aws

From Idea to Production:

Automate Your Machine Learning Workflows with Pipelines

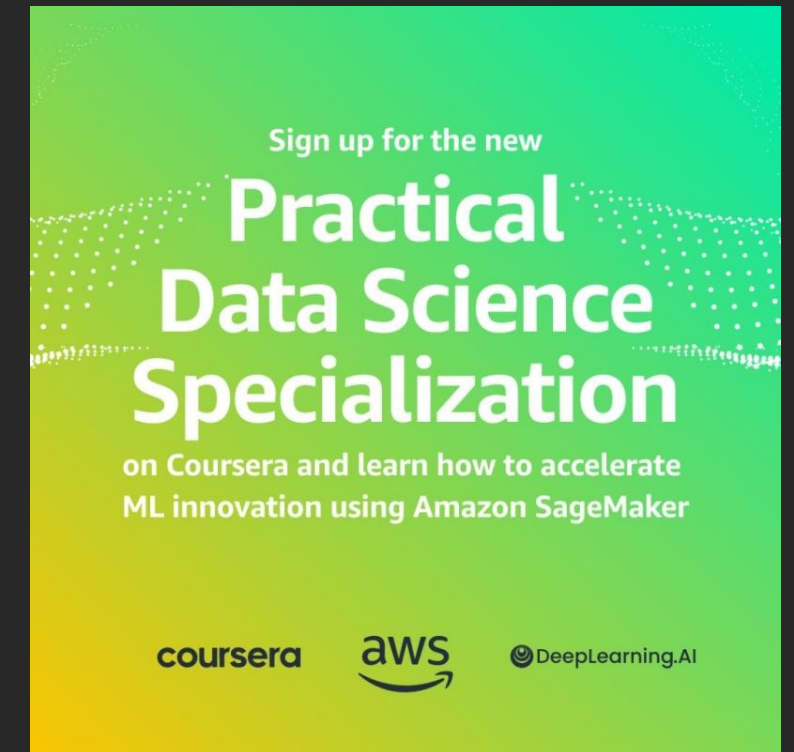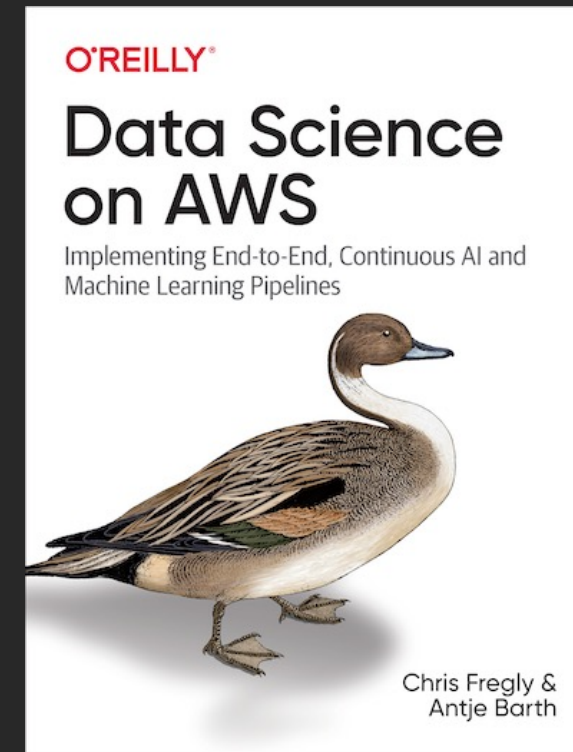**Antje Barth**
Senior Developer Advocate, AI and Machine Learning
Amazon Web Services (AWS)

A few words about me…

Twitter:
@anbarth

LinkedIn:
antje-barth

Developer Advocate AI/ML
O'Reilly Author
Coursera Instructor

O'REILLY®
Data Science on AWS
Implementing End-to-End, Continuous AI and Machine Learning Pipelines

Chris Fregly &
Antje Barth

Sign up for the new
Practical Data Science Specialization
on Coursera and learn how to accelerate ML innovation using Amazon SageMaker

coursera   aws   DeepLearning.AI

https://www.datascienceonaws.com/

https://www.coursera.org/specializations/practical-data-science

Automating ML Workflows?

