From Problem to Production: DevOps for AI with Kubeflow

Fabrizio Lazzaretti & Marco Crisafulli

Fabrizio Lazzaretti
Managing Consultant @ Wavestone |
CNCF Ambassador | Speaker | Author

Helping enterprises adopt cloud-native technologies while ensuring future-proof architectures and business-IT alignment

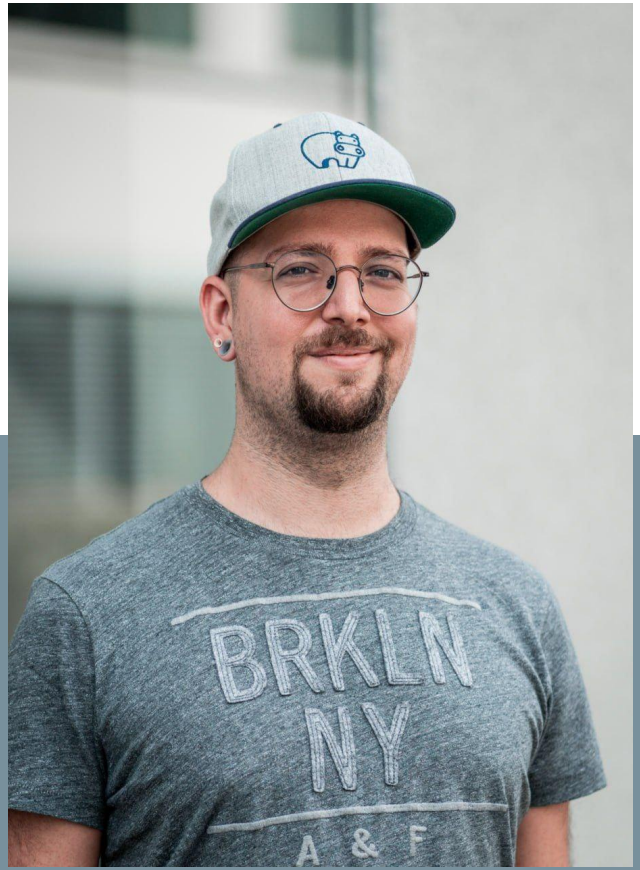

Marco Crisafulli
Co-Founder @ enki |
ML Engineer | Father | Sometimes funny

AI can solve real problems. Helping companies to understand and utilize this technology.

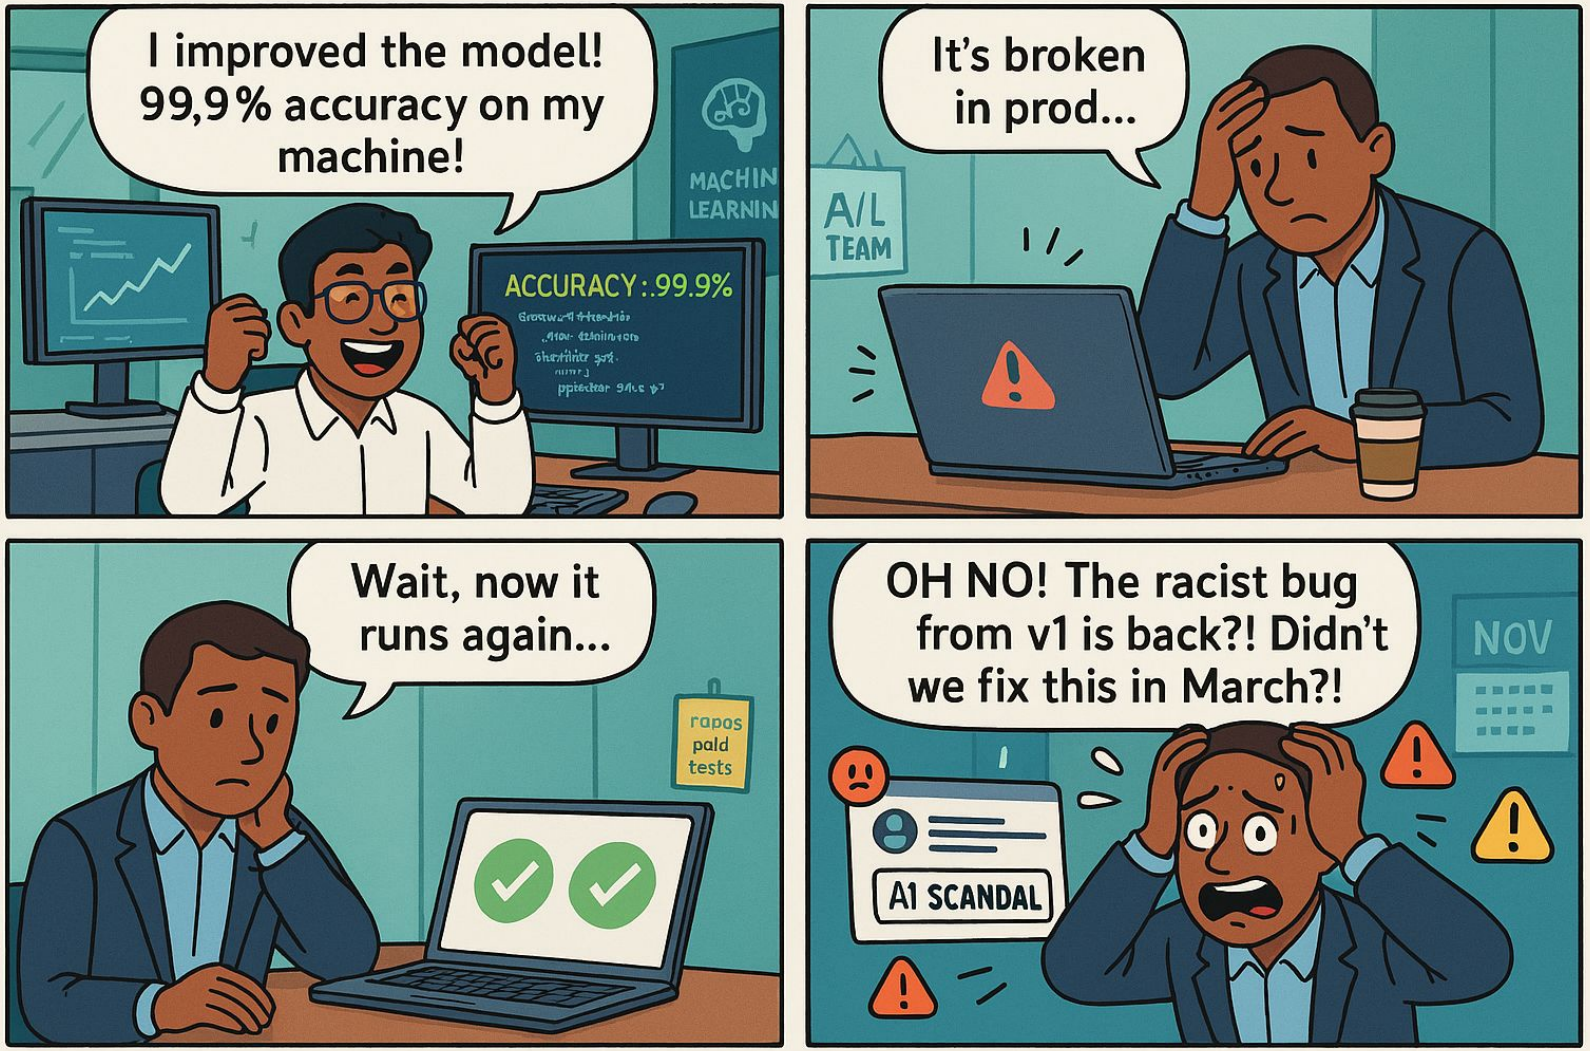
I improved the model! 99,9% accuracy on my machine!
ACCURACY:.99.9%
It's broken in prod...
A/L TEAM
Wait, now it runs again...
rapos pald tests
OH NO! The racist bug from v1 is back?! Didn't we fix this in March?!
A1 SCANDAL

