SKILUP SM DAYS

by: DevOps Institute
ADVANCING THE HUMANS OF DEVOPS

Taming Data Science Dragons with MLOps and Kubernetes

In Praise of Data Science - Helping solve complex challenges

Clinical diagnostics

Citizen interaction

Financial fraud

Data Science - harnessing today's tech makes it possible

"Data scientists create business value when they quickly develop AI/ML ...
...that value is only released when AI/ML runs securely in production"

Data Science - by the numbers

25%

Year over year increase in the use of AI in standard business processes*

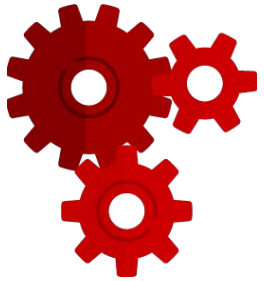

63%

Revenue increases from AI adoption in the business units where their companies use AI*

$98B

Projected worldwide spending on AI - 2023**

*McKinsey – Global AI survey, November 2019
**IDC -- Worldwide AI spending guide, 2020

THE BIG BUT....

87% ...of data science projects never make it into production

80% ...of Analytics projects will not deliver business outcomes through 2022

77% *...of businesses report that "business adoption" of big data and AI initiatives continues to represent a big challenge*

https://blogs.gartner.com/andrew_white/2019/01/03/our-top-data-and-analytics-predicts-for-2019/
https://venturebeat.com/2019/07/19/why-do-87-of-data-science-projects-never-make-it-into-production/
http://newvantage.com/wp-content/uploads/2018/12/Big-Data-Executive-Survey-2019-Findings.pdf

Data Science - Business enabler or Hocus Pocus?

80%

...of AI projects will...

"remain alchemy, run by wizards, whose talents will not scale in the organisation"