" Alexa,
    deploy my model! "
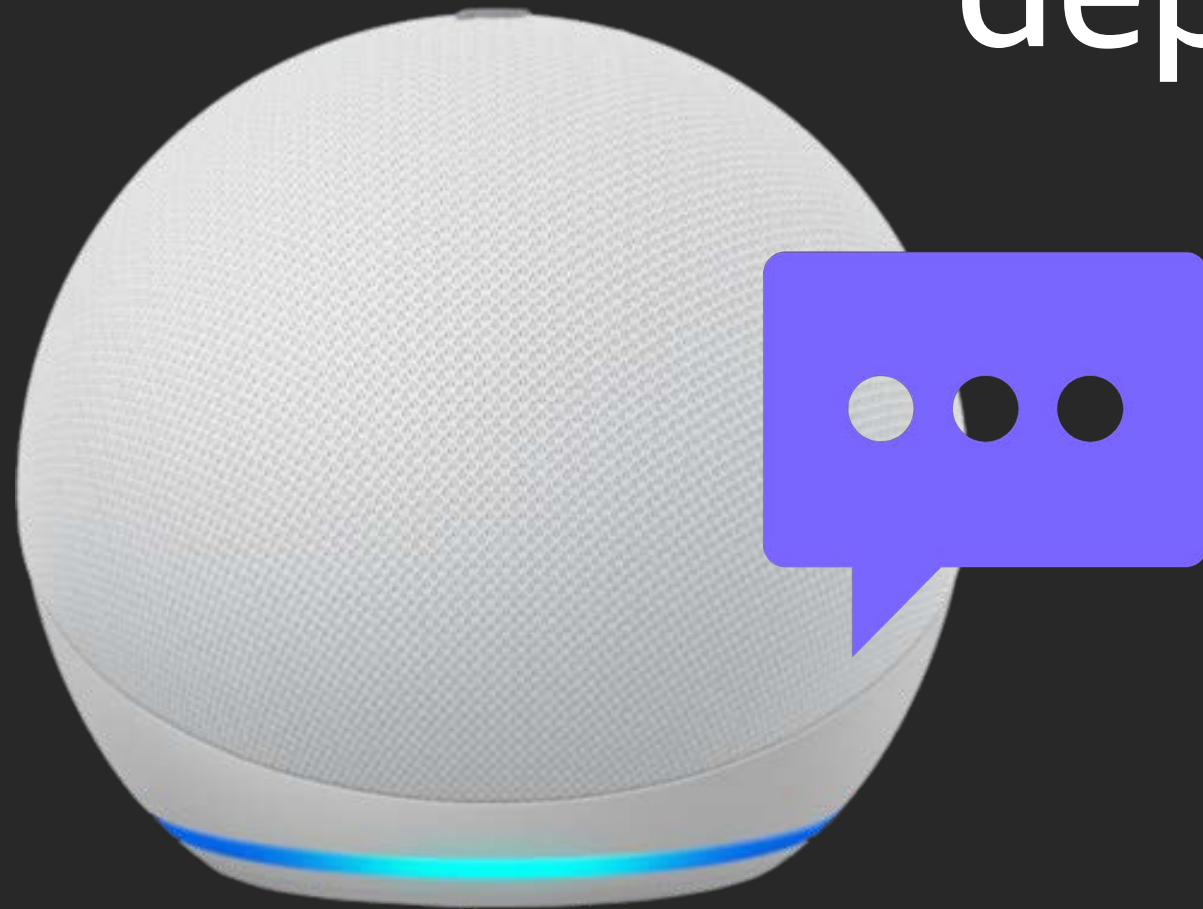
It's not just technology…