Why is this needed?

“The main challenges people face when developing ML capabilities are scale, version control, model reproducibility, and aligning stakeholders”

2020 State of Enterprise Machine Learning Report

From problem to solution

1. We will not solve all your AI problems in this talk
2. You will not outperform the big players, but you can learn from them

Our goal for today:

- Understand your business, avoid wrong expectations, and help the business understand IT
- Reproducibility, testing, and continuous deployment with a practical implementation in Kubeflow

Before we start…

CLOUD NATIVE COMPUTING FOUNDATION

We need AI!

- What **business problem** do we want to solve?
- Which **KPIs** can we define to determine if the project was successful?
- Do we have the **necessary data**, **know-how** and **legal framework** to do this?
- Does the problem **benefit** from applying **AI**?

enki

How can we find out if we need this AI project?

- **Qualify business ideas**
  Understand what we want to do, and what to prioritize
  - Business Model Canvas, Capability Map, Wardley Map
- **Gathering business requirements**
  Understand how we need to do it and get a common language and understanding
  - Domain Storytelling, Visual Glossary, Event Storming, Context Map
- **Tactical design** of the components
  Define how we structure the solution
  - Bounded Context Canvas, API Product Canvas

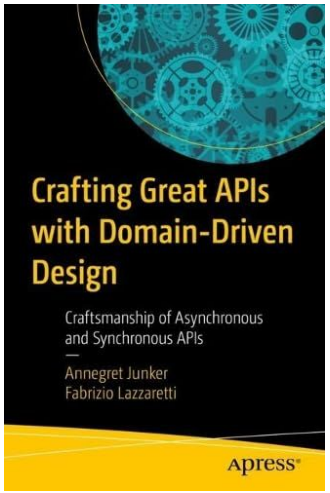

Getting started with the project

The 4 phases

Awareness & Active	Active	Operational	Systemic & Transformational	
PoC 🥰	MVP 🕵️	Product Strategy 🎯	Scaled Approach 📈	
Minimal Concept	Concept	Concept	Concept	Governance
	Data Quality	Data Quality	Data Quality	Disaster Recovery & HA
	Compliance	Compliance (NIST AI RMF, ISO 42001)	Compliance (NIST AI RMF, ISO 42001)	Cost Optimization
	Minimal SLAs	SLAs	SLAs (99.9%+ availability)	Training Plan
	Basic Monitoring	Monitoring	Monitoring	Observability
		Evolution (A/B testing infrastructure)	Evolution (Canary deployment)	

From concept to implementation

End-to-end ML, right?

Load Data

Magic (training)

Why is this needed?

