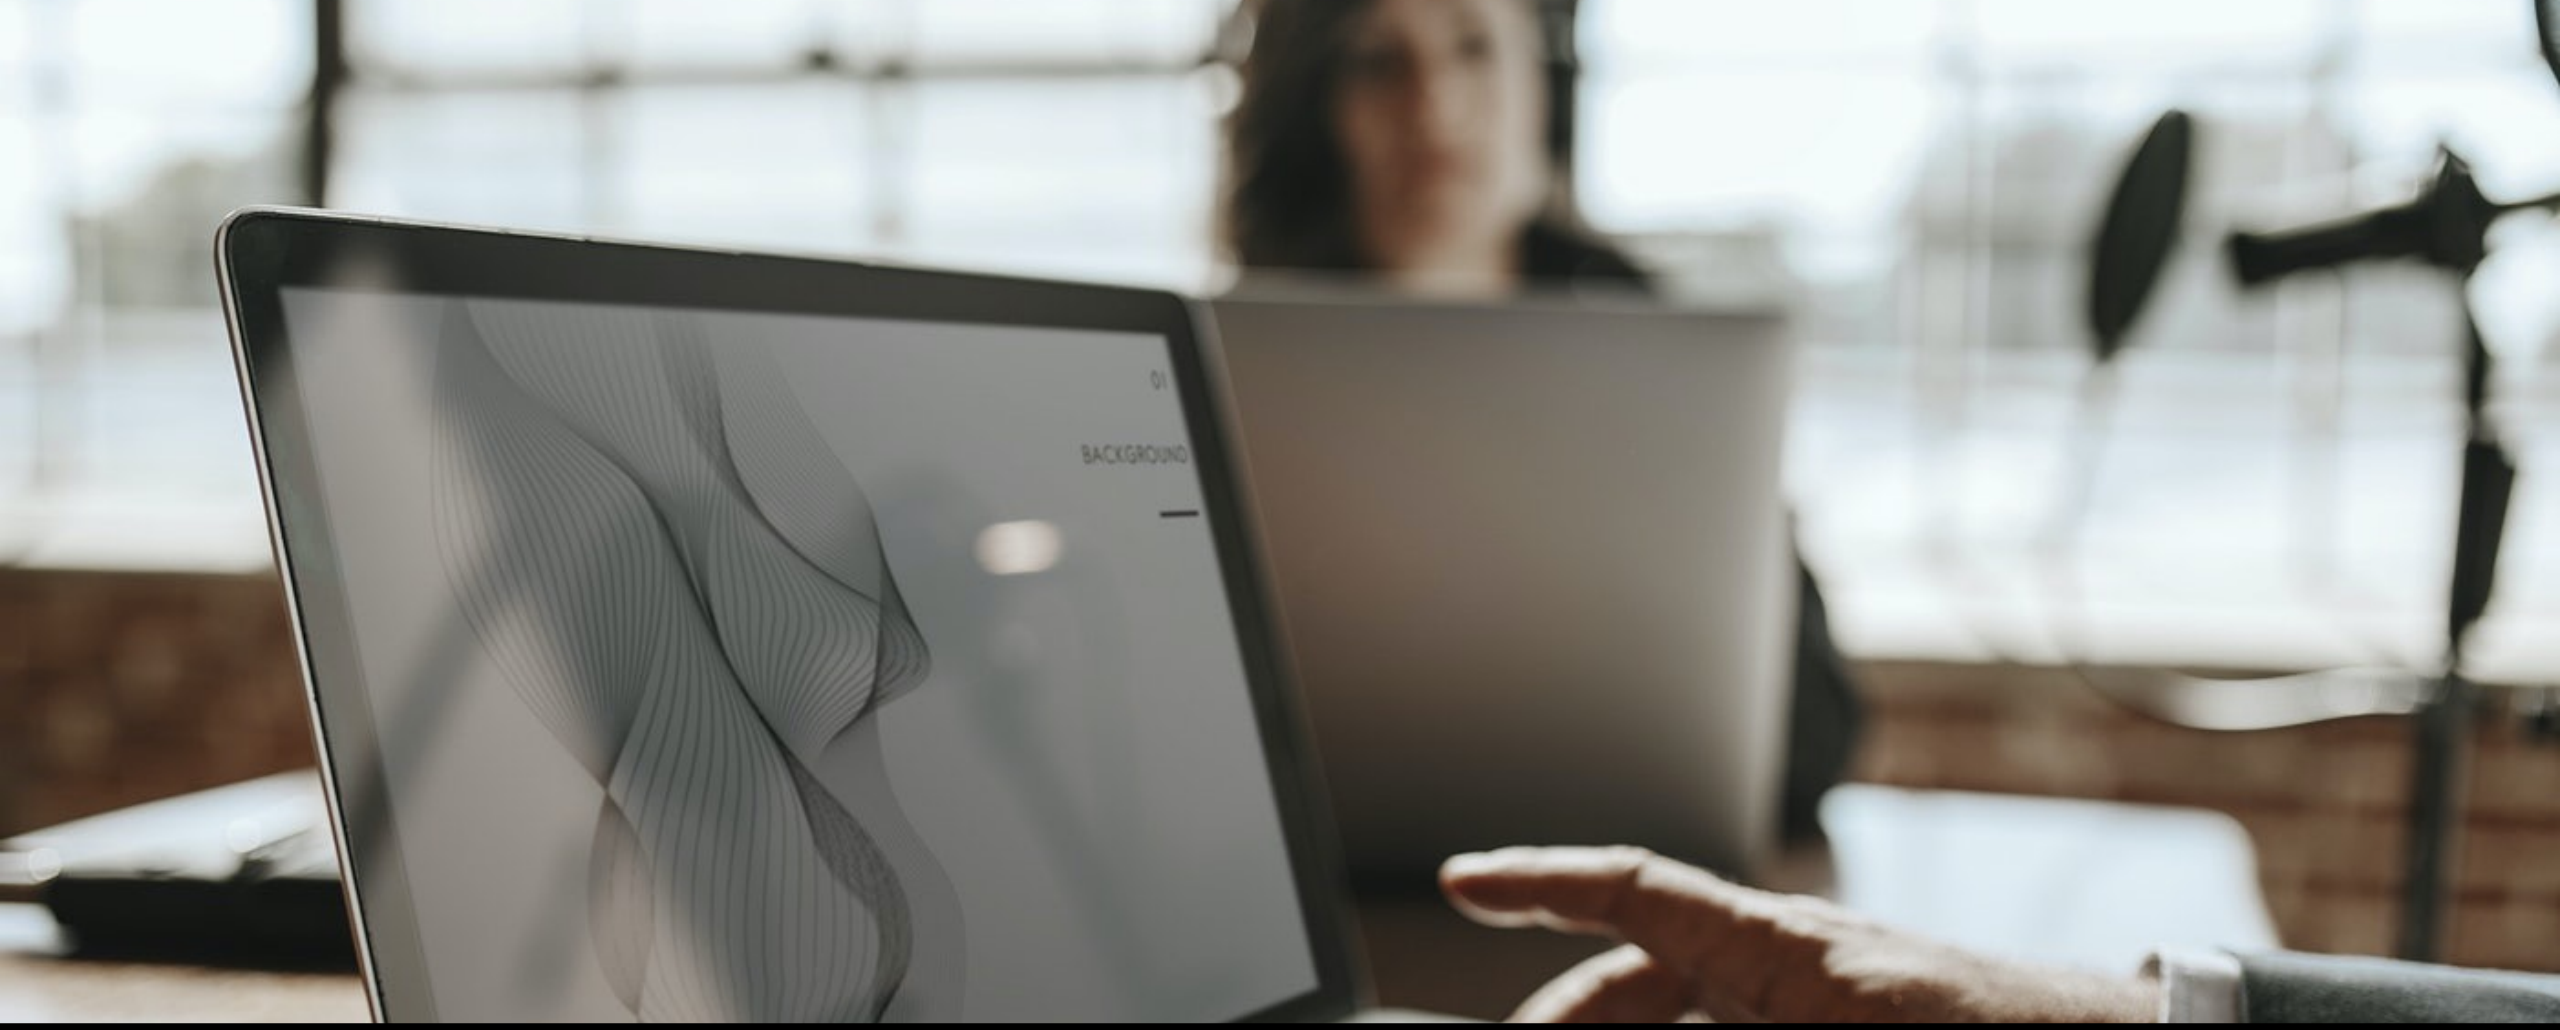

AI&ML CONF

aspitalia.com

UGIdotNET

Big thanks to our sponsors

Secure MLOps

Building secure MLOps platforms for regulated industries

Yevgeniy Ilyin
AWS, Senior Solutions Architect,

AI&ML CONF

Agenda

- MLOps – what problem we are solving
- MLOps – considerations for regulated industries
- MLOps design process, decisions, and key components
- Apache Airflow and Amazon MWAA
- MLflow
- Hands on: MLOps with Amazon SageMaker