aws

Get ML applications ready for production - MLOps
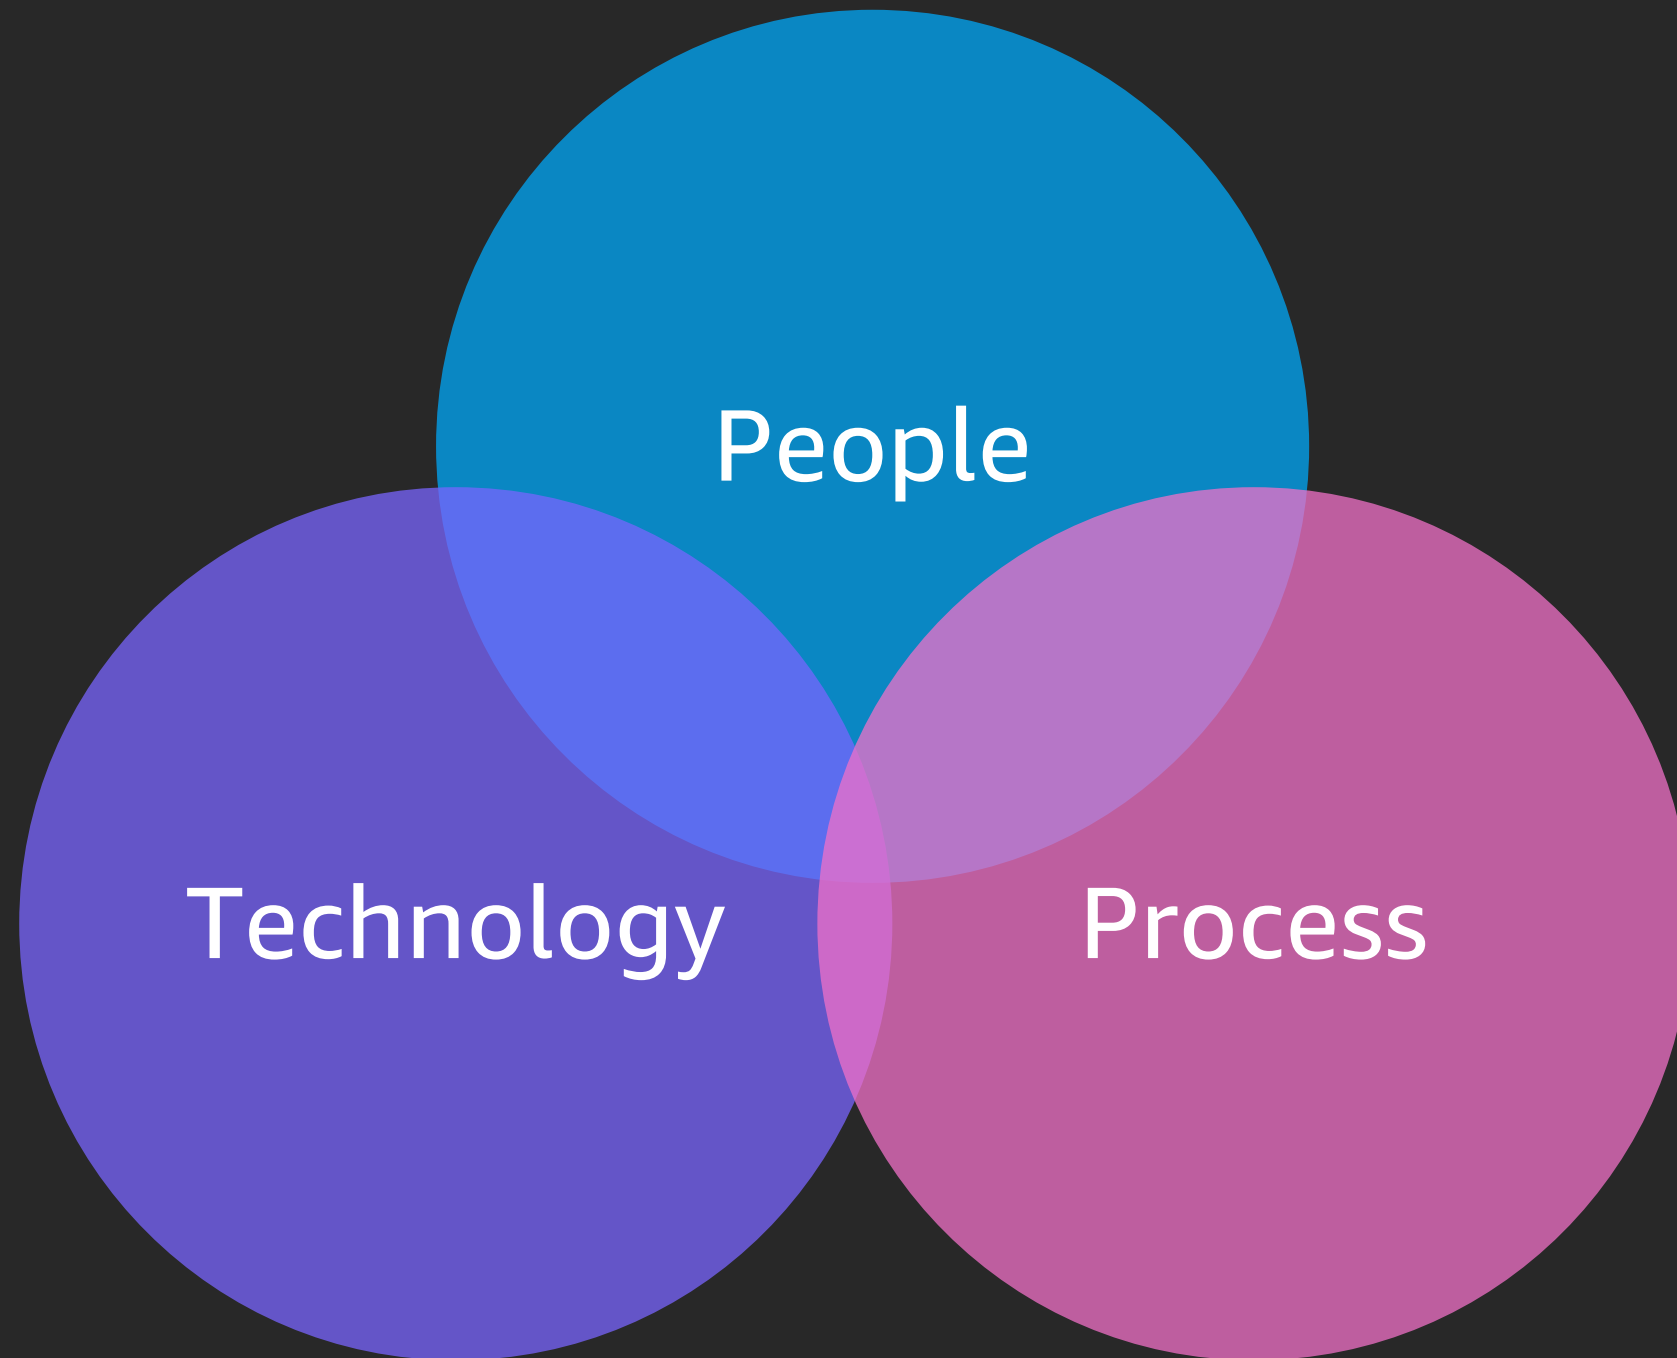
Considerations

Machine Learning Development Lifecycle (MLDC)
!= Software Development Lifecycle (SDLC)