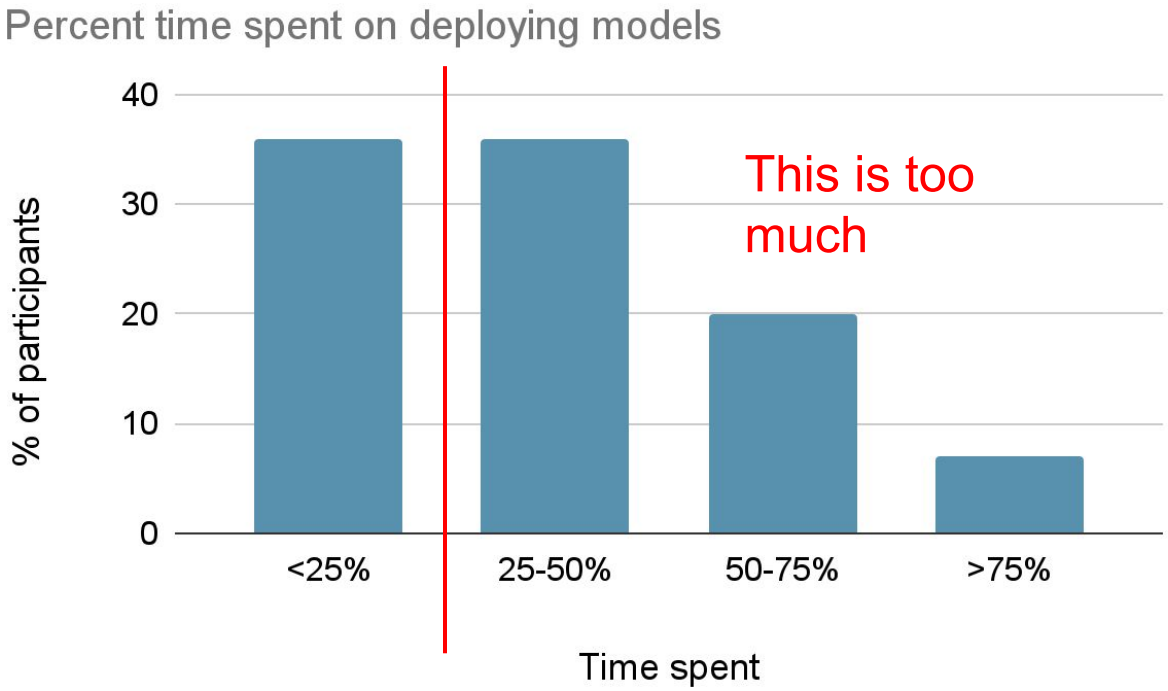

Src: 2020 State of Enterprise Machine Learning Report

MLOps is not DevOps

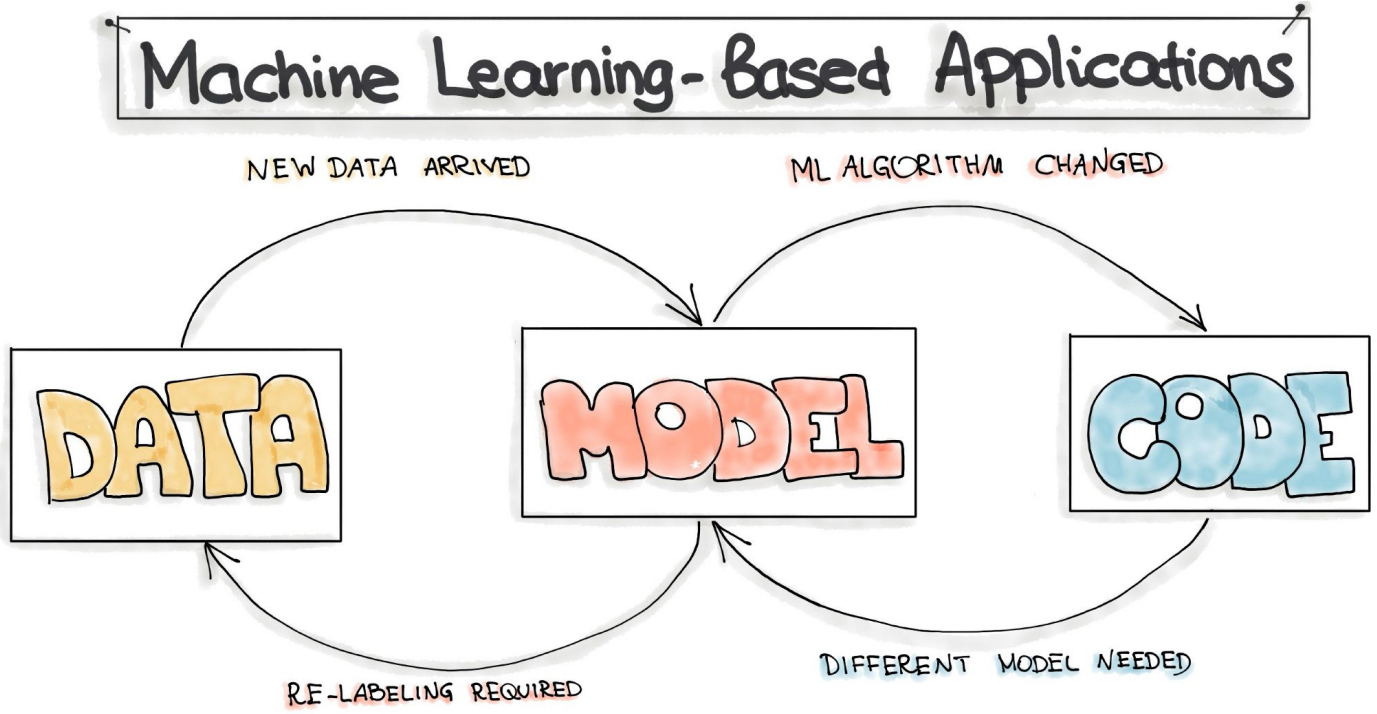

ImgSrc: https://ml-ops.org

MLOps in detail

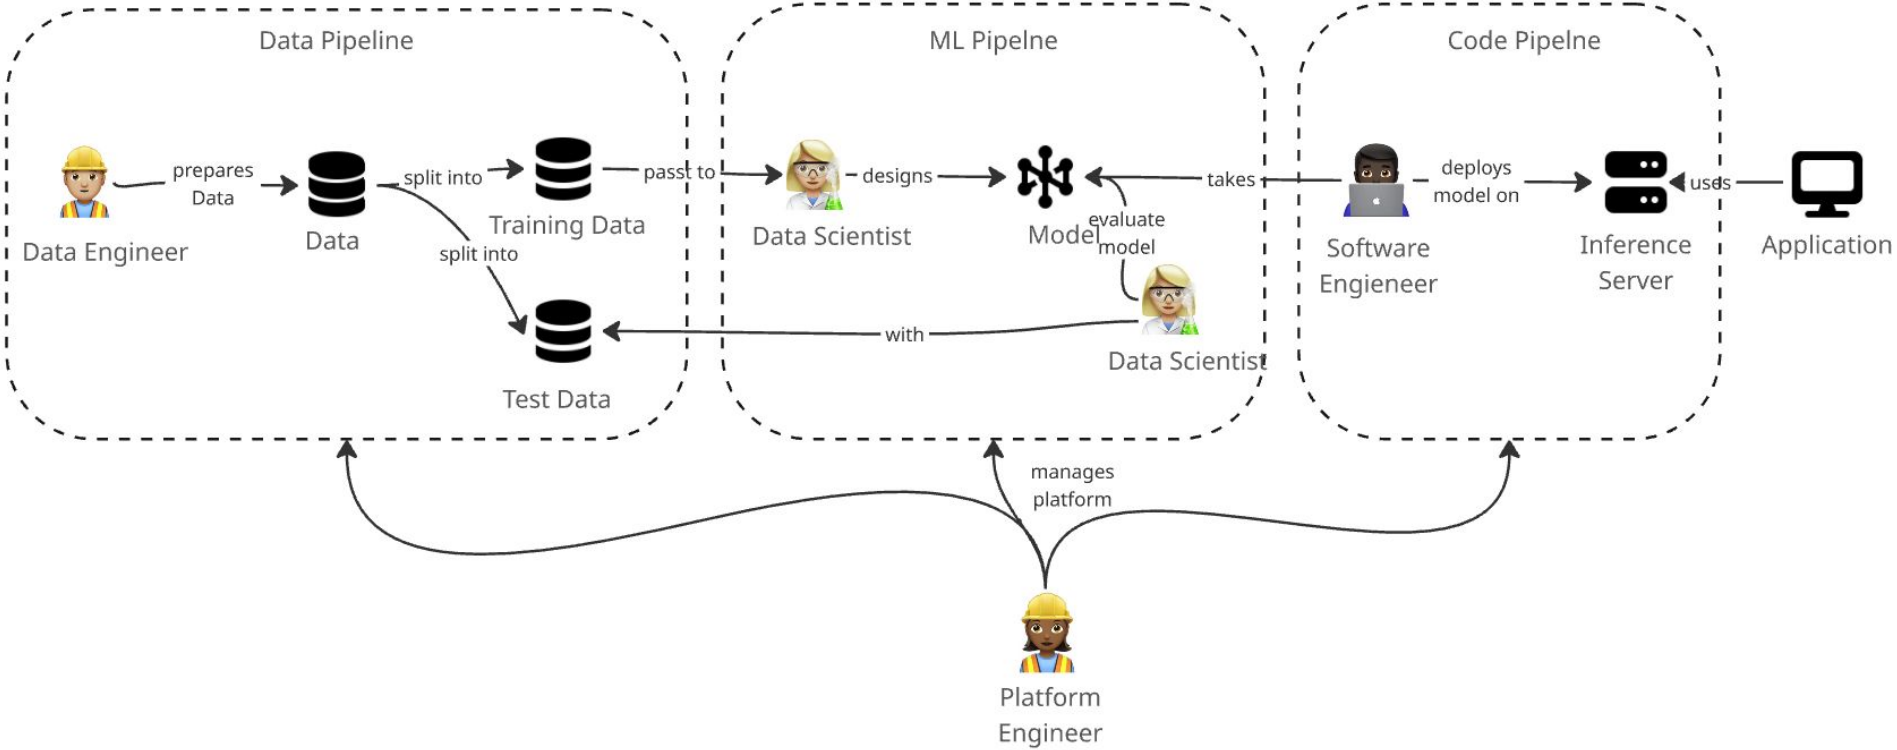

Applying MLOps with Kubeflow

Kubeflow an incubator project with traction

FIGURE 14

CNCF INCUBATED PROJECTS IN USE OR IN EVALUATION

Which of these incubating CNCF projects is your organization using in production or evaluating?