Gartner - Top Predictions for IT Organisations and Users in 2019 and Beyond

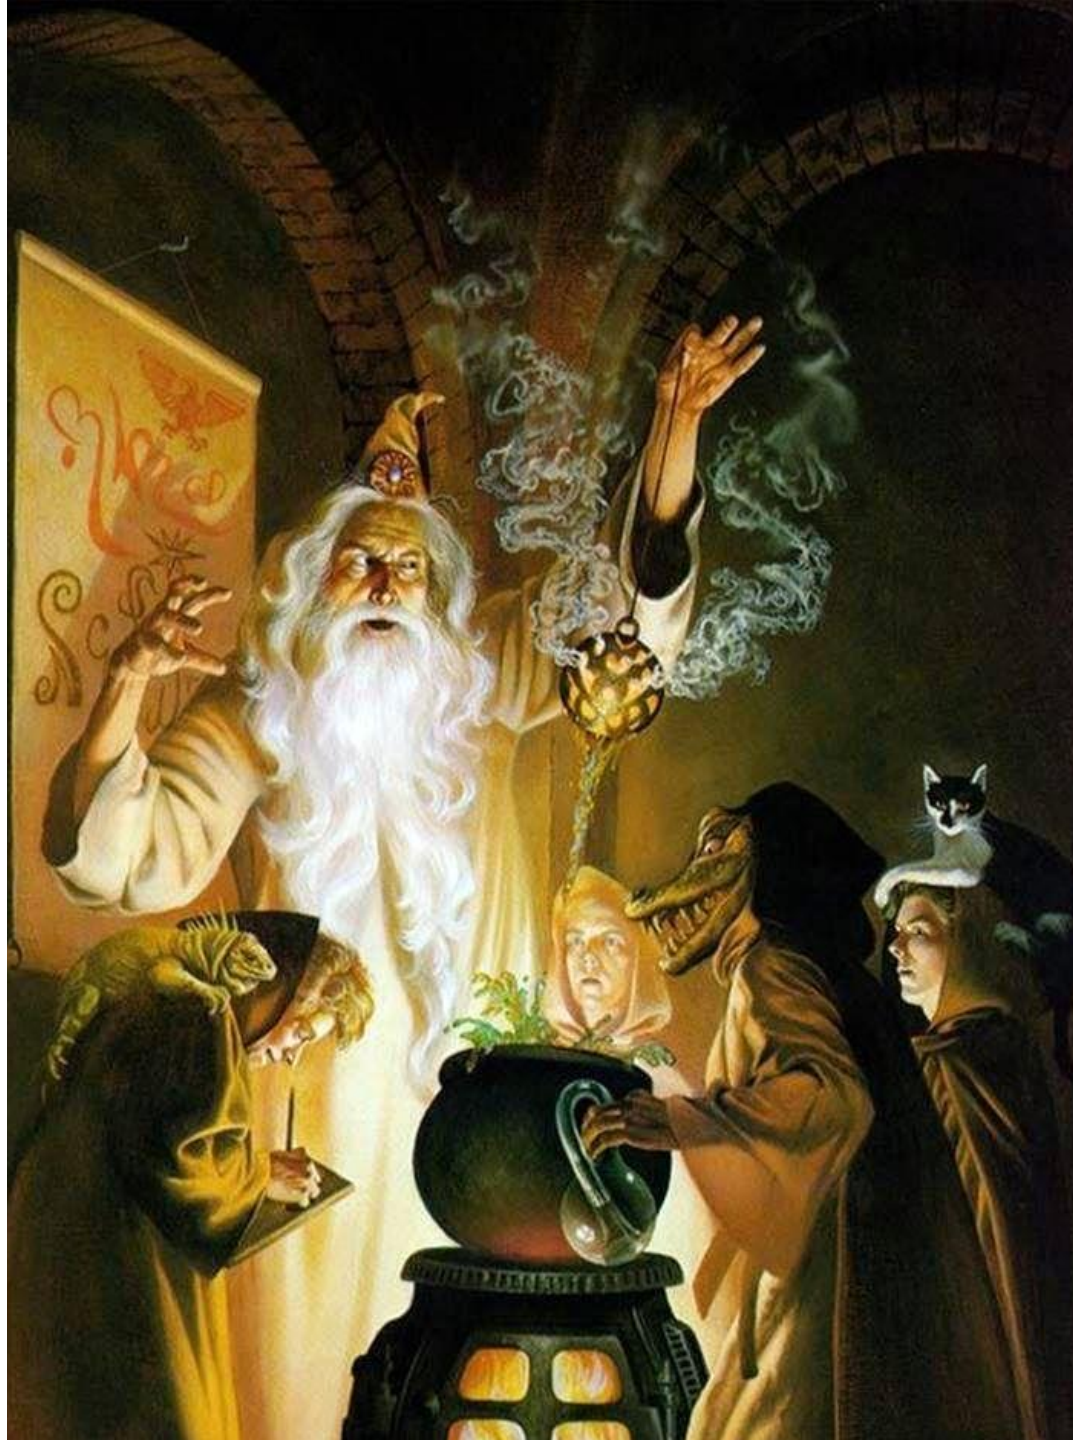

No shortage of product spells (sample list) …

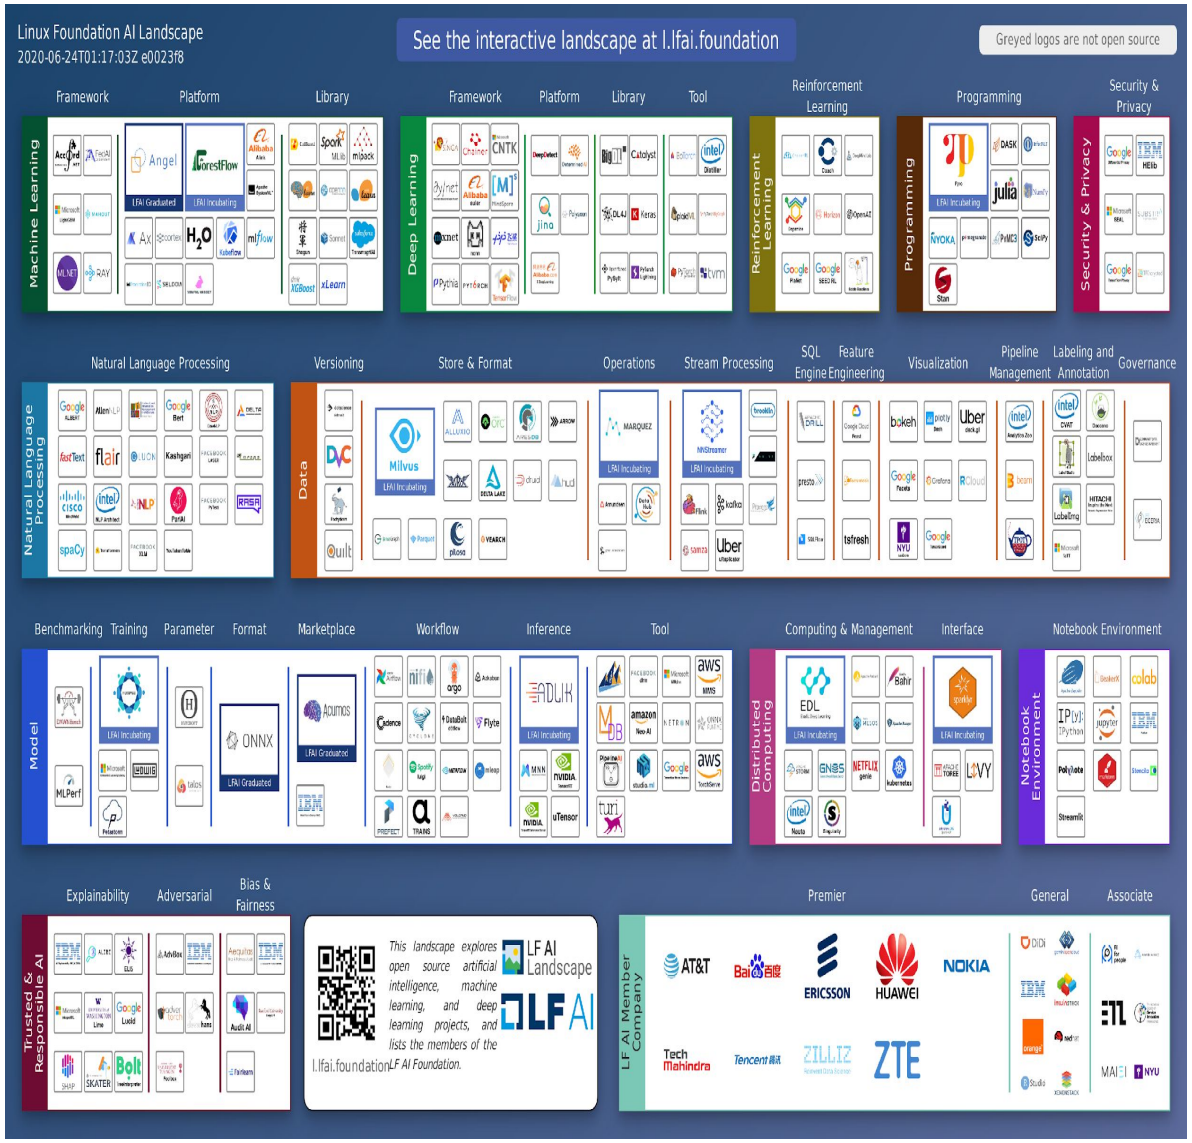

