

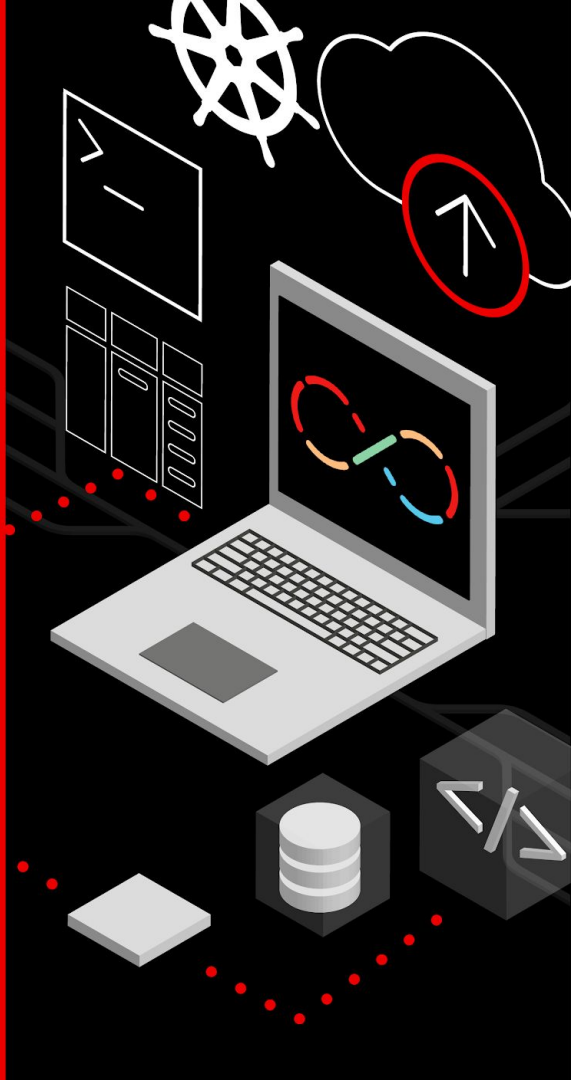
## MLOps with Kubernetes

Unleash the power of the  
ecosystem!

**Natale Vinto**

Developer Advocate Lead

@natalevinto





Natale Vinto @ Red Hat



**Global Developer Advocate Lead**

Previously EMEA OpenShift SSA  
CD Foundation Ambassador 2023

[twitter.com/@natalevinto](https://twitter.com/natalevinto)

[mastodon.uno/@natalevinto](https://mastodon.uno/@natalevinto)