MLOps – path to production

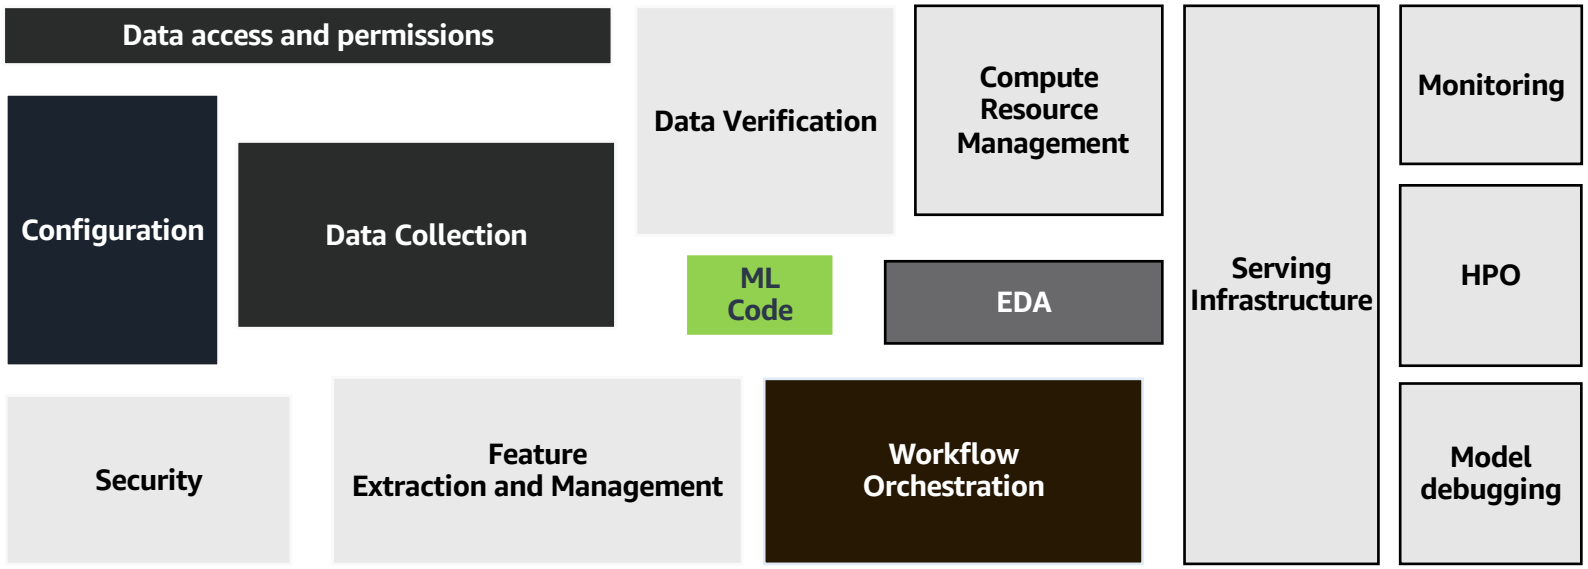

MLOps – what problems we are solving?

- Getting models to production
- Operationalization of AI/ML workloads and workflows
- Create secured, automated, and reproducible ML workflows
- Manage models with a model registry and data lineage
- Enable continuous delivery with IaC and CI/CD pipelines
- Monitor performance and feedback information to your models
- Providing compliance, security, and cost tools for ML development
- Increasing collaboration and experimentation
- Support diverse, multidisciplinary teams

ML development – a need for a structure

- Today:

- Only 53% of all pilot projects make it into production
- 9 months average time to deploy
- Focus mostly on developing ML models
- Operationalization is secondary

- Tomorrow:

- by the end of 2024, 75% of organisations will shift from piloting to operationalising AI, driving a 400% increase in streaming data and analytics infrastructures (Gartner)

https://www.idgconnect.com/article/3583467/gartner-accelerating-ai-deployments-paths-of-least-resistance.html

