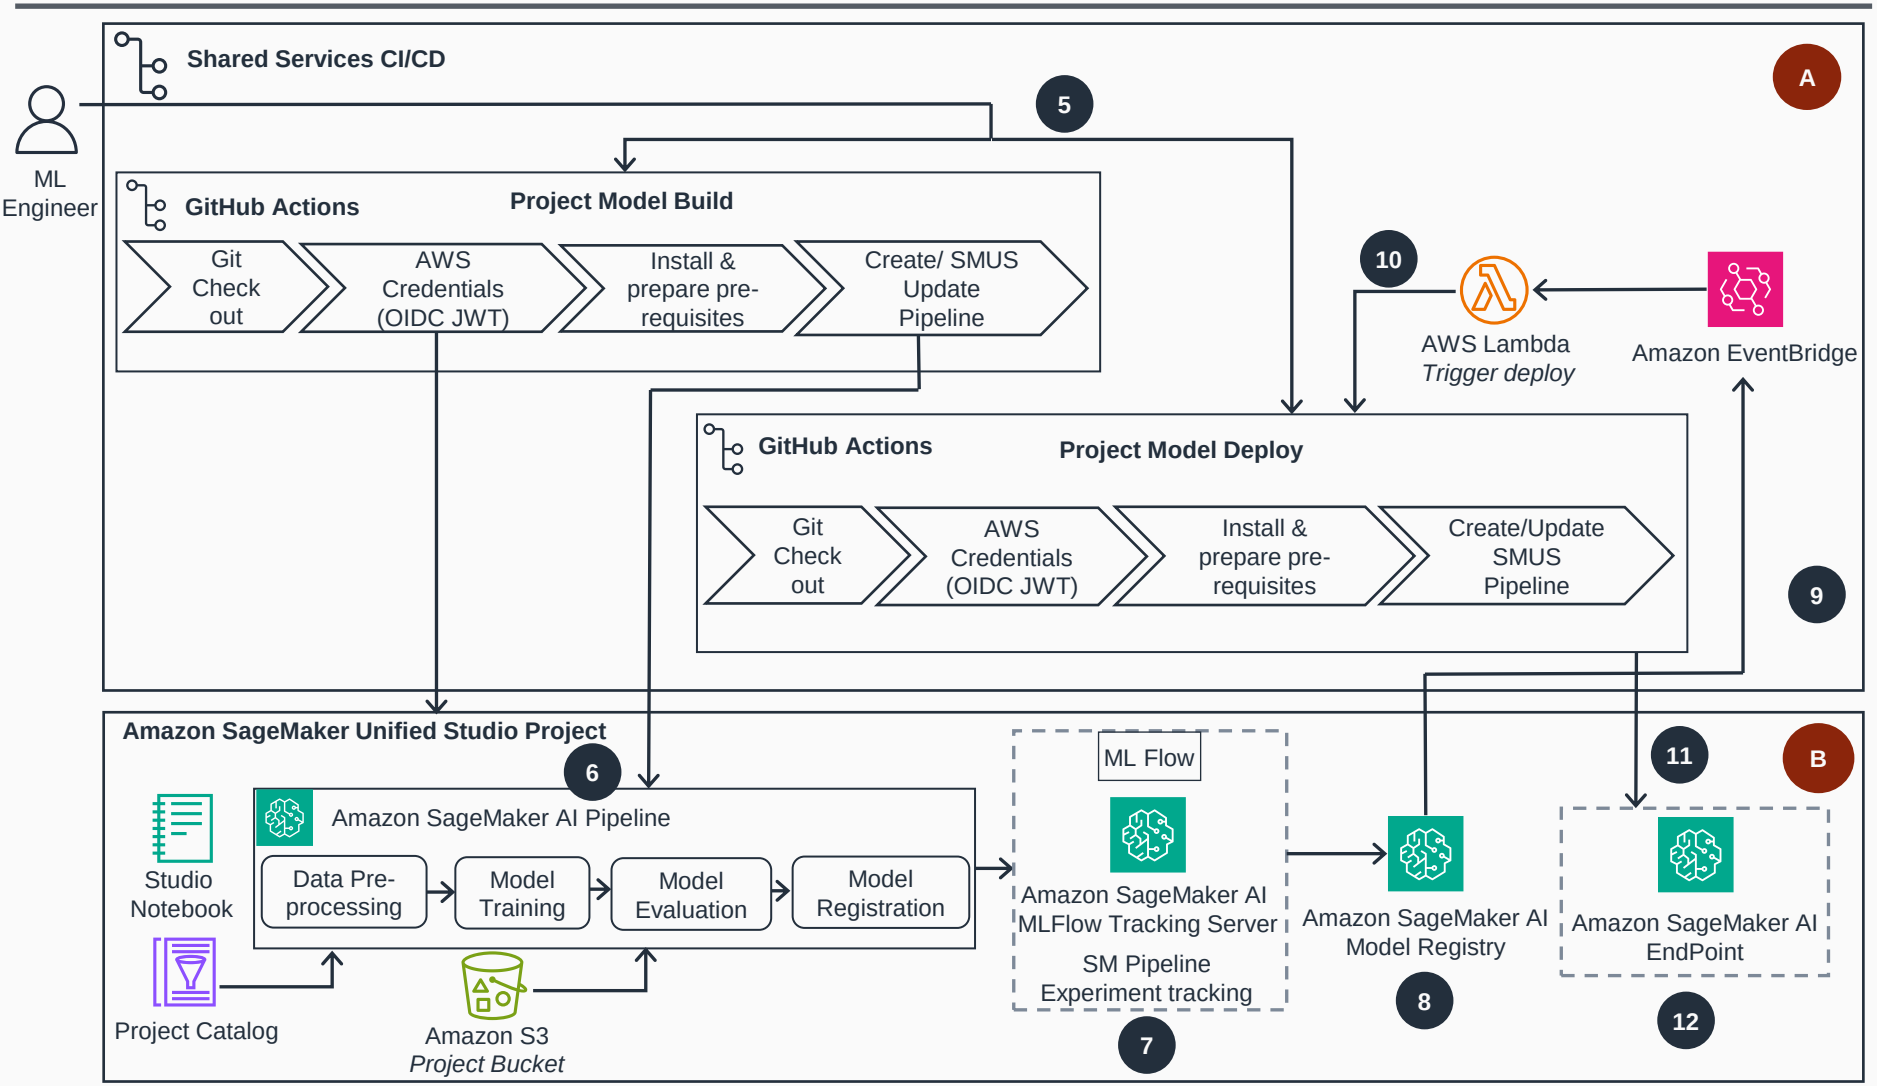

5. The project repositories are linked to the Shared Services CI/CD system, which manages the build and deployment pipelines through GitHub Actions.

6. During development, the **Amazon SageMaker AI** Pipeline runs within the project, executing a series of connected steps: Data Pre-processing, Model Training, Model Evaluation, and Model Registration. These steps are initiated from the Studio Notebook and integrate with the project's **Amazon S3** bucket.

7. All experiments are tracked using **Amazon SageMaker AI** MLflow integration through the MLflow Tracking Server, which logs metrics, artifacts, and experiment details for full traceability. This connects directly with the **Amazon SageMaker AI** Pipeline workflow.

8. When training and evaluation succeed, the model is registered in the **Amazon SageMaker AI** Model Registry through the Model Registration step in the pipeline, awaiting further validation and approval.

9. The registered model undergoes review in the **Amazon SageMaker AI** Model Registry. Once approved, it is marked ready for deployment, ensuring only validated models progress to the next phase.

10. An **Amazon EventBridge** event is emitted upon approval, which invokes a Deploy **AWS Lambda** function to trigger the deployment process through the GitHub Actions workflow.

11. The Project Model Deploy repository runs its GitHub Actions workflow, following specific steps: Git checkout, AWS credentials setup (OIDC JWT), prerequisite preparation, and SMUS Pipeline updates to handle the approved model.

12. The model is deployed as an **Amazon SageMaker AI** Endpoint through the final step of the deployment pipeline, completing the automated journey from development to production. The endpoint becomes live and ready to serve inference requests.

Reviewed for technical accuracy November 7, 2025

AWS Reference Architecture