<https://www.linkedin.com/in/natalevinto/>



# Who are you?



Data  
Science

Dev

Ops

<https://www.menti.com/eaaa3a8zab>



---

# AI/ML People

# People + Collaboration



Business  
leadership



Data  
engineer



Data  
scientists



ML  
engineer

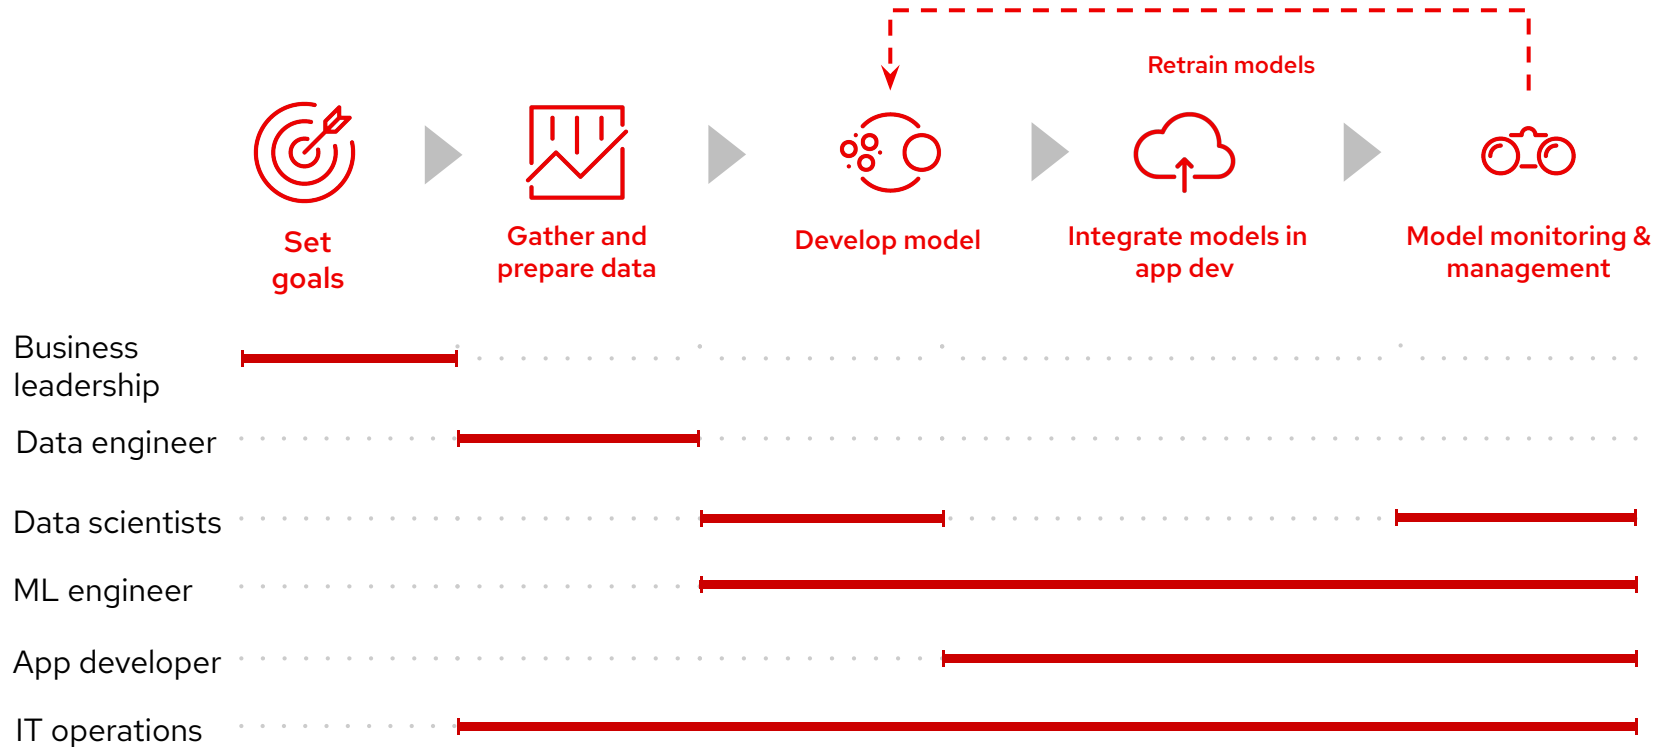


App  
developer



IT  
operations

# Operationalizing AI/ML is not trivial



# The ML lifecycle in technical detail



---