ML model is only a small part of an overall solution

Multiple personas and teams involved

Integration with IT systems and practices

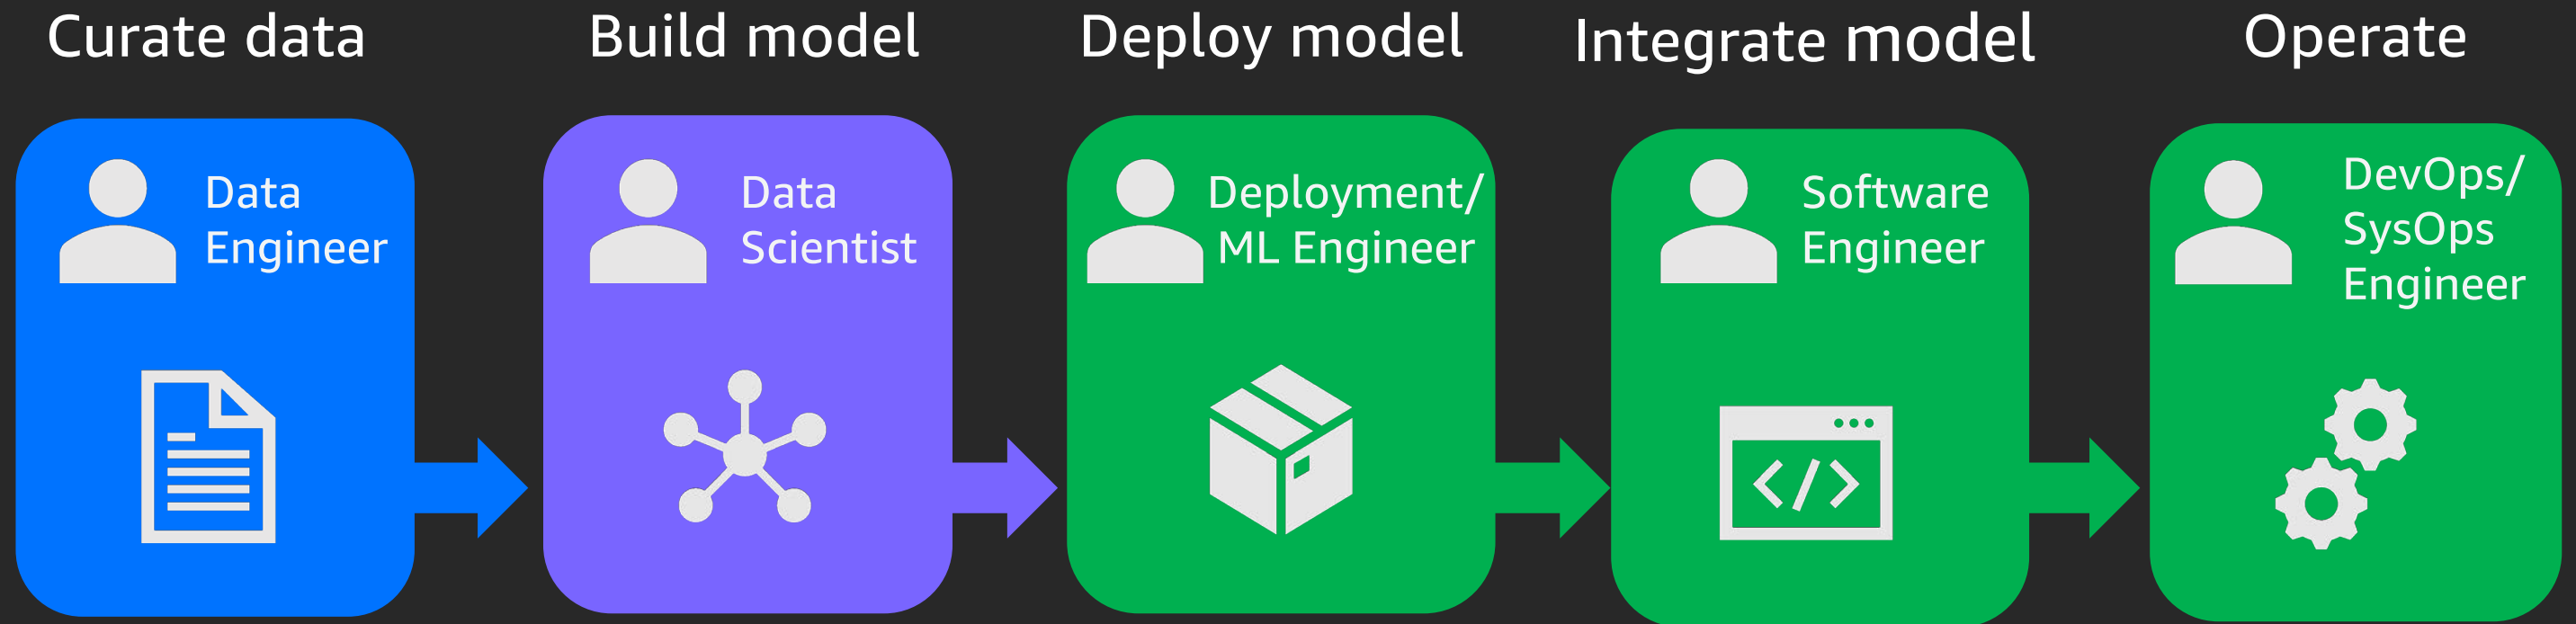# Operationalizing Machine Learning

Goals

- Move away from manually building models { Status Quo }
- Accelerate the path to production { Phase 1 }
- Improve the quality of deployed models { Phase 2 }

Ultimately

- Build resilient, secure, performant, operationally efficient and cost optimized AI/ML solutions