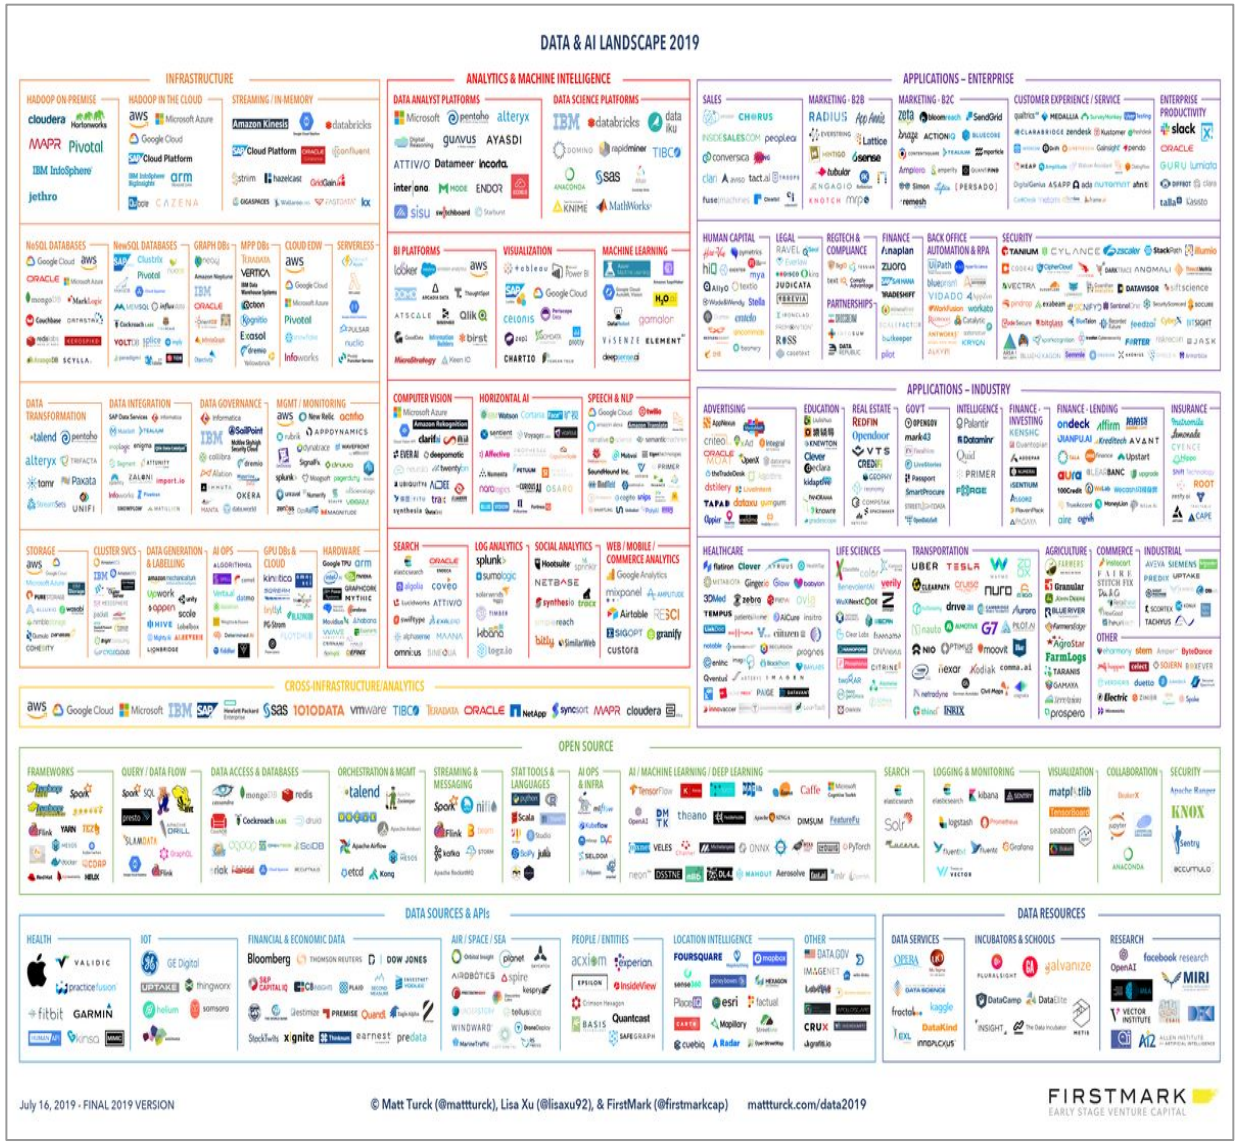

… with many potions

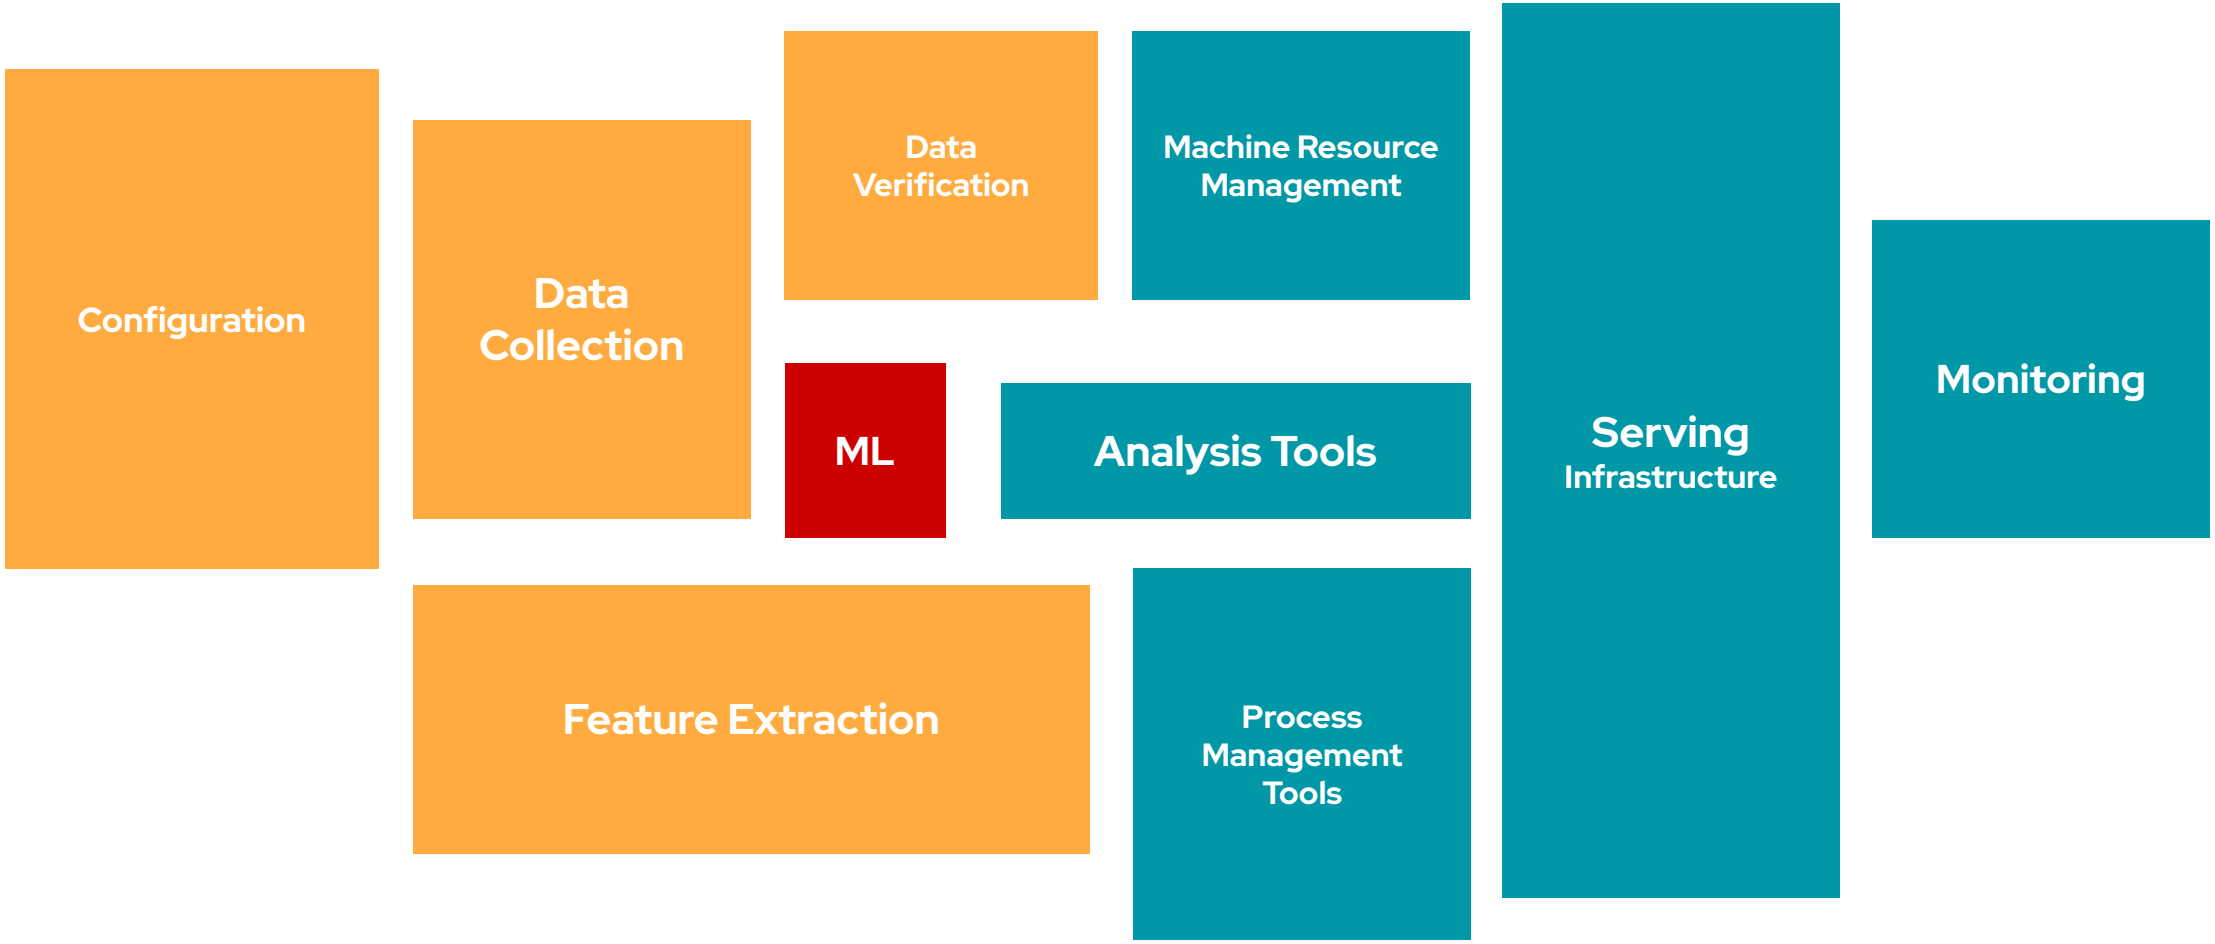

Provisioning; data collection, verification and analysis; data auditing, model monitoring, explainability, machine resource management