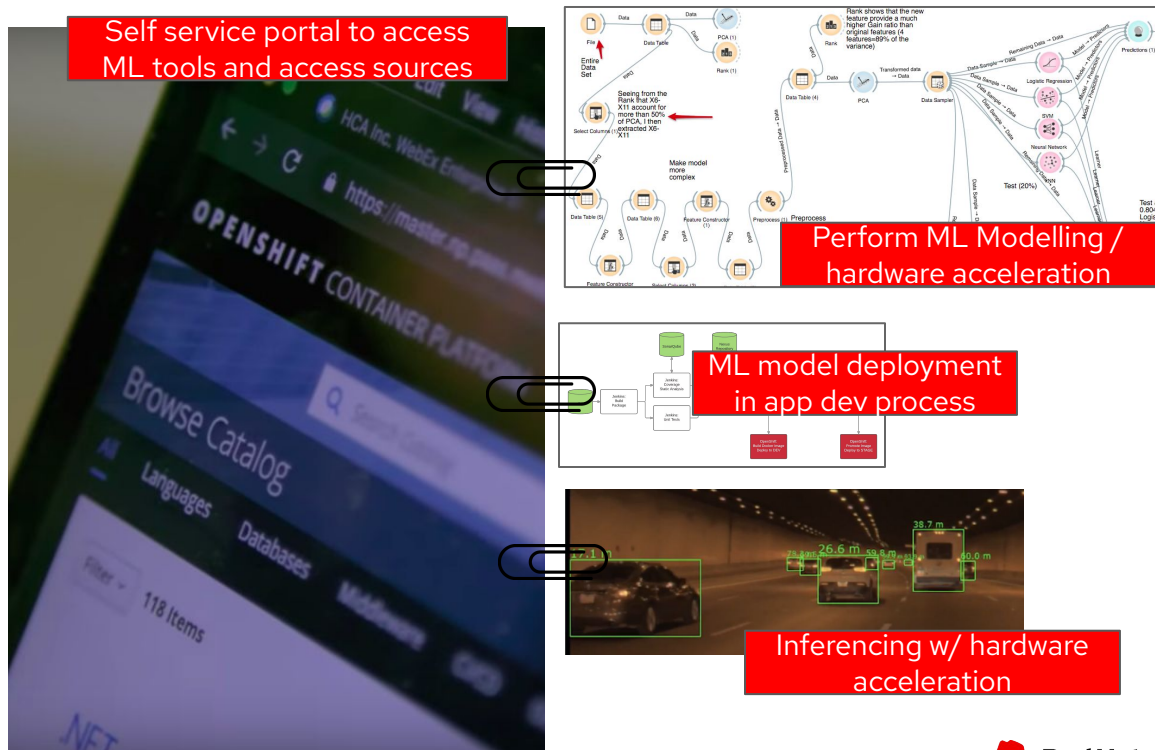
# A day in the life of a Data Scientist



# What does a Data Scientist care about?

As a Data Scientist, I want a **“self-service cloud like”** experience for my Machine Learning projects, where **I can access a rich set of modelling tools**, data, and computational resources, share and collaborate with colleagues, and deliver my work into production with speed, agility and repeatability to drive business value!