Status Quo:
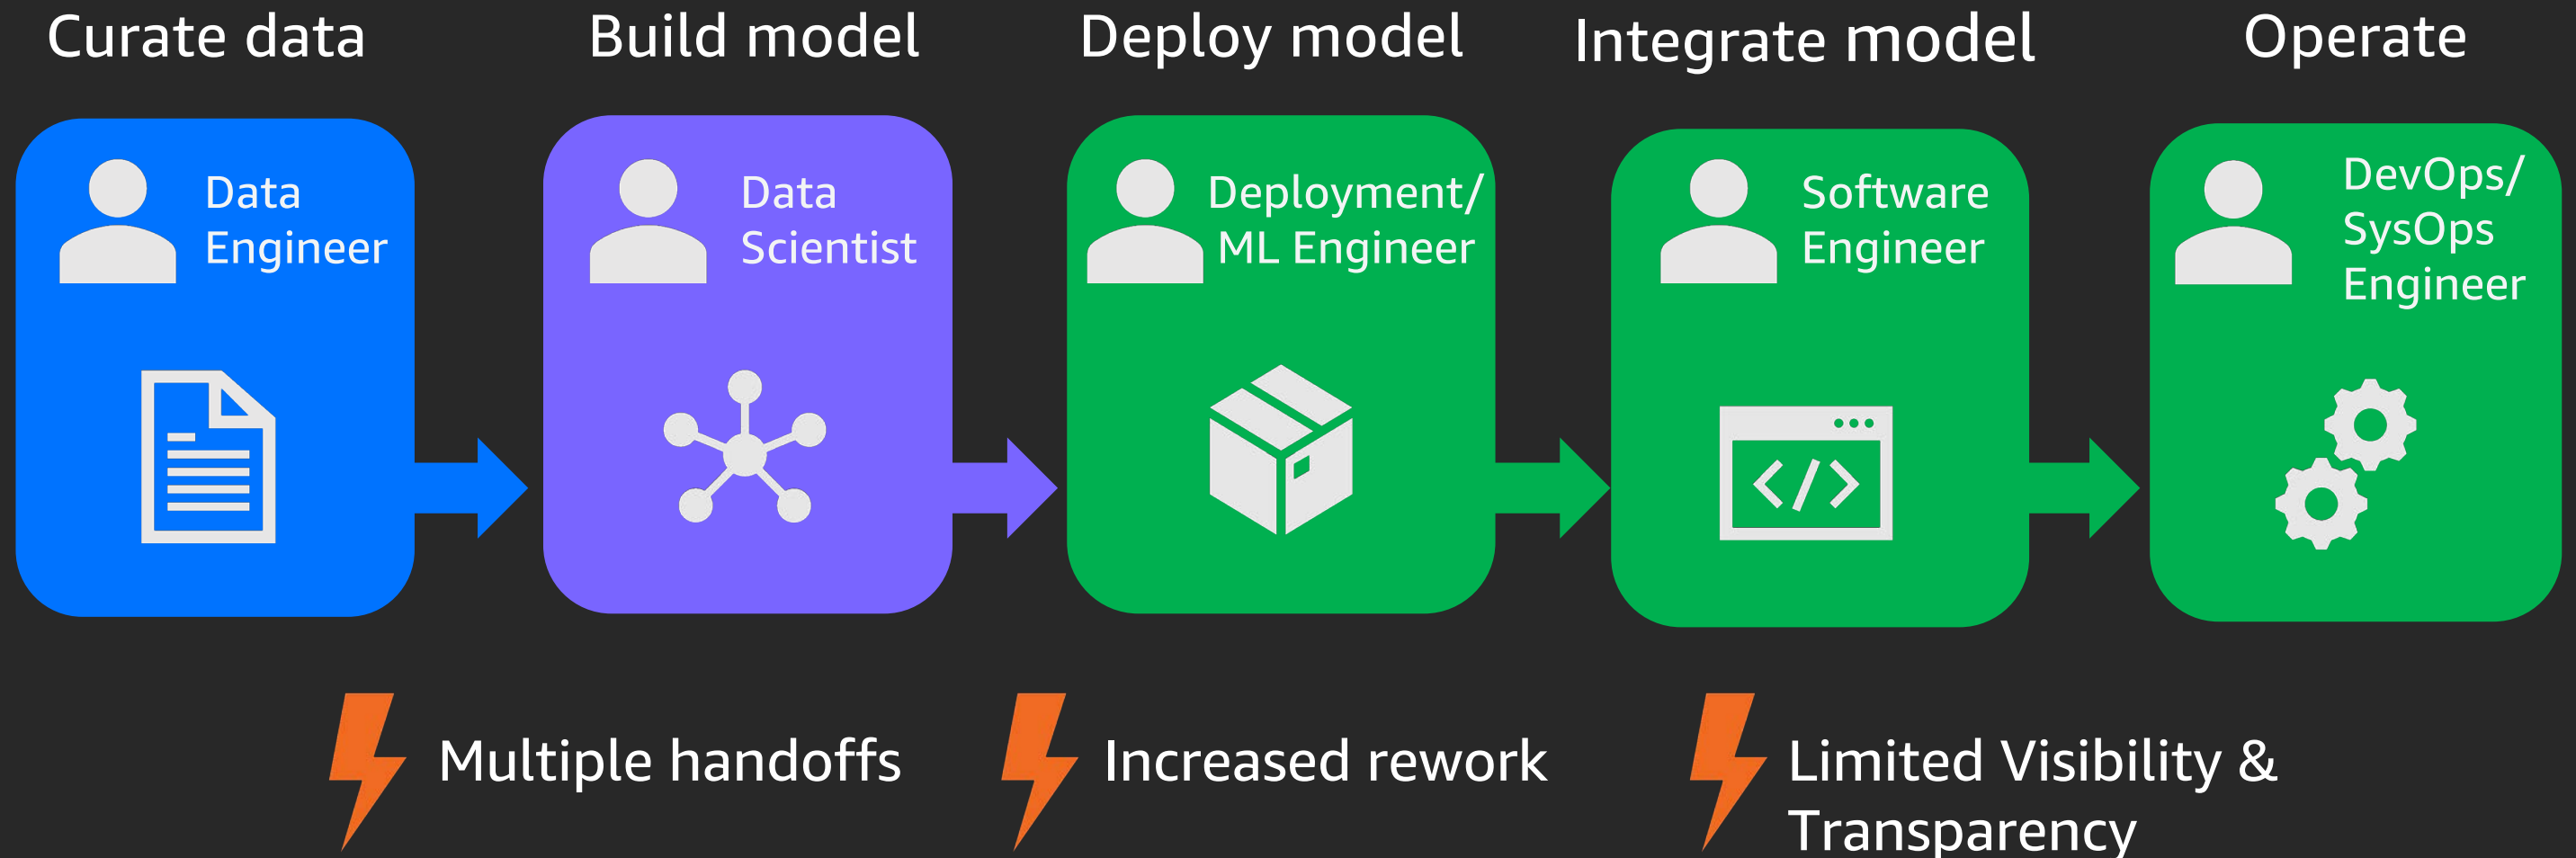## Manually build models, individually executed steps