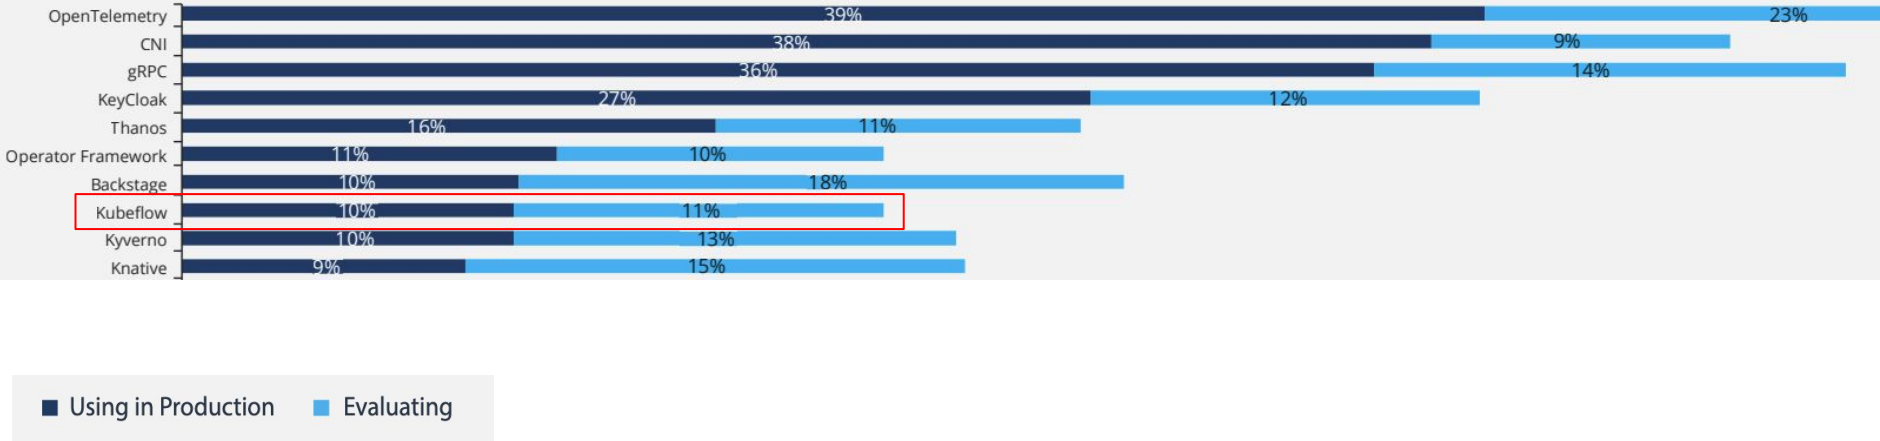

Src: https://www.cncf.io/wp-content/uploads/2025/04/cncf_annual_survey24_031225a.pdf

Who uses Kubeflow?

The Kubeflow user survey [2023] drew responses from 90 members of the community mostly made up of members from the United States (43%), Europe (34%), and Asia-Pacific (10%).

The majority of the respondents were from the Tech industry (49%), followed by Finance (13%) and Consulting (11%).