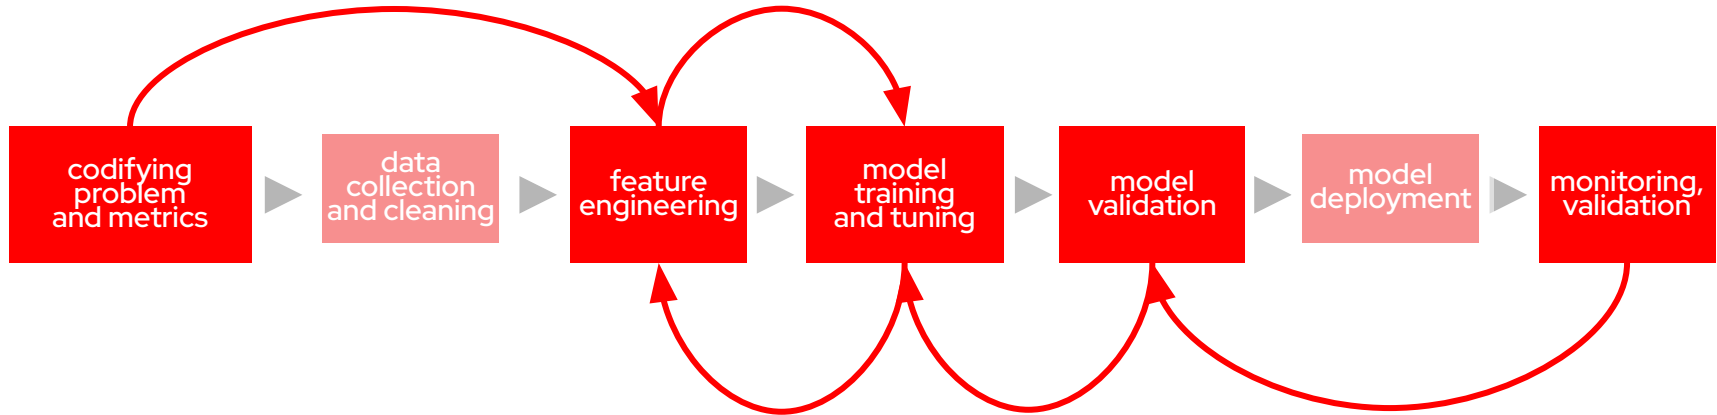
9



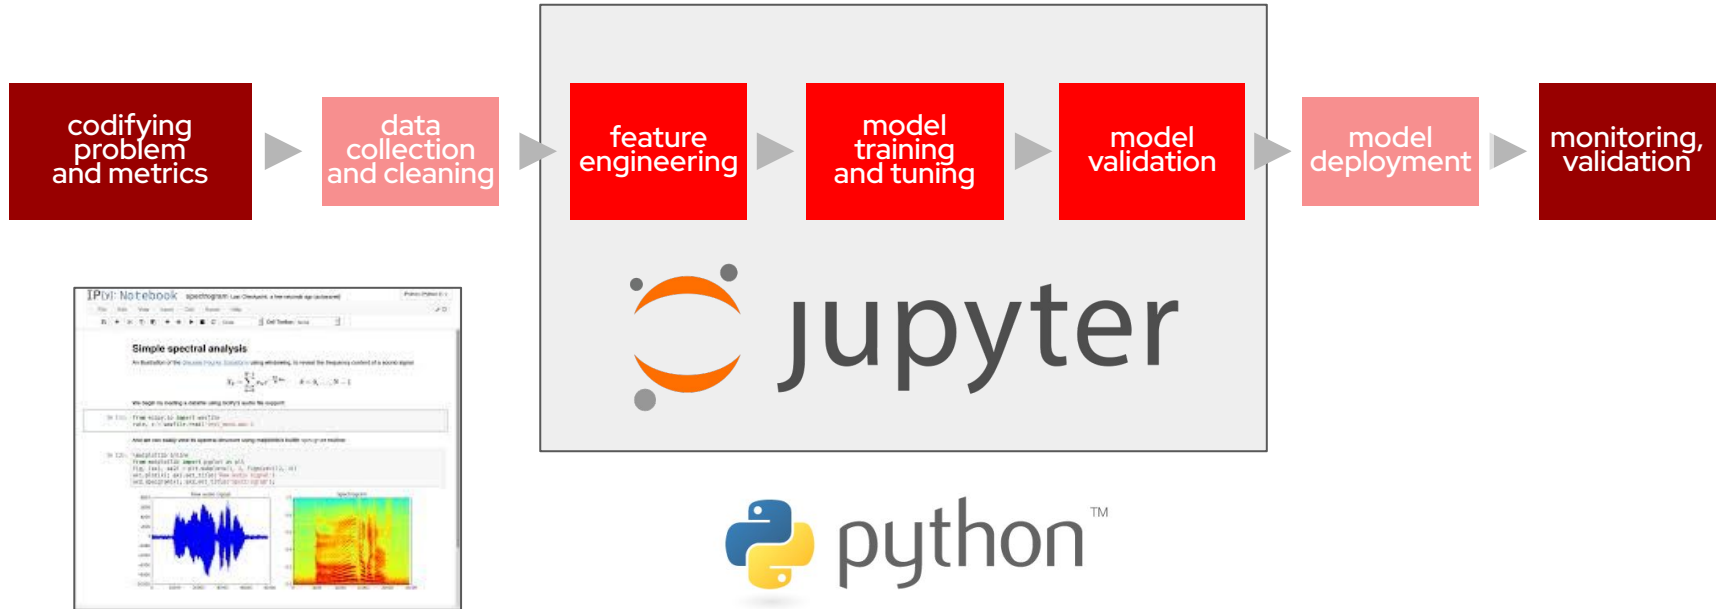
## The data scientist's domain



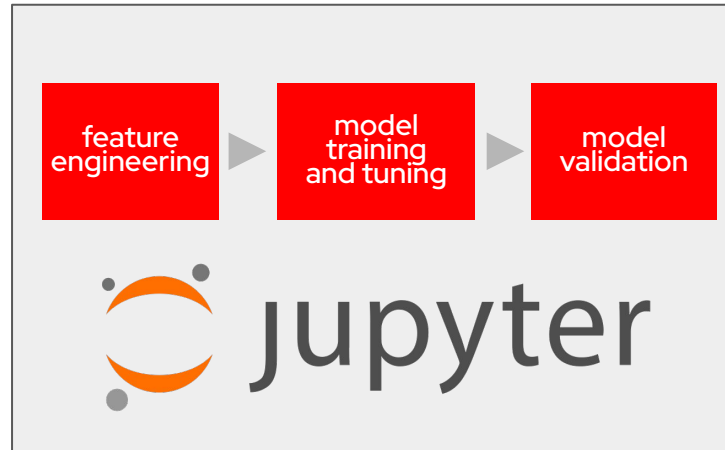
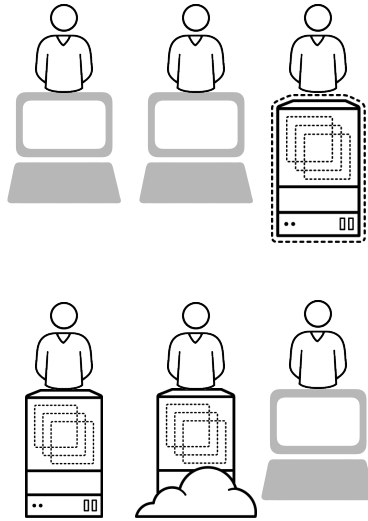
# The process is iterative




# Jupyter(Hub/Notebooks) in data science



# Everyone does the same thing... differently



## Two primary bottlenecks



- Getting started