Curate data — Data Engineer

Build model — Data Scientist

Deploy model — Deployment/ML Engineer

Integrate model — Software Engineer

Operate — DevOps/SysOps Engineer

Status Quo:
Manually build models, individually executed steps

Curate data	Build model	Deploy model	Integrate model	Operate
Data Engineer	Data Scientist	Deployment/ ML Engineer	Software Engineer	DevOps/ SysOps Engineer

Multiple handoffs   Increased rework   Limited Visibility & Transparency

Phase 1: Orchestrate steps as pipelines and automate tasks within each step

Curate data
Data Engineer

Build model
Data Scientist

- Prepare Data
- Train Model
- Evaluate Model

Deploy model
Deployment/ ML Engineer

Register models

Integrate model
Software Engineer

CI/CD

Operate
DevOps/ SysOps Engineer

Monitor model

Phase 1: Orchestrate steps as pipelines and automate tasks within each step

Curate data

Build model

Deploy model

Integrate model

Operate

Data Engineer

Data Scientist

Deployment/ ML Engineer

Software Engineer

DevOps/ SysOps Engineer

Prepare Data

↓

Train Model

↓

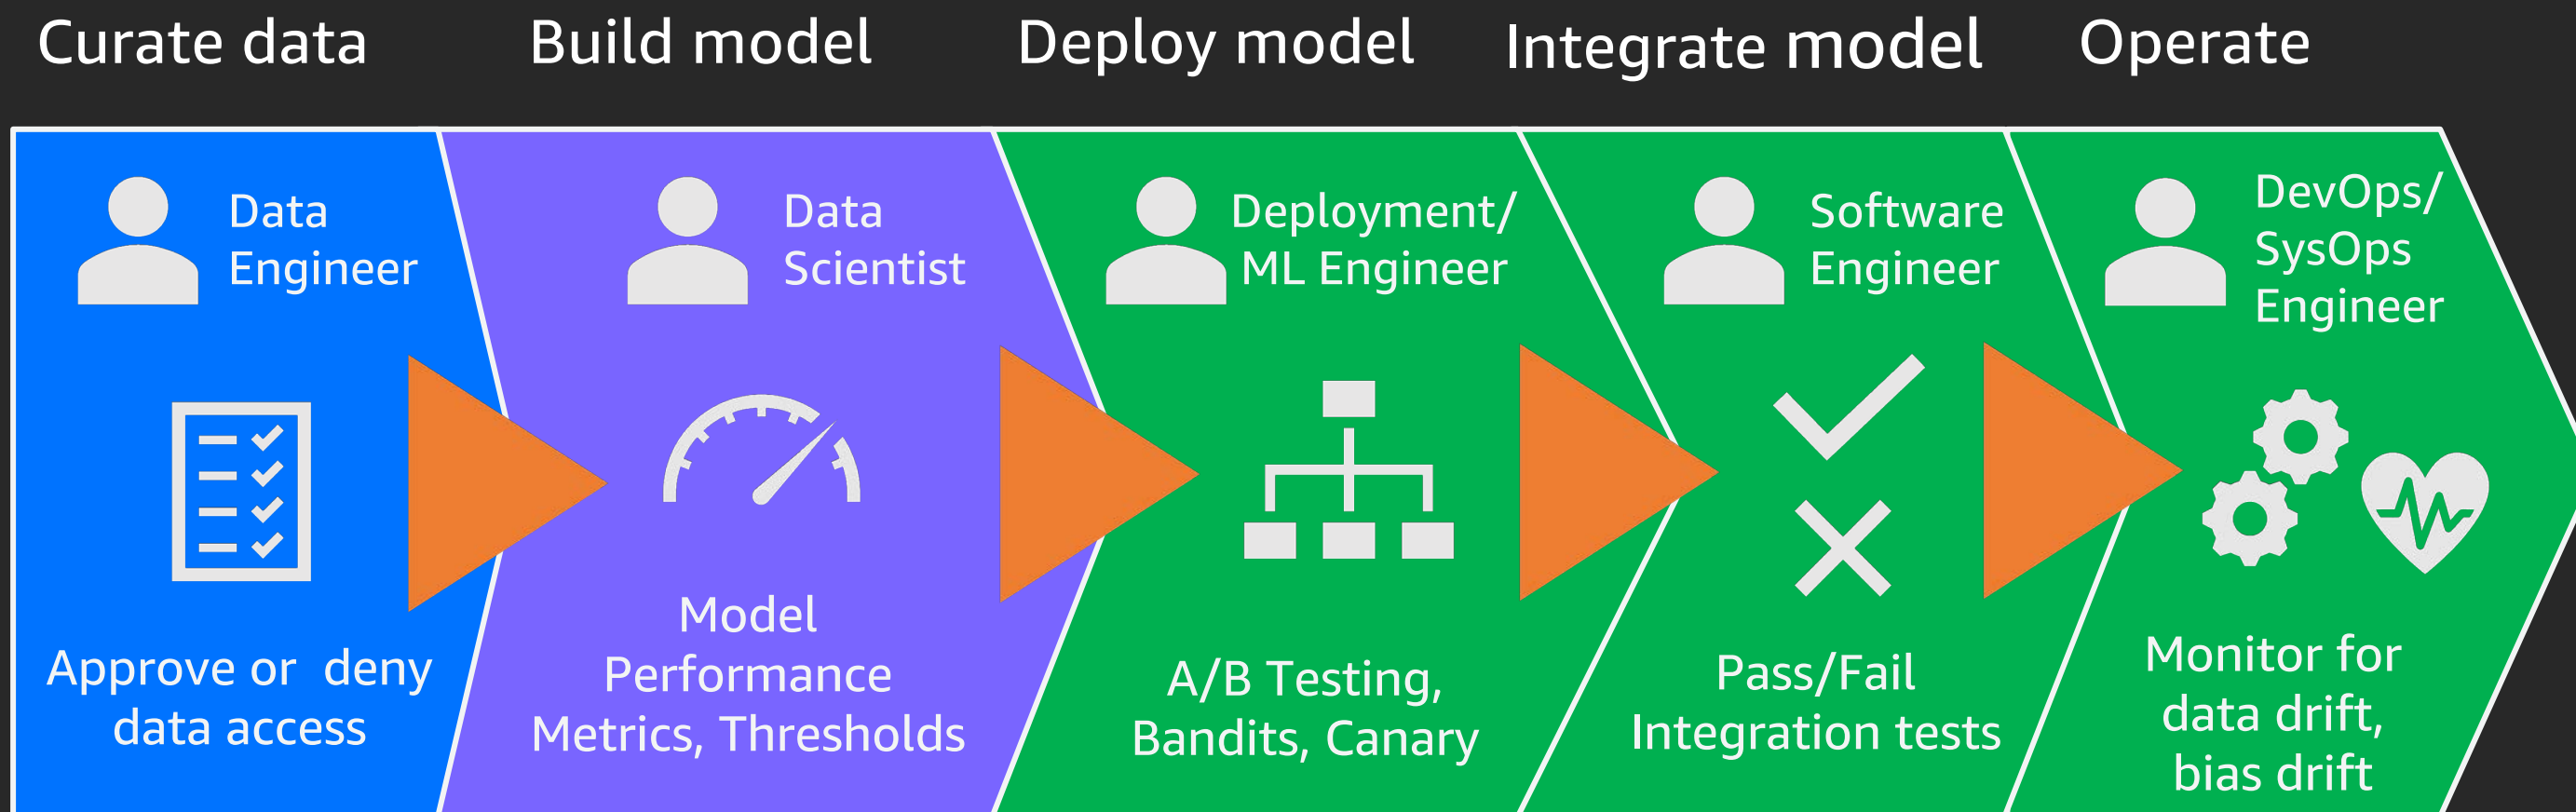Evaluate Model

Register models

CI/CD

Monitor model

⟳ Deployment Pipeline

⟳ Model Training/Retraining Pipeline