The main barriers to AI/ML implementation

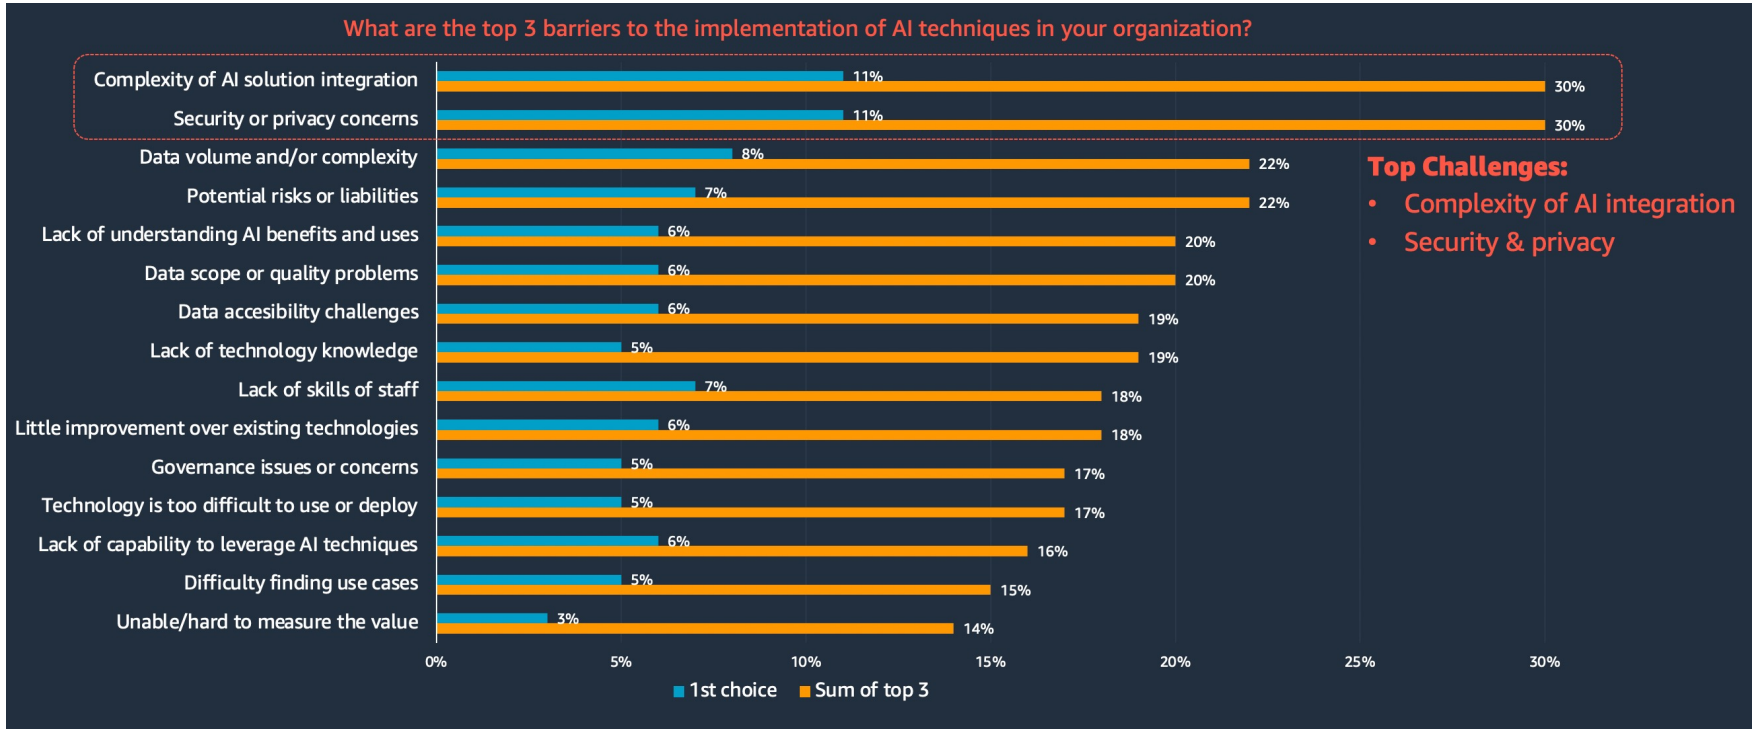

2019 Gartner AI in Organizations Survey

MLOps for regulated industries

Two main challenges in the focus:

1. Security: Create secure, compliant, and resilient ML workloads
2. Integration: Effectively support integration and operationalization complexity in the existing development organization, governance, and operational model

Design and process decisions to be made

Before implementing your MLOps platform, you need to answer the following design questions:

- Centralized or federative access control and permission approach
- Who are your key user personas
- Data access control end-to-end
- Single or multi-account cloud setup
- Framework/technology for:
  - AI/ML development
  - CI/CD automation
  - ML workflow and automation
  - Data storage, processing, and access

MLOps technology components

Consider and evaluate the following technology topics for implementing MLOps:

- ML development/experimentation/collaboration
- Compute/training environment and infrastructure
- Model registry
- Feature store
- Model deployment
- Monitoring in production
- Hyperparameter optimization
- Dataset management
- Workflow orchestration and pipelines

MLOps platform key features and categories

Data Acquisition and Preparation					
Data Acquisition	Data Visualization and Exploration	Data Preparation and Transformation	Data Versioning	Data Pipelines	Data Labeling
Model Development and Training					
AutoML	Feature Extraction and Engineering	Feature Store	ML Pipelines or Workflows	Development Environment Support	Model Repository
Model Marketplace	Model Training	Distributed Model Training	Model Debugging	Experiment Management	Deep Learning Support
Reinforcement Learning Support	Bias Detection and Mitigation	Model Explainability	Tuning and Optimization		
Model Deployment and Operations					
Model Portability and Compression	Model Deployment	Edge ML Support	Model Monitoring	Cost Management	
System-wide features					
ML Infrastructure Orchestration	Accelerator Support	Kubernetes Support	Teamwork and Collaboration	Enterprise Security	Governance

https://twimlai.com/solutions/introducing-twiml-ml-ai-solutions-guide/

MLOps platform key requirements

A good MLOps platform must provide the following key features:

- Reusability: once created, any MLOps component must be reusable
- Reproducibility: same data, same initial conditions must produce same result
- Security: any MLOps component must be integrated in end-to-end security concept
- Auditability: any material change on data or model must be logged and versioned
- Governance: MLOps components are approved, compliant, and monitored
- Scalability: ability to scale and grow with use cases, data volume, and organizational changes
- Flexibility: MLOps platform must accommodate any ML framework

Popular frameworks to build MLOps pipelines

MLflow

Open source platform for the ML lifecycle

Apache Airflow

Platform to author, schedule and monitor workflows

Kubeflow

ML toolkit for Kubernetes

AWS Step Functions

Serverless pipeline orchestration

Amazon SageMaker Pipelines

Managed ML pipelines in SageMaker Studio

Apache Airflow and Amazon MWAA

Main reasons to chose Apache Airflow:

- Extensibility: Airflow operators are re-usable and can be developed based on your requirements. Multi-cloud support.
- Directed Acyclic Graph (DAG) workflow management: simple mechanism for defining and running complex workflows with dependencies.
- Python-based: imperative (how) programming paradigm. DAG file is a simple Python file
- Open-source and active community

Apache Airflow DAG

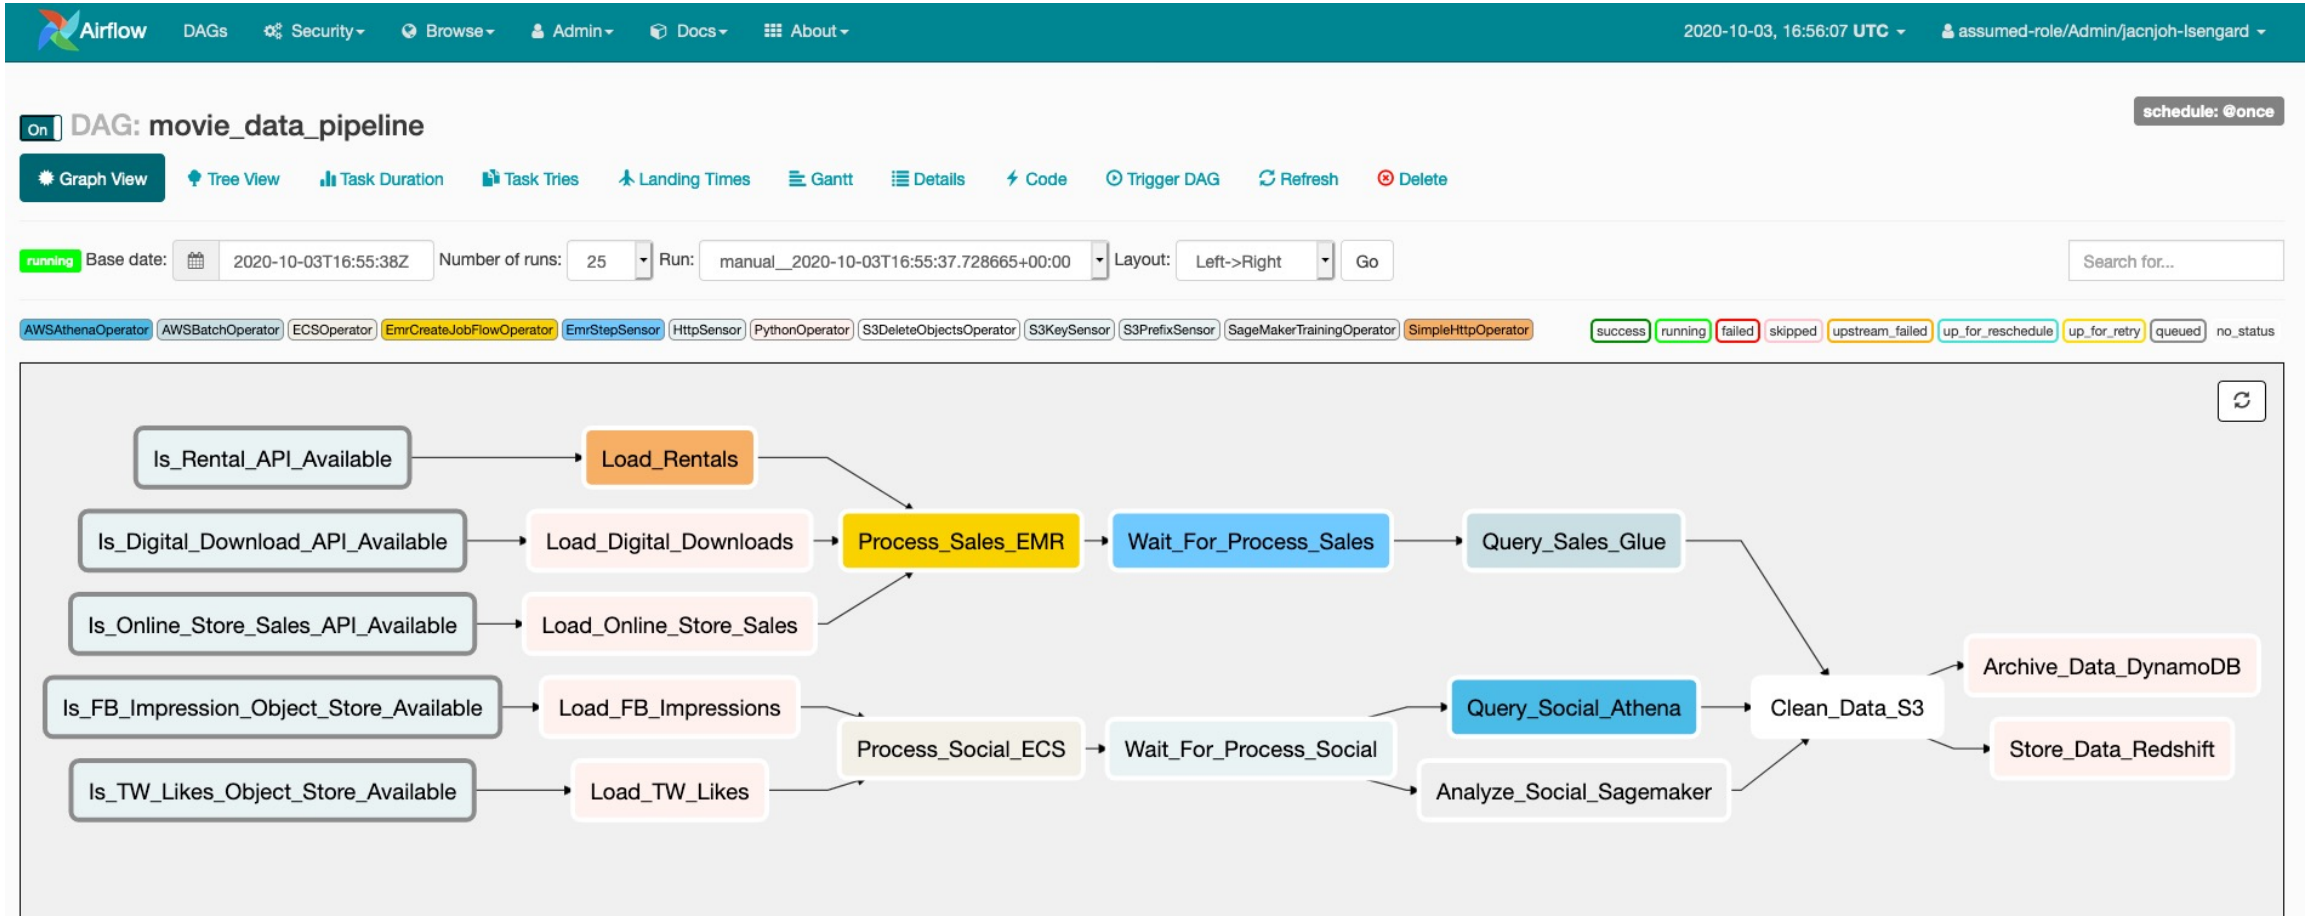

MLflow

Open-source platform to manage the ML lifecycle, including experimentation, reproducibility, deployment, and a central model registry

- Tracking: record and query experiments: code, data, configuration, and results. Provides API and UI for logging parameters, code versions, metrics, and artifacts when running your machine learning code and for later visualizing the results
- Projects: package data science code in a format to reproduce runs on any platform
- Models: deploy ML models in diverse serving environments. Offer a convention for packaging machine learning models in multiple flavours
- Registry: store, annotate, discover, and manage models in a central repository. It provides model lineage (which MLflow experiment and run produced the model), model versioning, stage transitions (for example from staging to production or archiving), and annotations.
- Supports multi-cloud, library-agnostic, all functions are accessible through a REST API and CLI
- Designed to scale to large data sets, large output files, and large number of experiments