Src:https://blog.kubeflow.org/kubeflow-user-survey-2023/

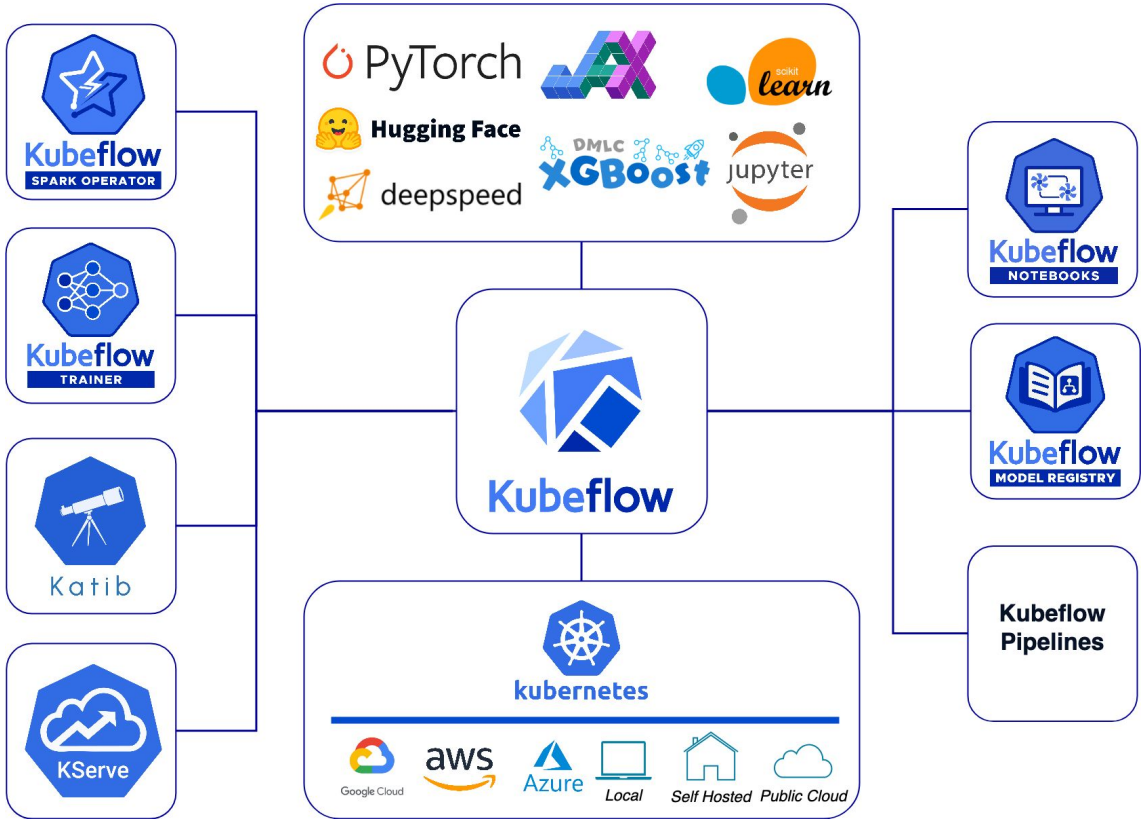

ImgSrc: https://www.kubeflow.org/docs/started/introduction/