Demo
Build a model training pipeline

Phase 2: Automatically run pipelines and include automated quality gates

Curate data

Data Engineer

Approve or deny data access

Build model

Data Scientist

Model Performance Metrics, Thresholds

Deploy model

Deployment/ ML Engineer

A/B Testing, Bandits, Canary

Integrate model

Software Engineer

Pass/Fail Integration tests

Operate

DevOps/ SysOps Engineer

Monitor for data drift, bias drift

Phase 2: Automatically run pipelines and include automated quality gates

Curate data

Data Engineer

Approve or deny data access

Build model

Data Scientist

Model Performance Metrics, Thresholds

Deploy model

Deployment/ ML Engineer

A/B Testing, Bandits, Canary

Integrate model

Software Engineer

Pass/Fail Integration tests

Operate

DevOps/ SysOps Engineer

Monitor for data drift, bias drift

Demo
Re-run pipeline on code change

Build AI/ML solutions

Production-grade ML platform and systems integration

Security

Reliability
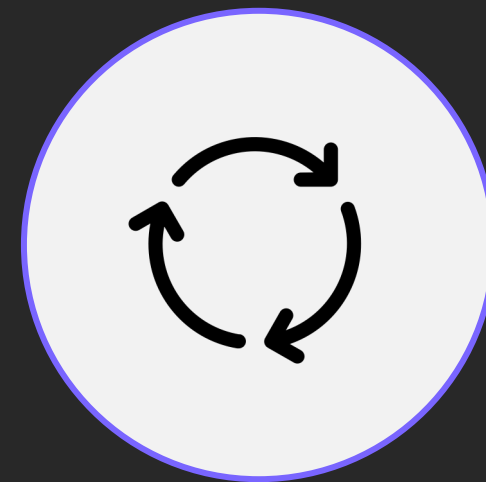
Cost Optimization

Performance Efficiency

Operational Excellence

Get Started

Data Science on AWS
https://www.datascienceonaws.com
https://github.com/data-science-on-aws/workshop