MLflow Tracking

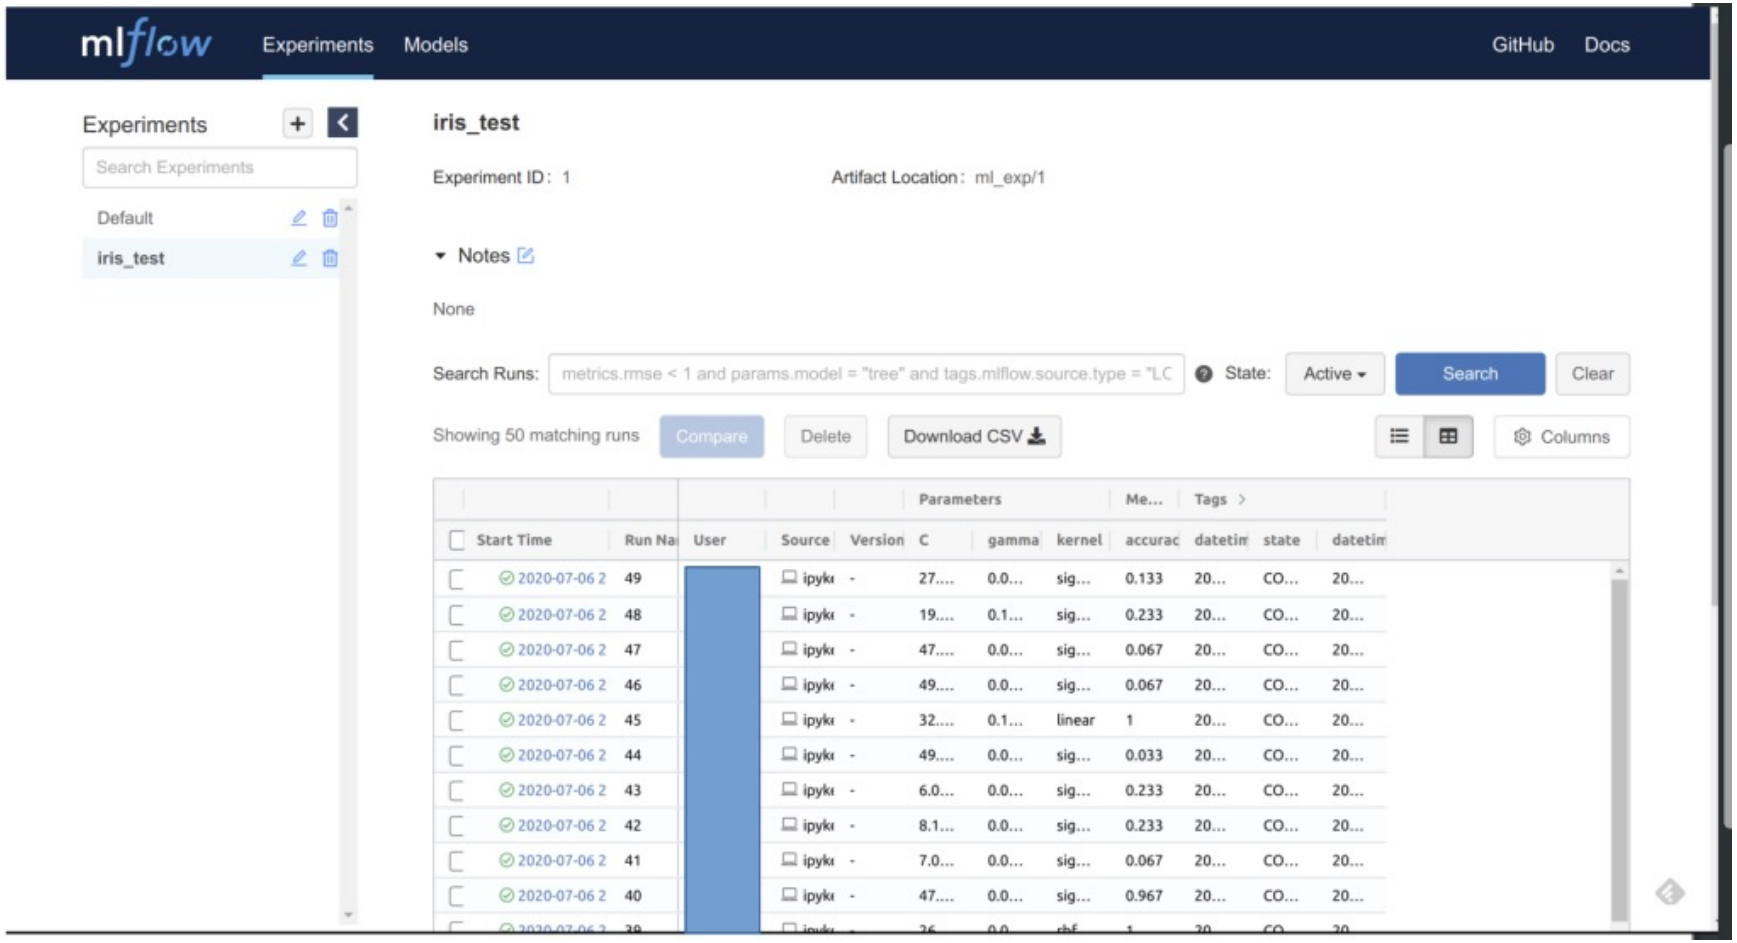

MLflow Model Registry

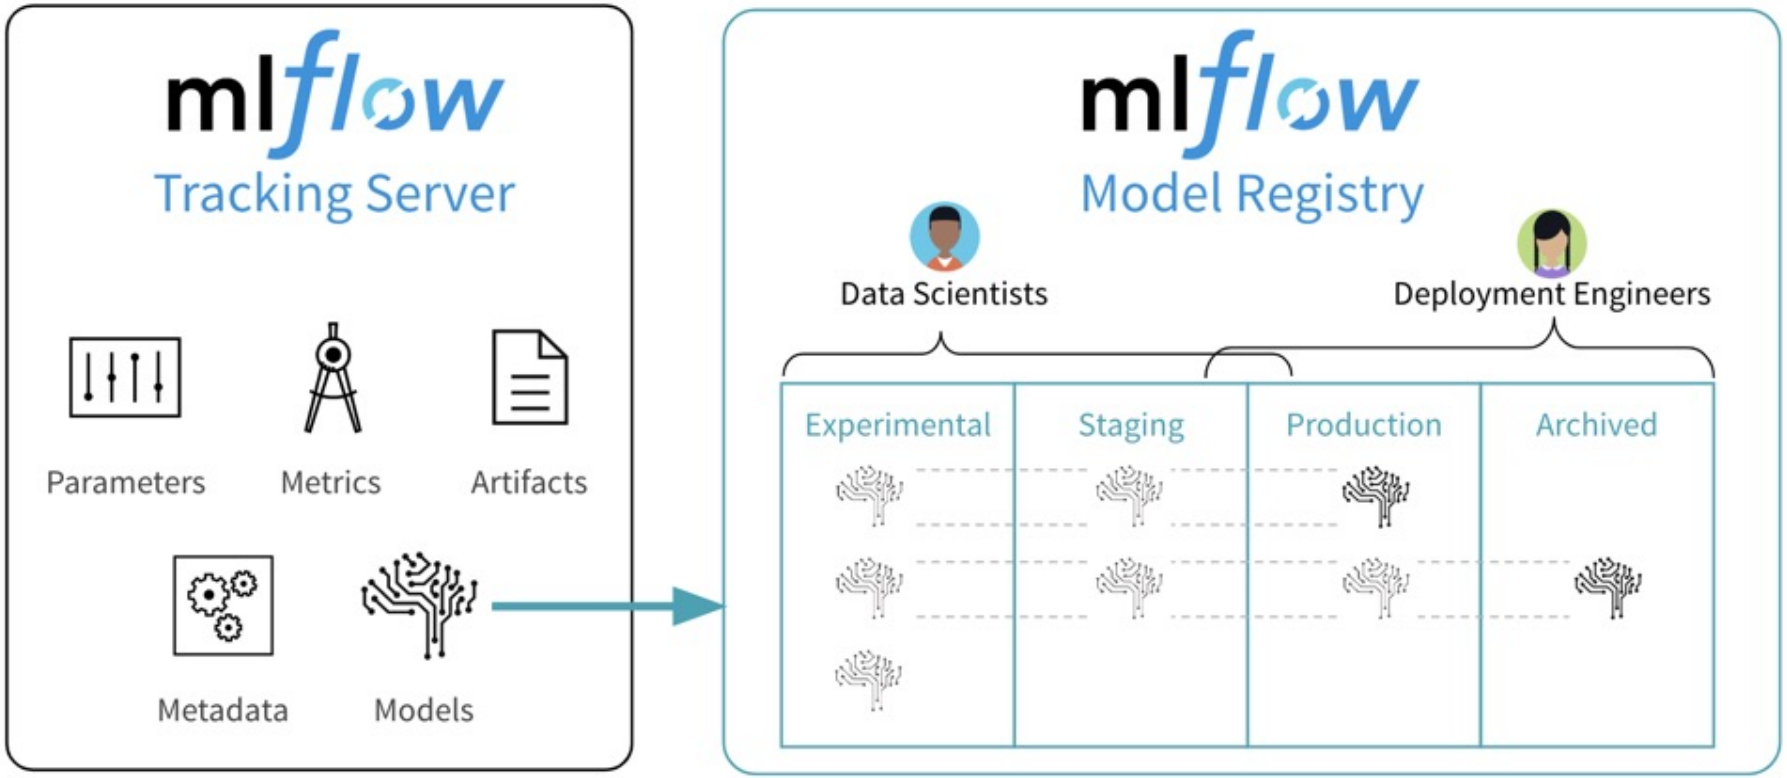

https://databricks.com/blog/2019/10/17/introducing-the-mlflow-model-registry.html

Amazon SageMaker (with hands-on demo)

Amazon SageMaker and SageMaker Studio provides:

- Integrated collaborative environment for all key ML personas
- Managed compute resources for training, processing, and inference
- Integrated end-to-end security and access control
- ML workflows with SageMaker Pipelines
- CI/CD automation with GitHub, Gitlab, Jenkins, AWS CodePipeline
- Reusable governed components with AWS Service Catalog and SageMaker Projects
- Model Registry, Model Explainability, Model Debugger
- Feature Store
- Hyperparameter tuning
- Trials and Experiments with data lineage

Demo

GitHub repo

Secure MLOps using Amazon SageMaker and SageMaker Studio

MLOps resources

Blog posts

- Building secure machine learning environments with Amazon SageMaker
- Secure multi-account model deployment with Amazon SageMaker Series
- Architect and build the full machine learning lifecycle with AWS: An end-to-end Amazon SageMaker demo
- Building, automating, managing, and scaling ML workflows using Amazon SageMaker Pipelines
- Build a CI/CD pipeline for deploying custom machine learning models using AWS services
- Build a Secure Enterprise Machine Learning Platform on AWS
- Machine Learning Best Practices in Financial Services: Whitepaper
- Orchestrate XGBoost ML Pipelines with Amazon Managed Workflows for Apache Airflow
- Managing your machine learning lifecycle with MLflow and Amazon SageMaker
- Create Amazon SageMaker projects using third-party source control and Jenkins
- Model and data lineage in machine learning experimentation
- 5 Lessons Learned Building an Open Source MLOps Platform

Workshops:

- Amazon Sagemaker MLOps workshop GitHub
- Amazon Managed Workflows for Apache Airflow workshop
- Operationalizing the ML pipeline workshop
- Safe MLOps deployment pipeline
- Secure Data Science with Amazon SageMaker Studio Workshop
- MLOps and integrations

Solutions:

- AWS MLOps Framework
- Enterprise AI/ML solutions
- Amazon SageMaker secure MLOps

Others:

- Continuous Delivery for Machine Learning

yevgeniy ilyin
ilyiny@amazon.com

AI&ML CONF