Kubeflow and an own model

MLOps with Kubeflow

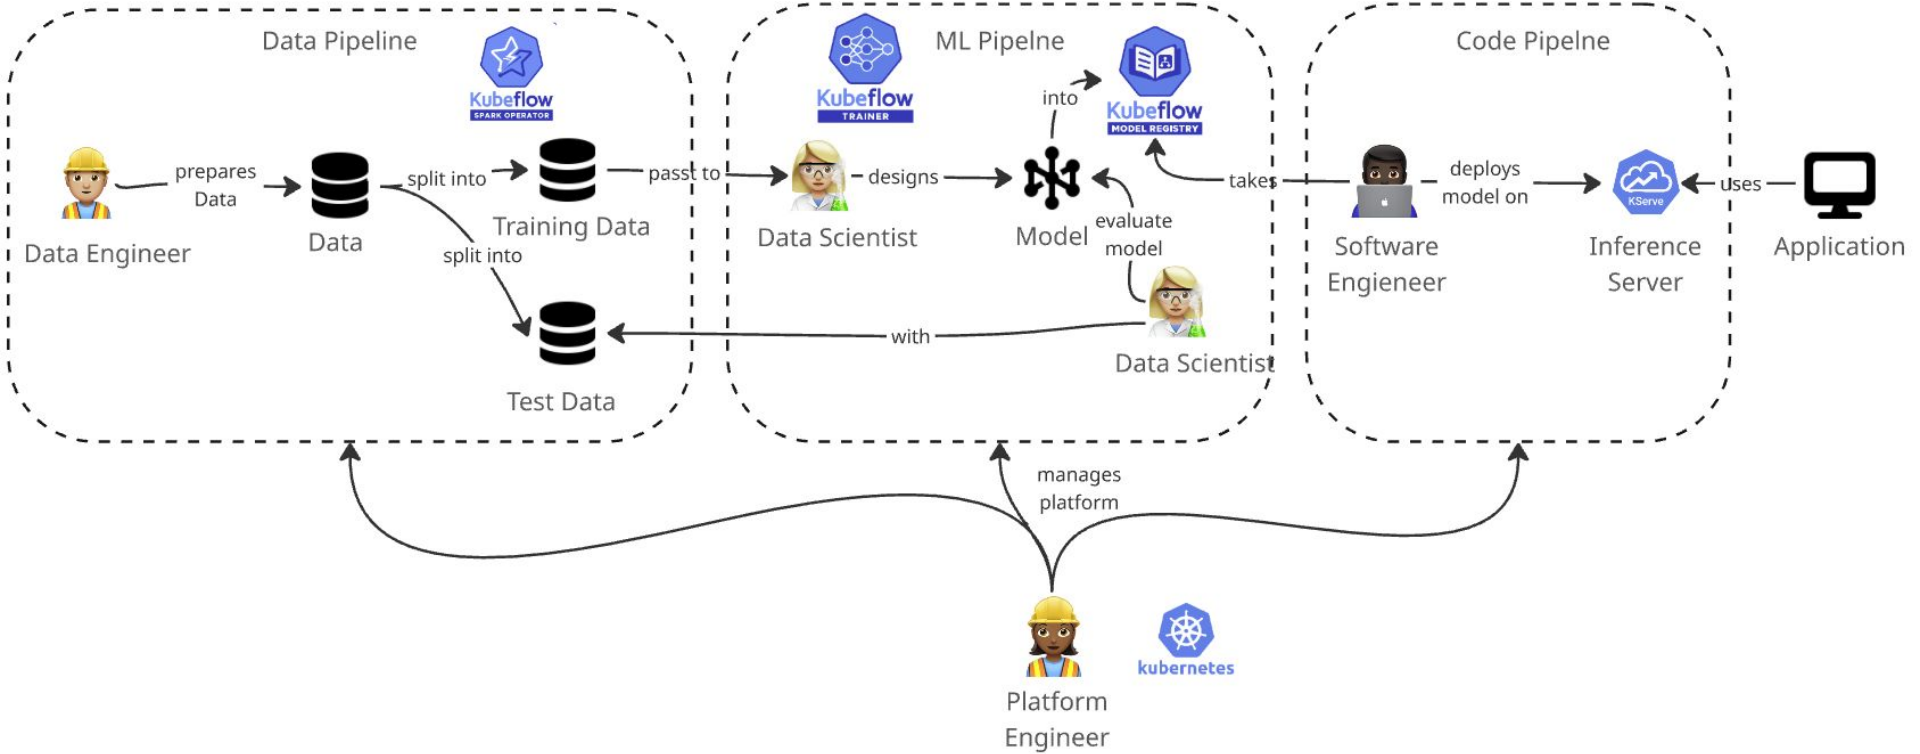

MLOps with Kubeflow