Amazon SageMaker Pipelines
https://aws.amazon.com/sagemaker/pipelines/

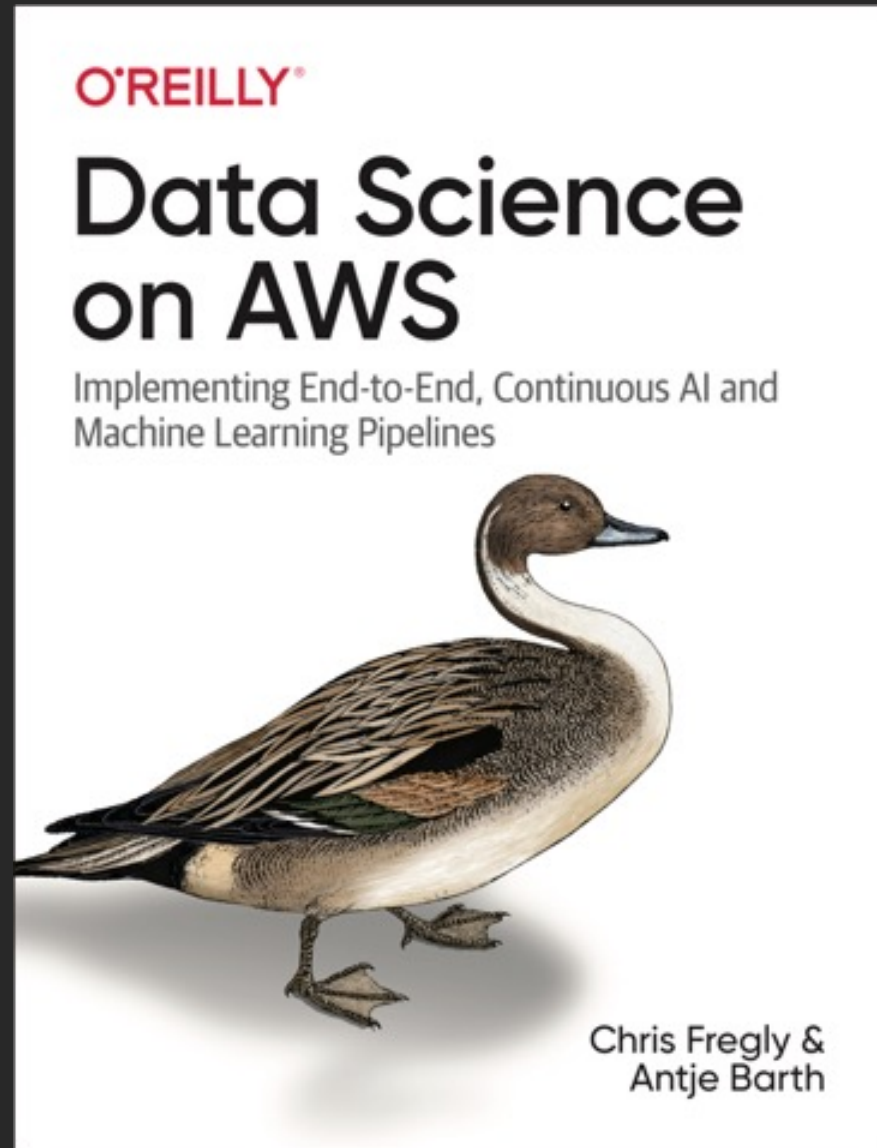https://github.com/aws/amazon-sagemaker-examples/tree/master/sagemaker-pipelines

Blog
https://aws.amazon.com/blogs/machine-learning/building-automating-managing-and-scaling-ml-workflows-using-amazon-sagemaker-pipelines

Learn more

O'Reilly Book "Data Science on AWS"

Implementing End-To-End, Continuous
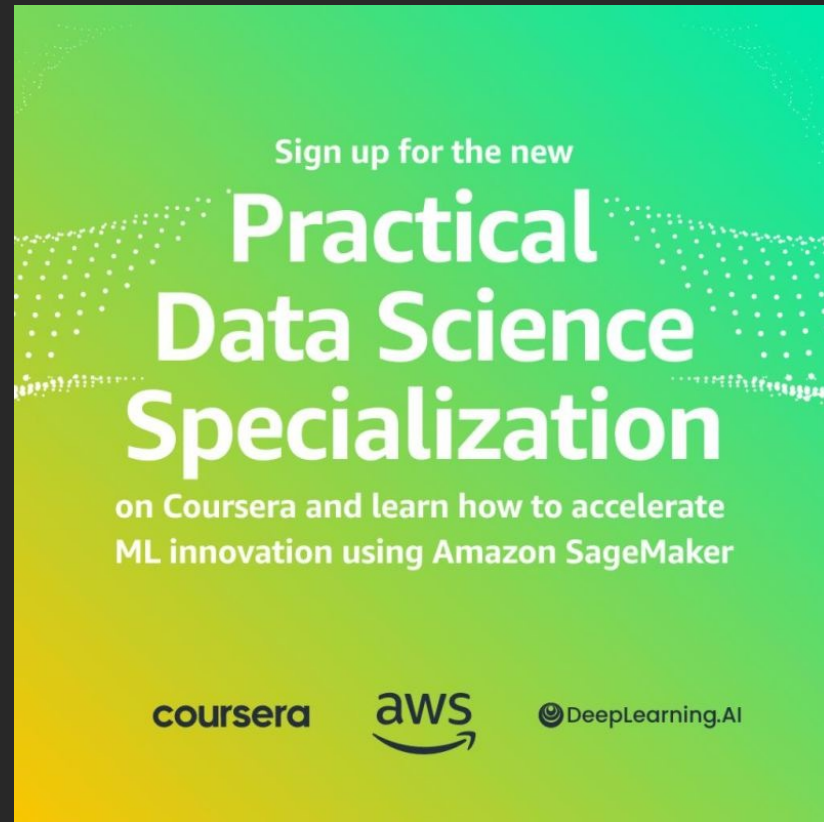AI and Machine Learning Pipelines

    12 chapters
    500 pages
    100's of code samples

https://www.datascienceonaws.com/

Learn more

DeepLearning.AI

Practical Data Science Specialization on Coursera

3 courses
10 weeks of content
Video lectures, quizzes, hands-on labs

https://www.coursera.org/specializations/practical-data-science

Thank you!

Antje Barth

[in] antje-barth

[twitter] @anbarth

aws