https://papers.nips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf

… and silos of sorcery

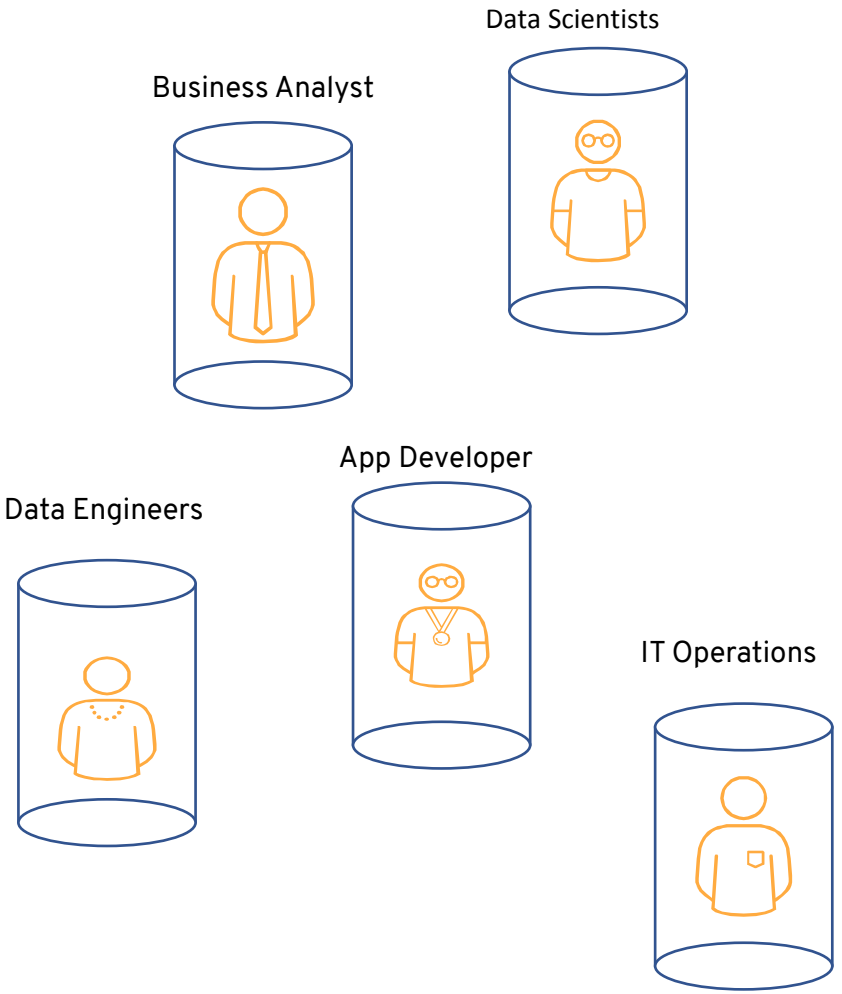

- Limited collaboration & model sharing
- Infrastructure provisioning
- Scale-out issues; resource blowouts
- Manual, error-prone handoffs
- Governance, compliance, security
- Lack of reuse, inconsistencies
- Complexities - training, testing, retraining ML models for prediction accuracy

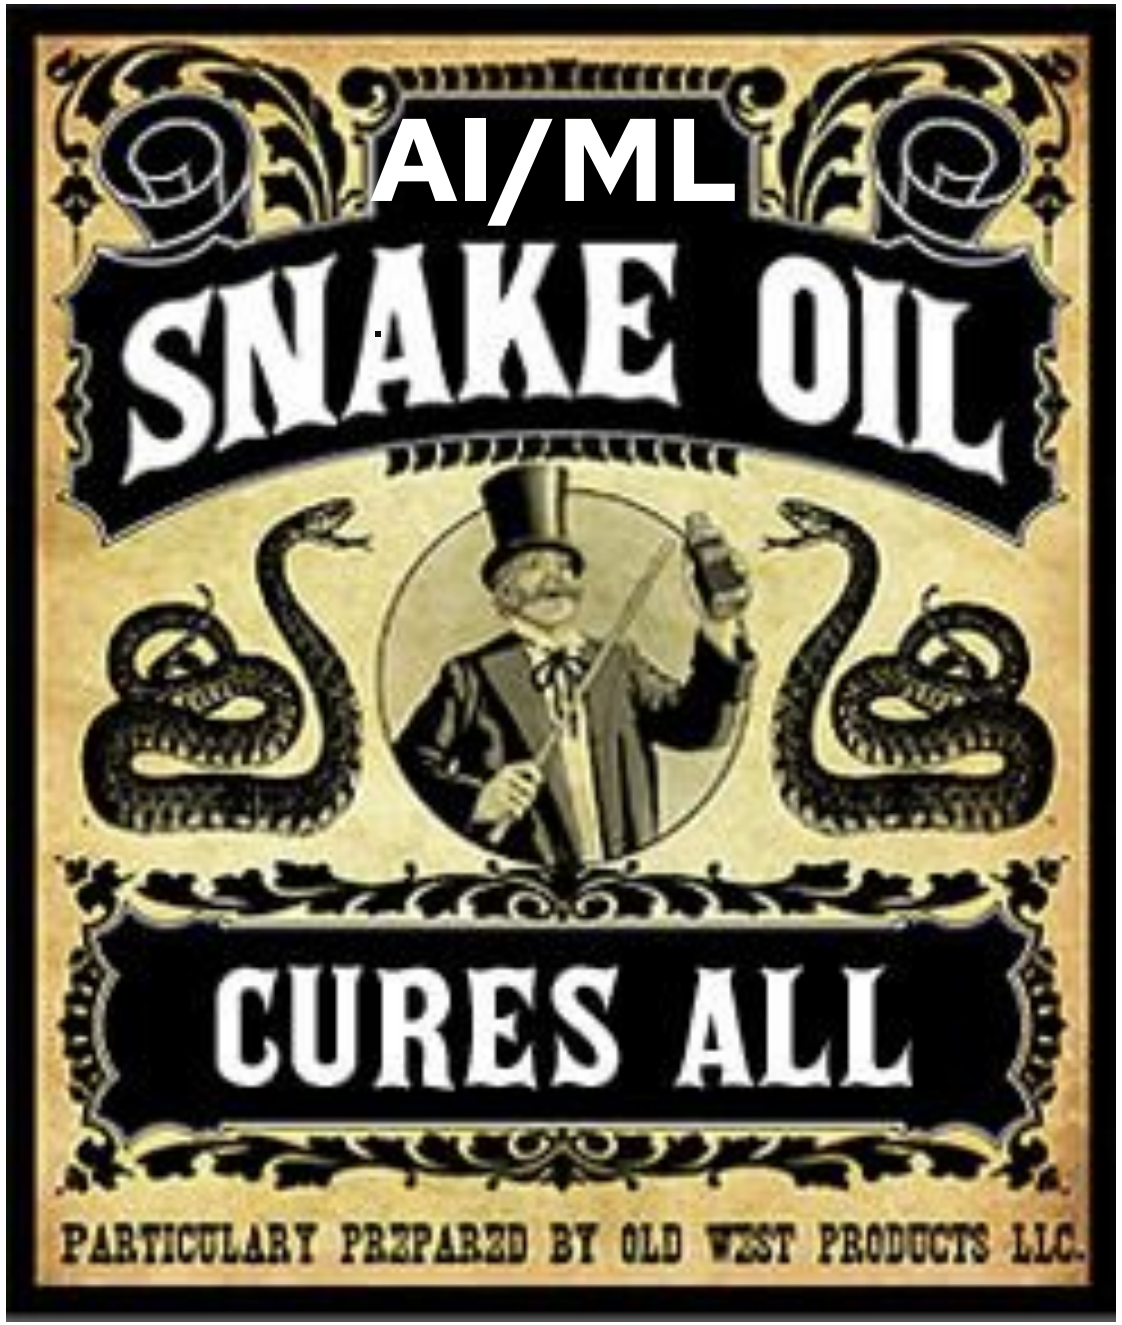
AI/ML
SNAKE OIL
CURES ALL
PARTICULARY PRZPARZD BY OLD WZST PRODUCTS LLC.


WINTER IS HERE

Dragon Taming: Essential Criteria

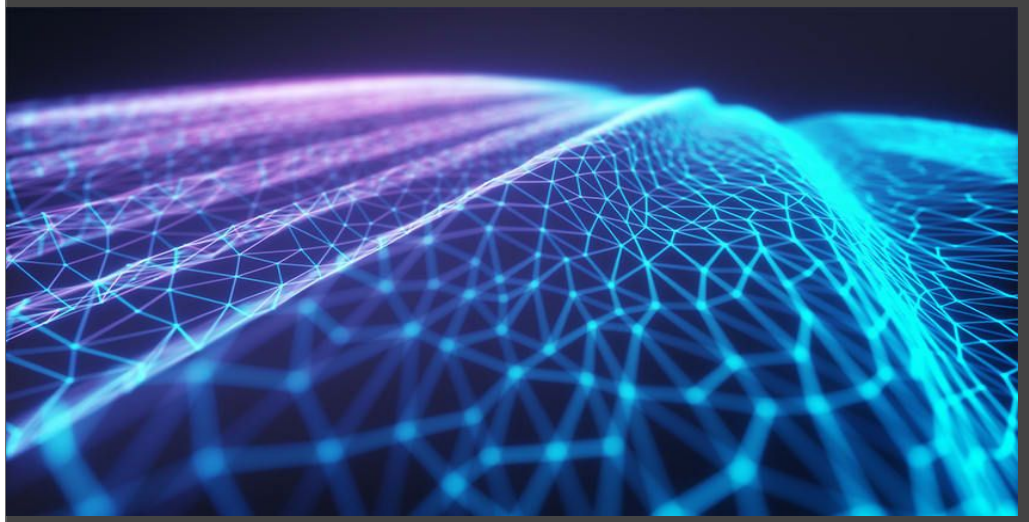

- Scalable, production-ready platform
- Automated, workflow-driven
- Built for collaboration and reuse

Scalable, production-ready platform

Containers and Kubernetes for Data Science Workloads

Agility

Respond quickly with automated compute resource management.