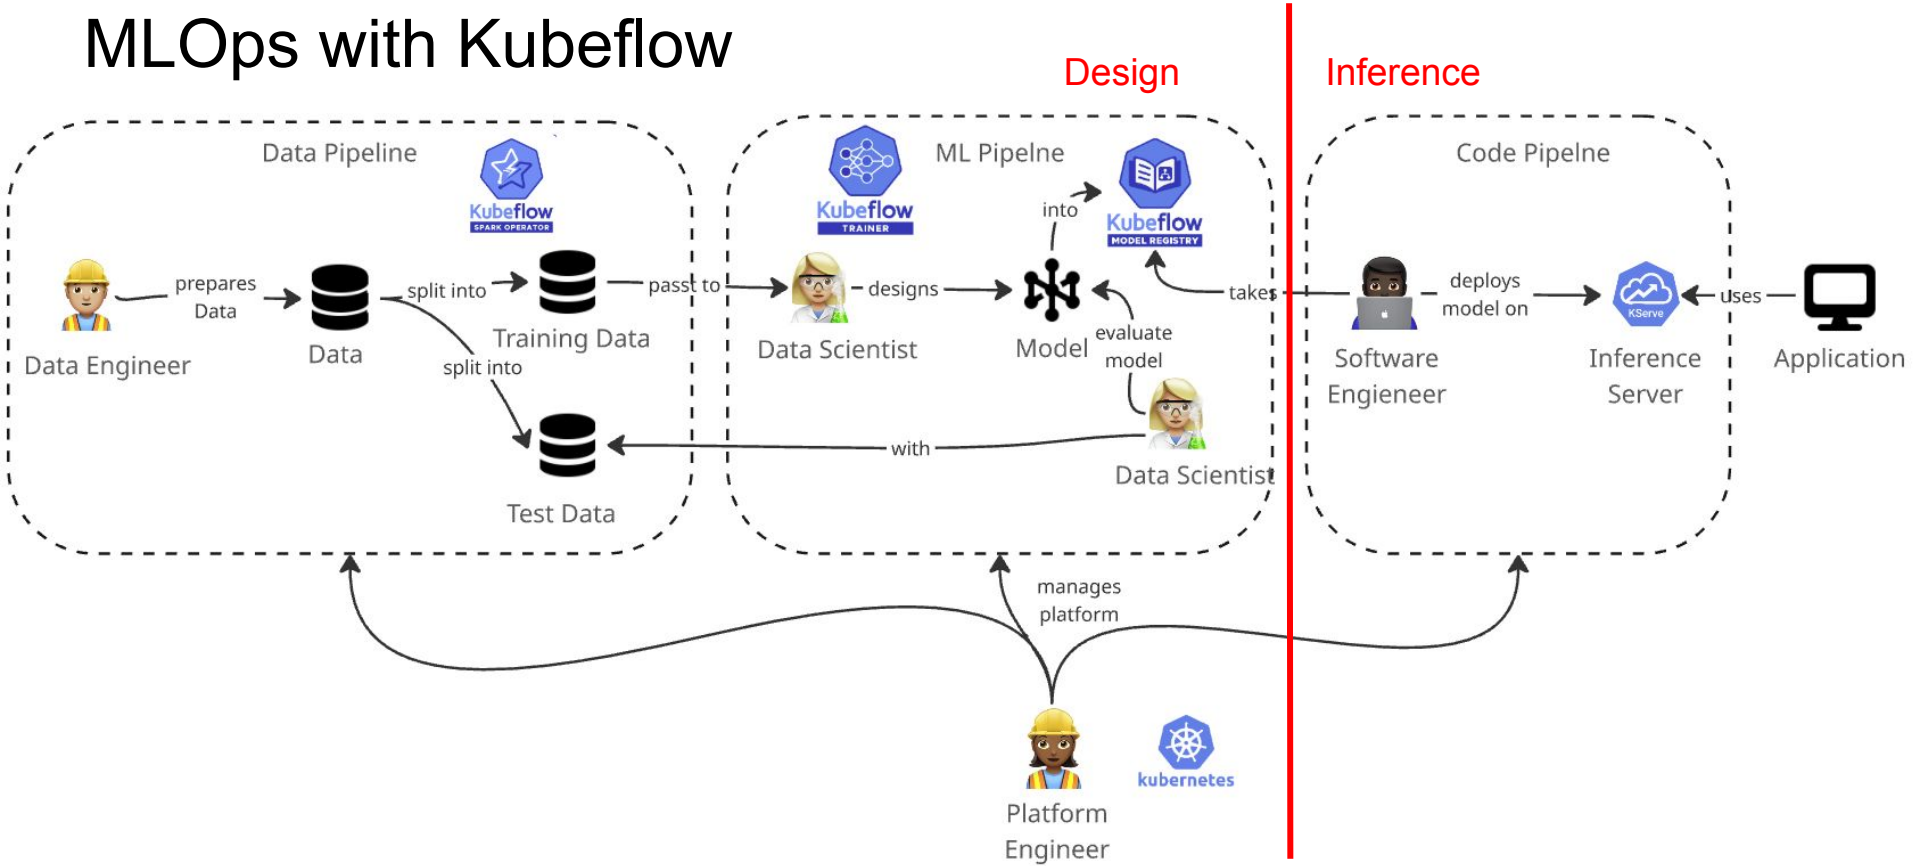

MLOps with Kubeflow

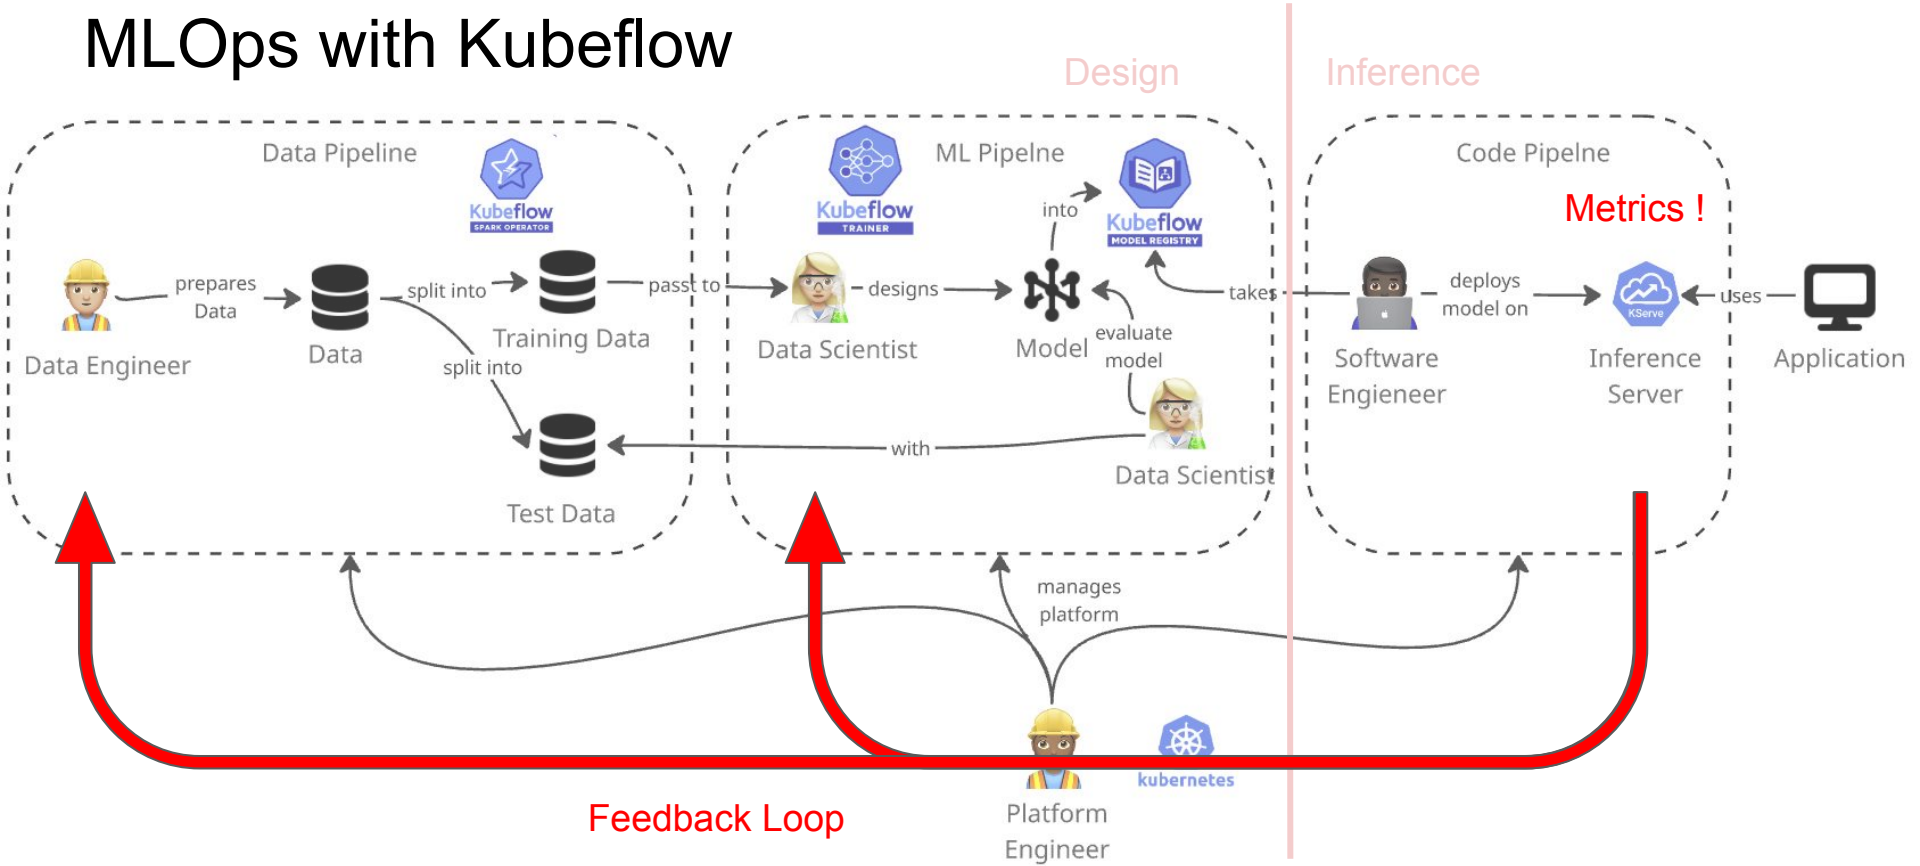

I want an off the shelf model (ChatGPT, ..)