Portability

Develop and deploy AI/ML models consistently across data center, edge, and public clouds.

Flexibility

Provision AI/ML environments as and when you need them.

Scalability

Autoscaling and high availability of the AI/ML solution stack.

Scalable, production-ready platform

Set goals

Gather and prepare data

Develop ML model

Deploy ML models in app dev process

Implement Apps & inference

ML models monitoring & management

ML Software Tools (e.g. TensorFlow, Jupyter Notebooks, Python, etc.)

ML data services - databases (SQL, NoSQL, etc.), data lake, etc.

Kubernetes platform with self service capabilities

Compute acceleration (GPU, FPGA, TPU)

Infrastructure

Physical

Virtual

Private

Public

Hybrid

Edge

Sepsis detection with predictive analytics

"Our existing data infrastructure was designed for large-scale BI and reporting....We need to gather, analyze, and share real-time data from all our facilities so that life-saving action can be taken quickly"

"To support our care providers with modern tools like ML and cloud computing, we needed a platform that is flexible, scalable, and fast"

VP & Chief Data Scientist

- Building clinician "trust"
- New insights from data
- 18 hours early warning
- 8,000 lives saved

Healthcare

Using data insights to save lives

Scaleable, production-ready platform

Supporting key stakeholders

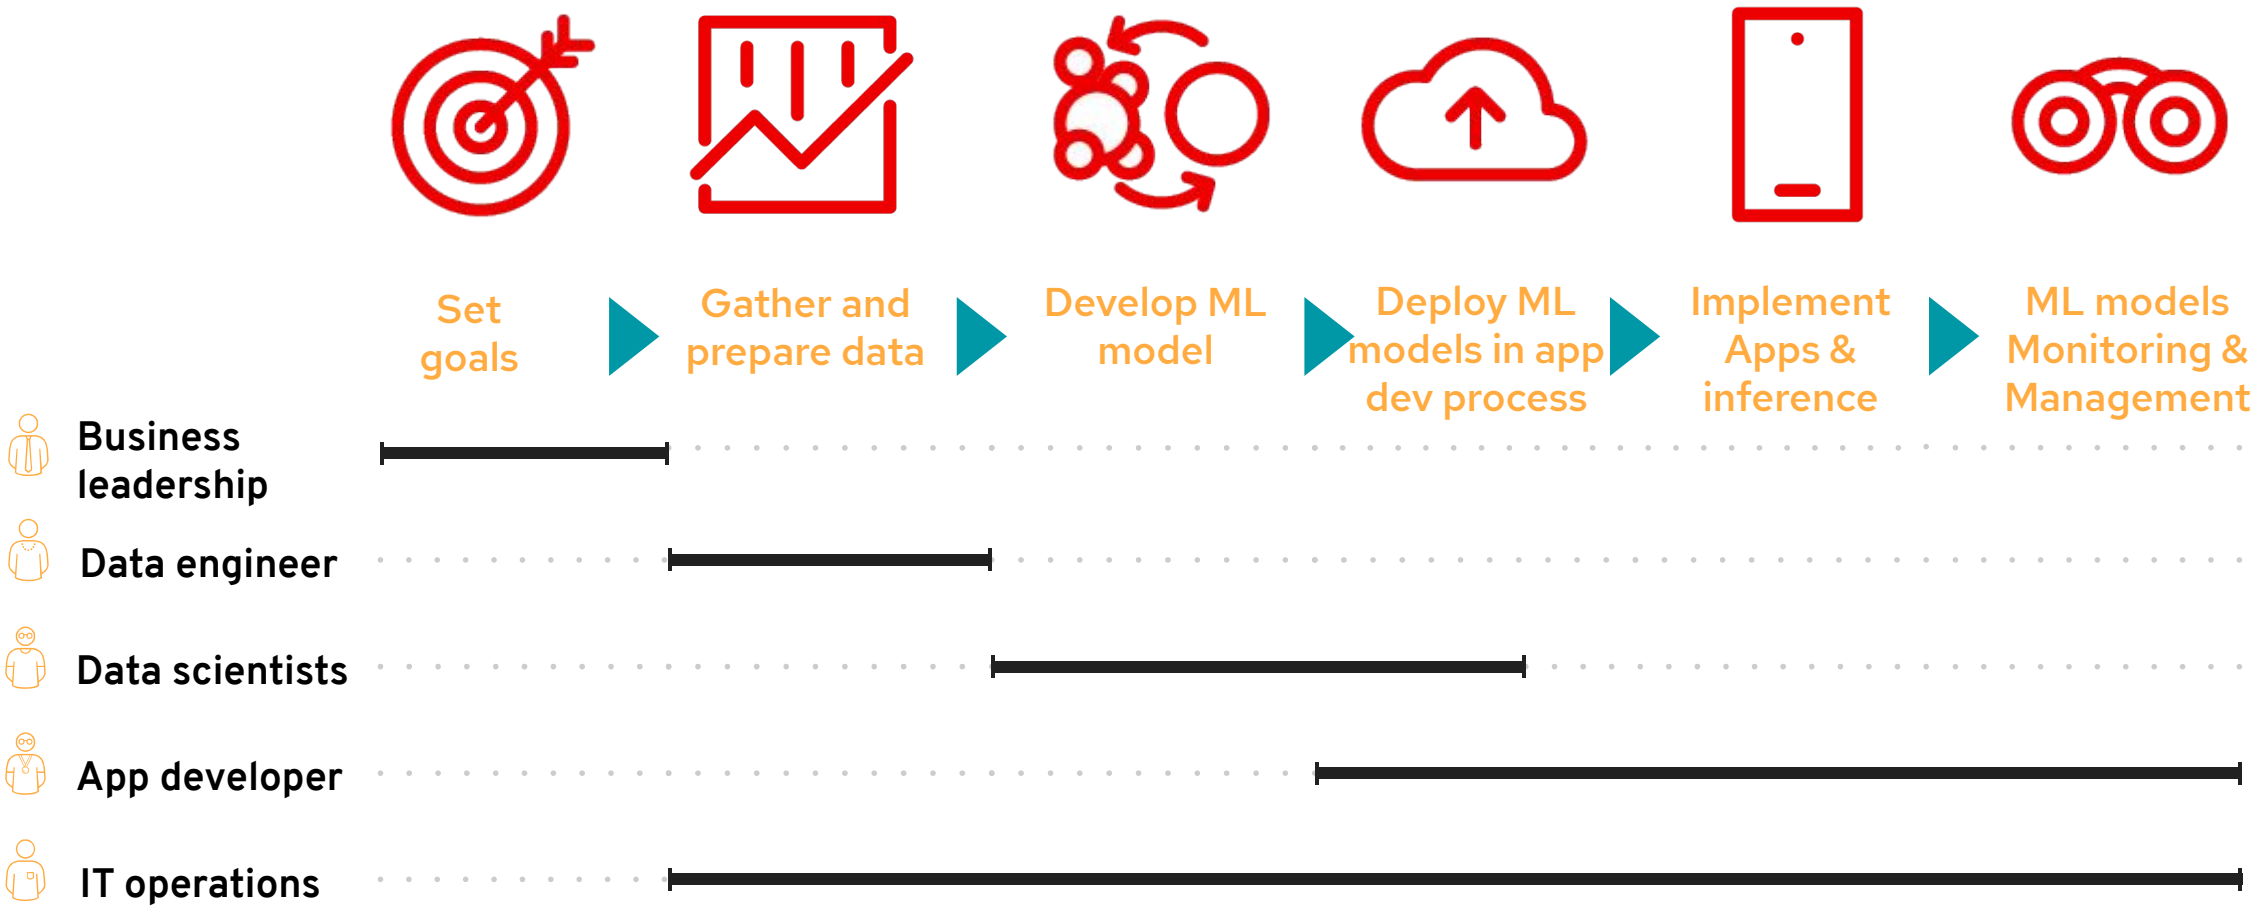

Automated, workflow-driven

Open Data Science Platform Solution

Architected to leverage DevOps and continuous delivery

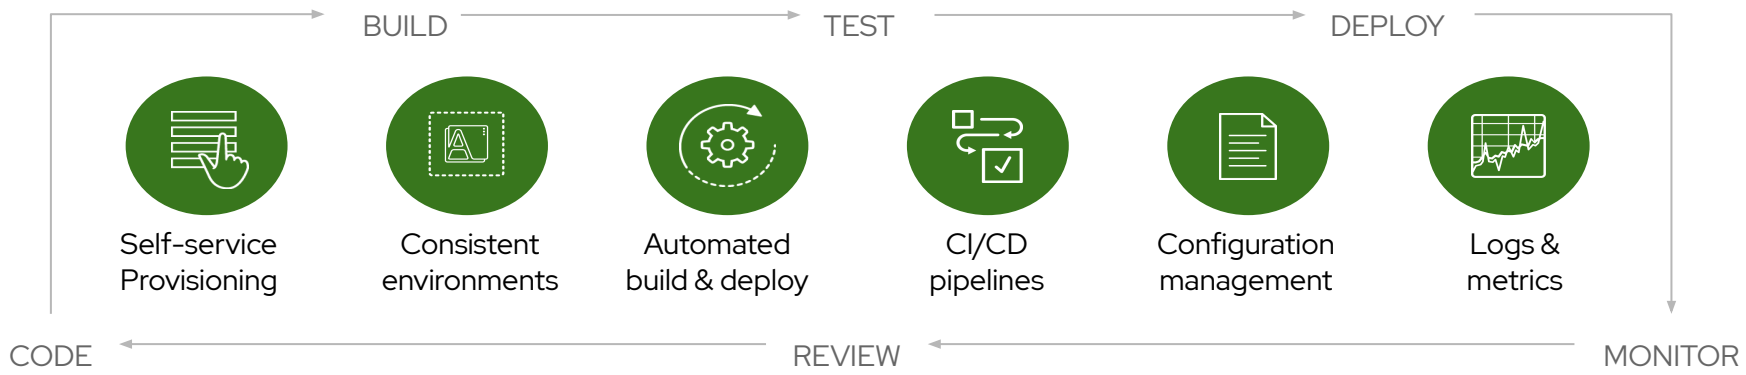

Kubernetes Platform

Automated, workflow-driven

End-to-end workflow - fraud example

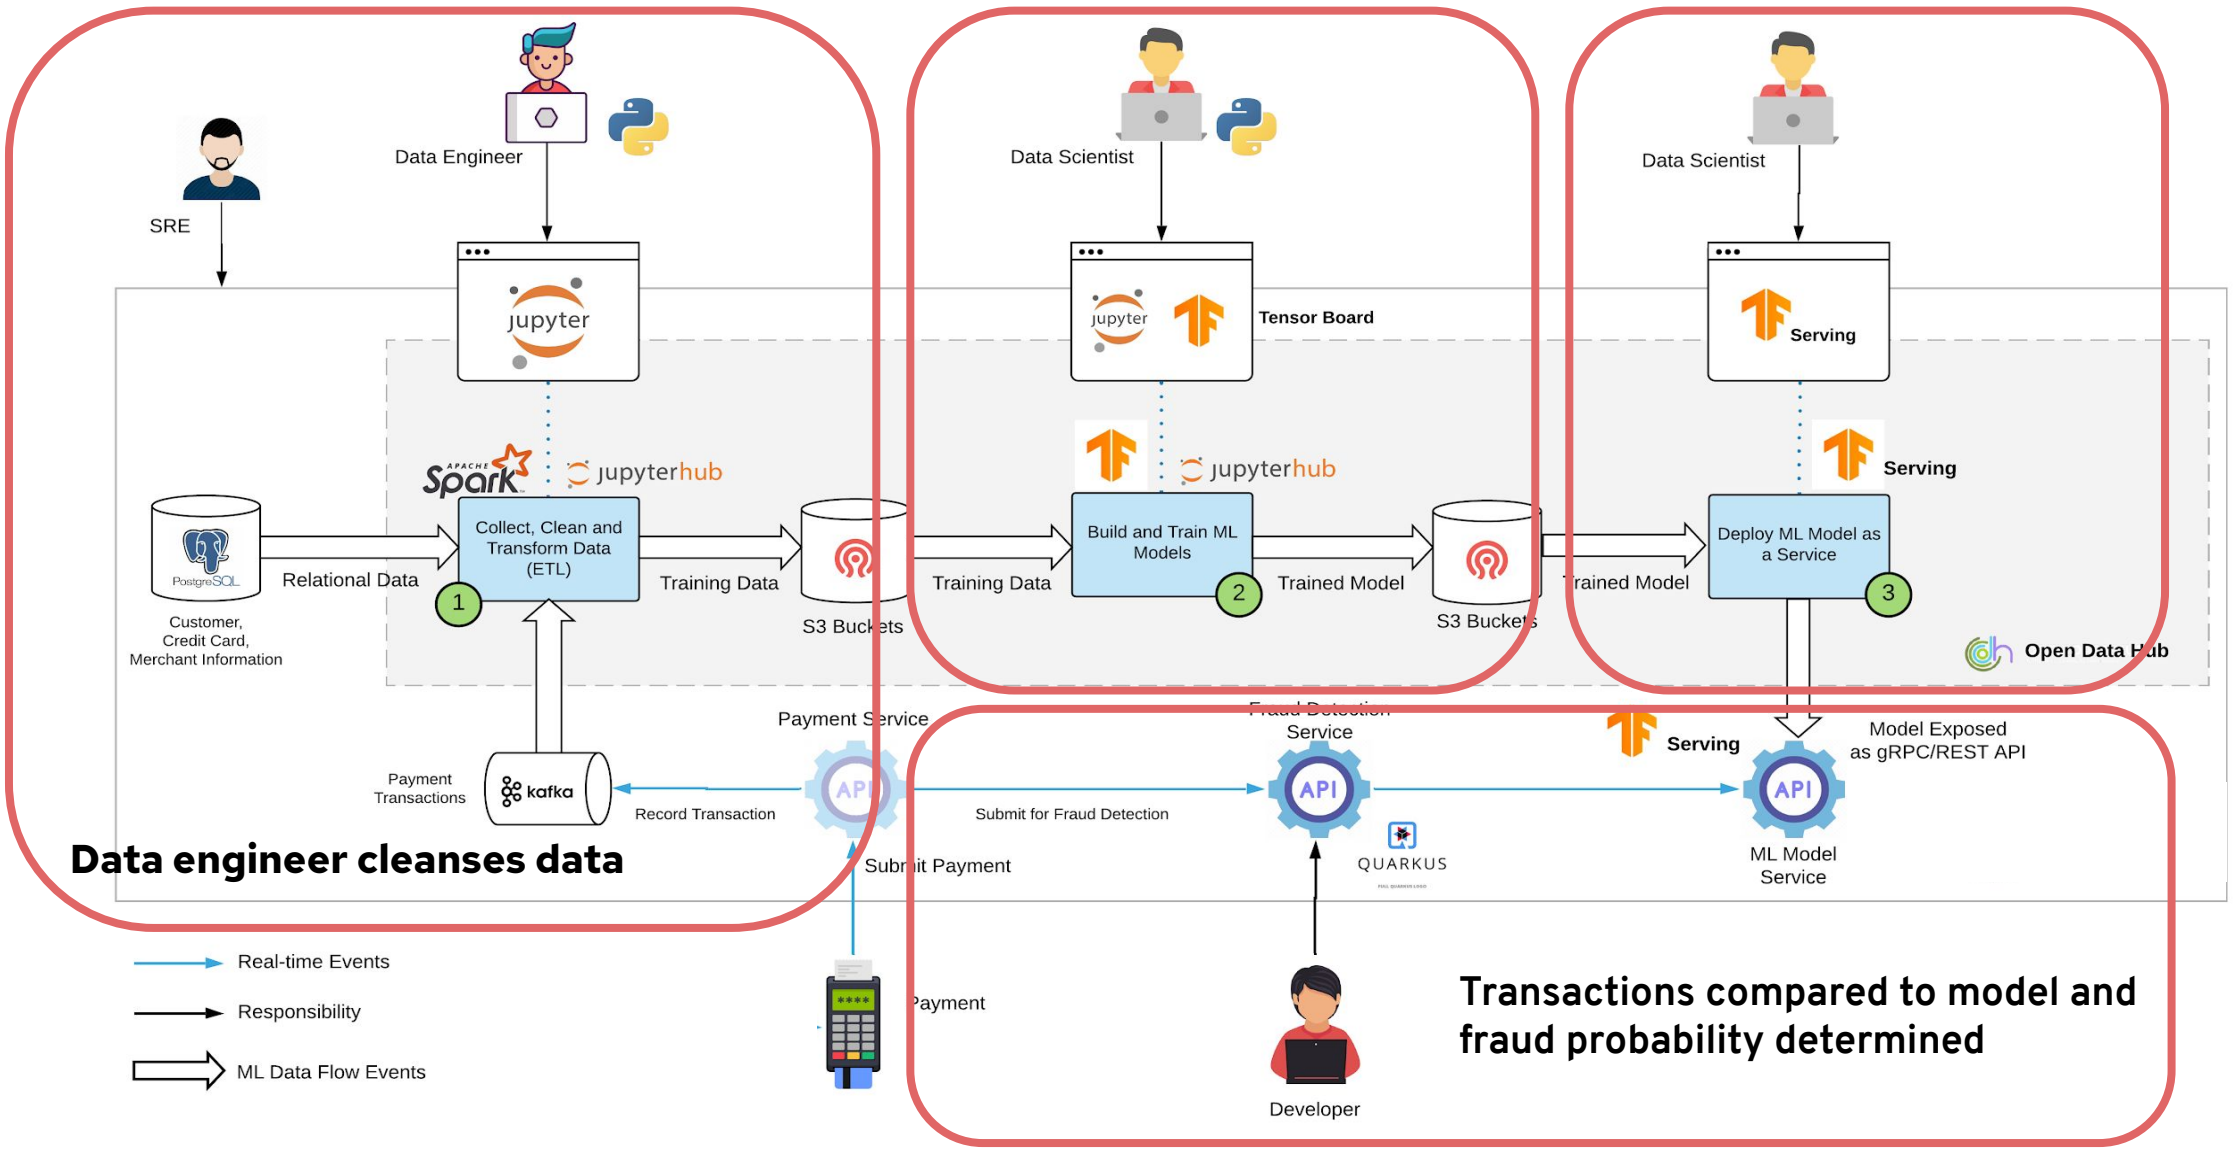

Automated, workflow-driven

Example Data Science Environment Before and After Platform Approach

Local PC Setup

Local admin access required

Access to latest source code

Jupyter Notebook Python 3.x (loaded onto PC – or setup on server)

pip install required Anaconda libraries

SQL Server

OS??

Complex setup requires data scientists to install software on laptops. No version control, high overhead.

Containerised Environment

Jupyter Notebook Python 3.x (image)

Libraries

- Numpy
- Pandas
- Matplotlib
- IPyWidgets
- SciPy

SQLite

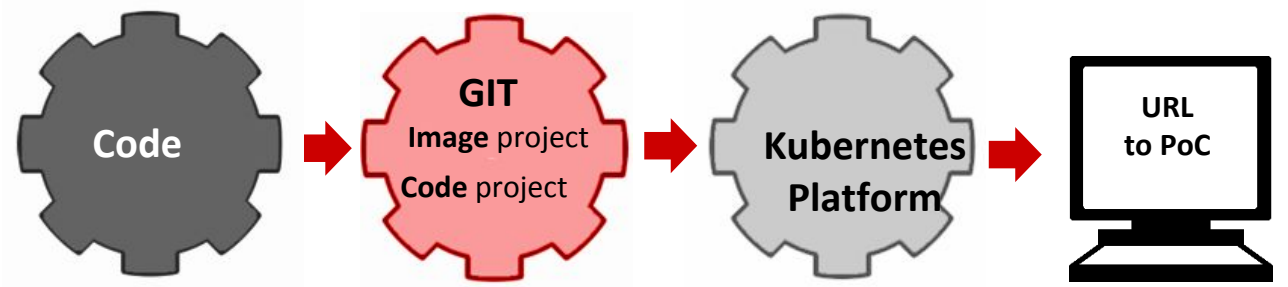

Reproducible Data Science environment that users interact with via chrome, no laptop or server configuration required.

Built for collaboration and reuse

"As a data scientist, I want a ***self-service cloud-like*** *experience for my projects, where I can access a rich set of modelling frameworks, data, and computational resources, where I can* ***share and collaborate*** *with colleagues, and deliver my work into production, with speed, agility, and* ***repeatability*** *to drive* ***business value****"*

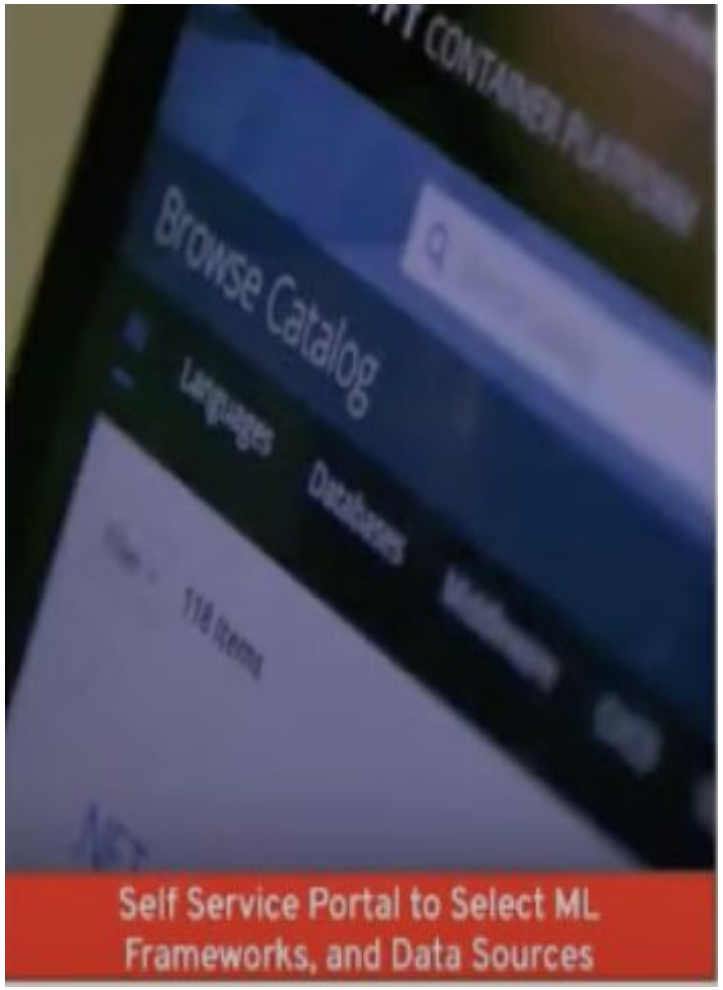

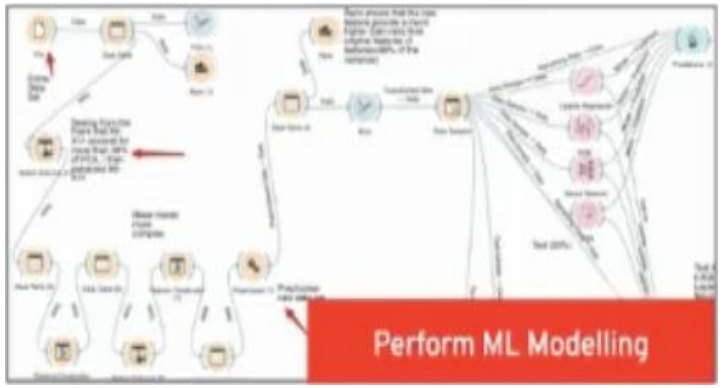

Built for collaboration and reuse

- Meta-project to integrate open source projects
- Collaboration between communities, vendors, user-enterprises, and academics
- Open source best practices (Kubeflow, Kafka, etc)
- Focus on integration, service abstraction, and continuity to reduce complexity
- Ensure reproducibility