MLOps with off the shelf models

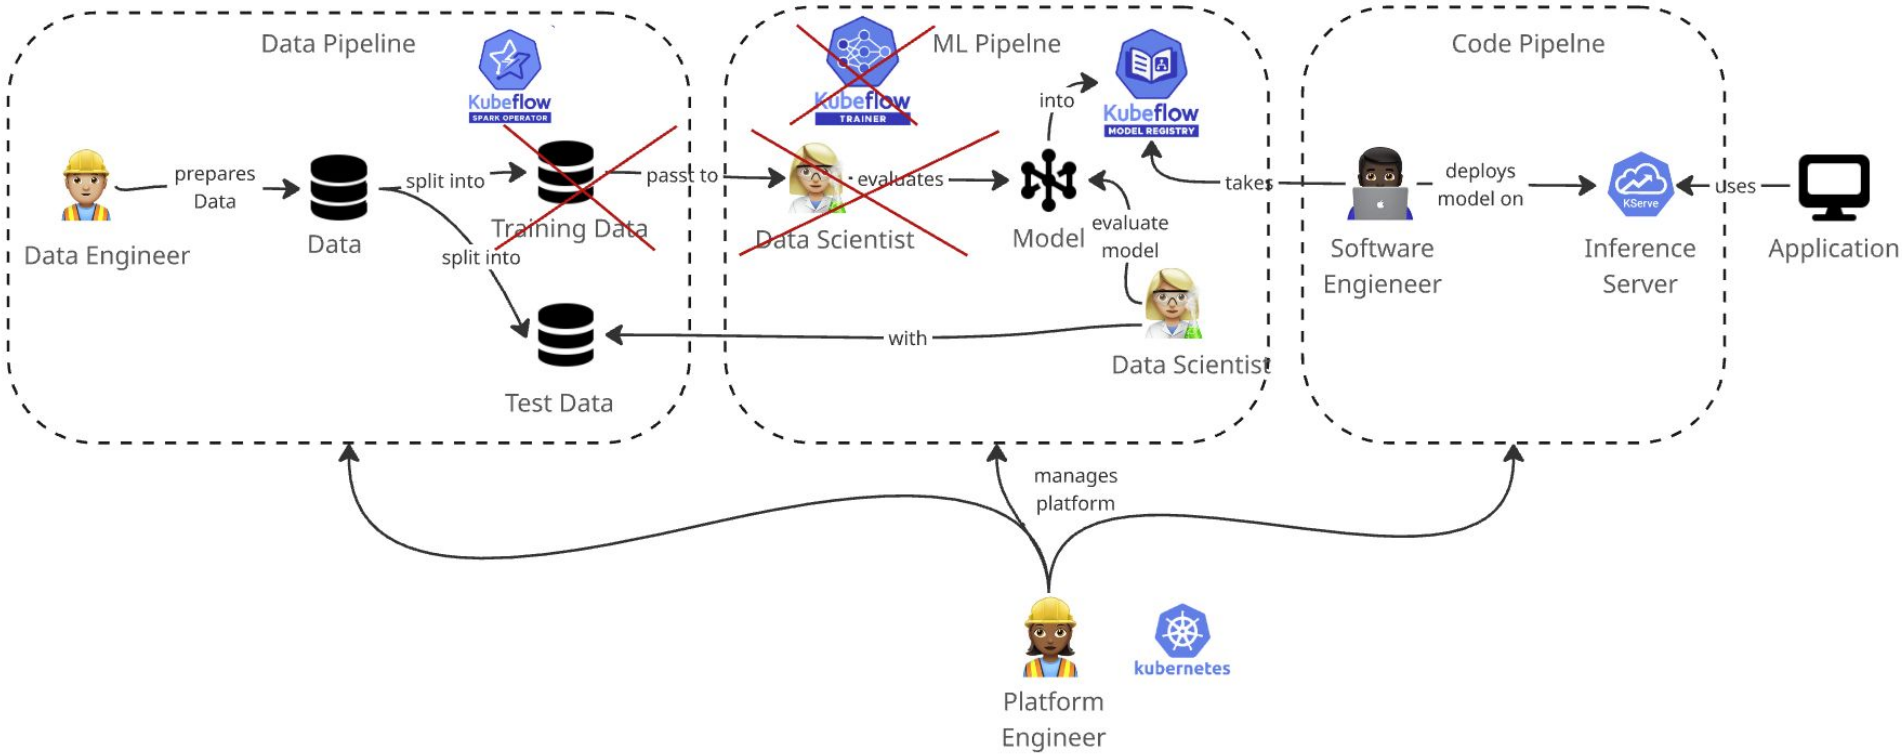

When to use Kubeflow?

Pros and cons

Pros	Cons
Cloud-native, Kubernetes-integrated	High operational complexity
Highly scalable & production-ready	Steep learning curve
Modular, extensible components	Difficult upgrades/maintenance
Vendor-neutral, open source	Heavy infrastructure overhead
Works across clouds & on-prem	Slow for rapid prototyping
Strong community & ecosystem	Documentation needs improvement

Lessons learned

- Start small
  - Kubeflow has many components -> Start with one
  - Be careful about over-engineering
- Run a dedicated cluster
  - Use managed services if possible
- Manage Kubernetes manifests with DevOps practices
  - e.g kustomize
- It is not a magic platform
  - You still need to think about the process

Conclusion

Don’t throw out what you learned from classical IT projects

- Requirement engineering
- Phases
- KPIs

At the end people need to talk!
DevOps + MLOps + Any-Ops

Always evaluate your results, outcome can change dramatically after small changes!

https://lazzaretti.me/cnc26

Fabrizio
https://www.linkedin.com/in/fabrizio-lazzaretti/

Marco
https://www.linkedin.com/in/mcrisafu/

Follow up

Wavestone: Beyond traditional change management whitepaper:
https://www.wavestone.com/en/insight/beyond-traditional-change-management/

Kubeflow Introduction:
https://www.kubeflow.org/docs/started/introduction/

Kubeflow training from CNCF:
https://trainingportal.linuxfoundation.org/courses/introduction-to-aiml-toolkits-with-kubeflow-lfs147