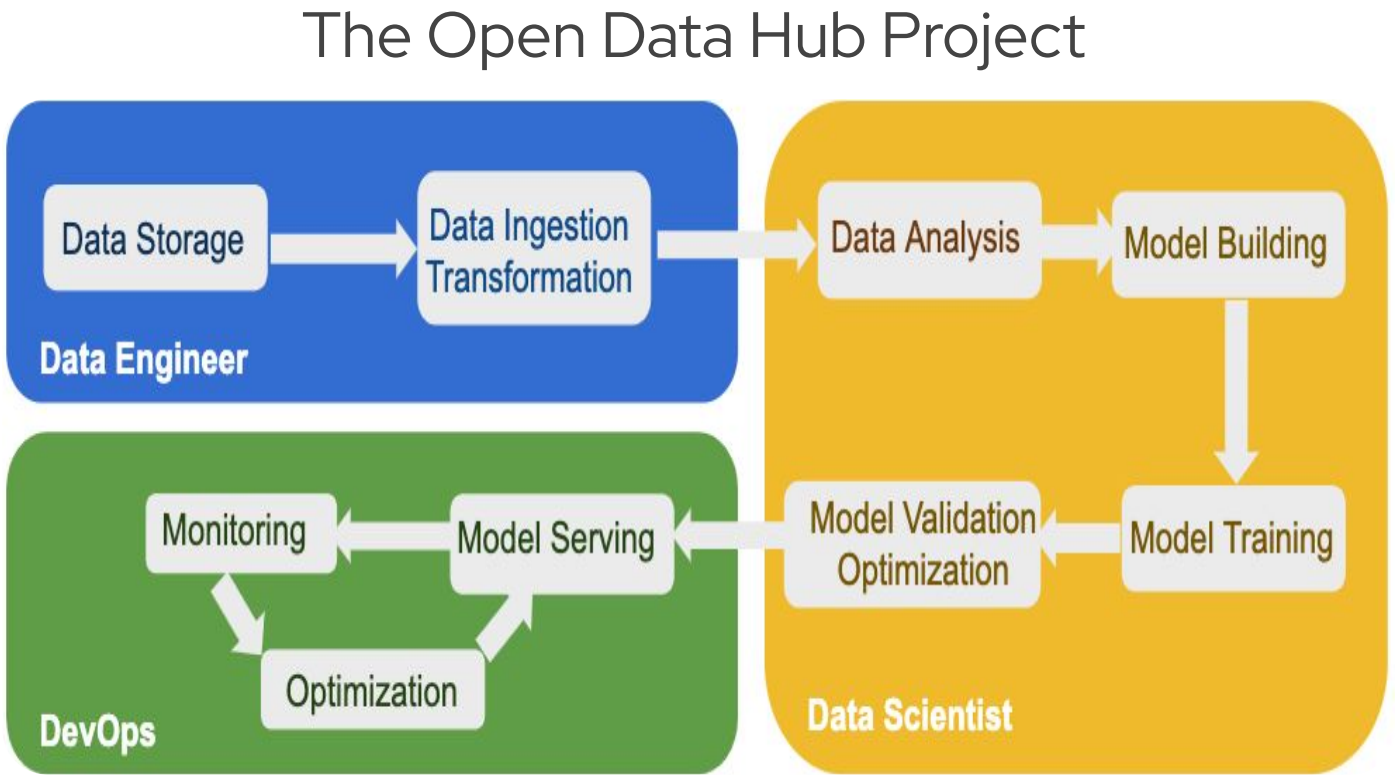

Open Data Hub Architecture:
https://opendatahub.io/docs/architecture.html

Built for collaboration and reuse

Open Data Hub Operator

DATA SCIENTIST

DATA SCIENTIST

DATA SCIENTIST

Prometheus

- Monitoring and alerting toolkit
- Used to diagnose problems

Grafana

- Analytics platform for all metrics
- Query, visualize and alert on metrics

SELDON

- Deploying machine learning models as micro-services
- Full model lifecycle management

APACHE Spark

- Unified analytics engine
- Large-scale data access

jupyterhub

- Multi-user Jupyter
- Used for data science and research

ceph

- Distributed Object Store
- S3 Interface

STRIMZI kafka

- Distributed event streaming
- Pub/Sub Messaging

argo

- Container-native workflow engine
- Declaratively deploy ML pipelines and models

Noteworthy mentions

Dedicated to making deployments of machine learning (ML) workflows on Kubernetes simple, portable and scalable.

Open source project that provides a framework to create cloud-native CI/CD pipelines quickly. As a Kubernetes-native framework, Tekton makes it easier to deploy across hybrid cloud environments

In Summary

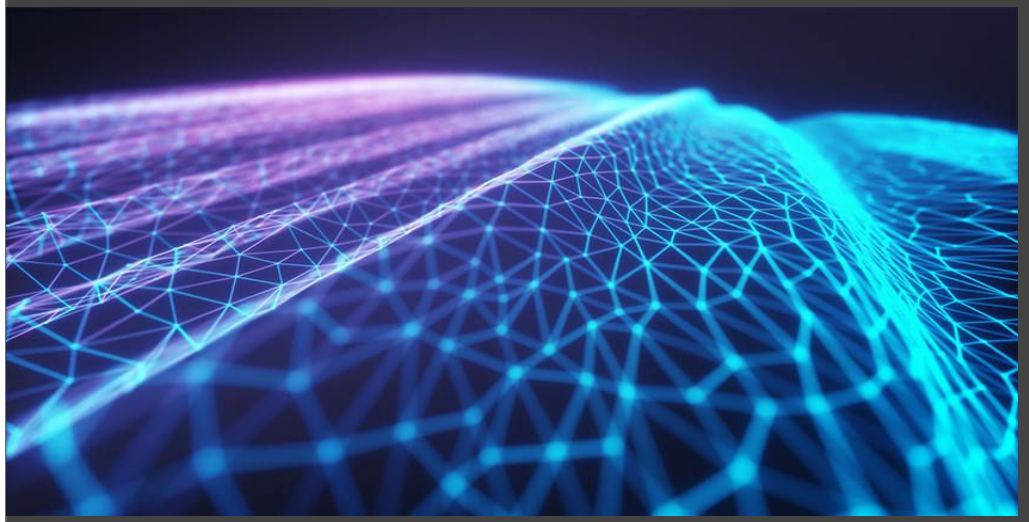

- Data Science - enormous business potential
- Complex challenges and constraints
  - data deluge, science silo's, disparate technology
- Kubernetes and MLOps approach needed: reduce cost, accelerate delivery, increase skills

Taming Data Science Dragons with MLOps and Kubernetes

Scalable, Production-ready platform

Optimally scale data science across multiple environments -- on-premise, cloud, edge

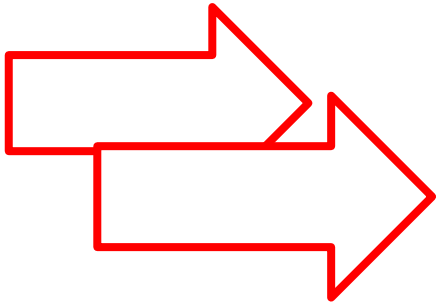

Automated, workflow-driven

Rapidly deploy AI models, integrating multiple technologies and processes

Collaborative, built for reuse

Unify data science teams, developers and operations towards accelerating business outcomes

SKILUP DAYS
by: DevOps Institute
ADVANCING THE HUMANS OF DEVOPS

Pete Waterhouse
Solution Specialist - AppDev
Red Hat
petew@redhat.com

THANK YOU!

Meet me in the Network Chat Lounge for questions

SKILUP DAYS by: DevOps Institute ADVANCING THE HUMANS OF DEVOPS