- Getting into production

---

# Kubernetes



# What is Kubernetes?

An open source orchestration system for managing containerized workloads across a cluster of nodes.



# Container Technology

Easy to scale up

Mature technology


Supporting also Data Science workloads



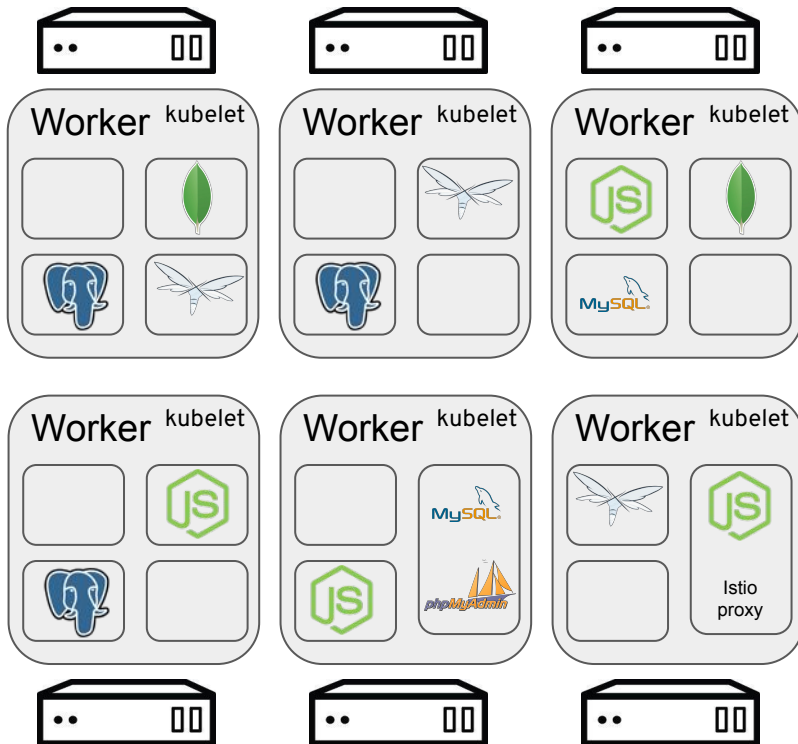
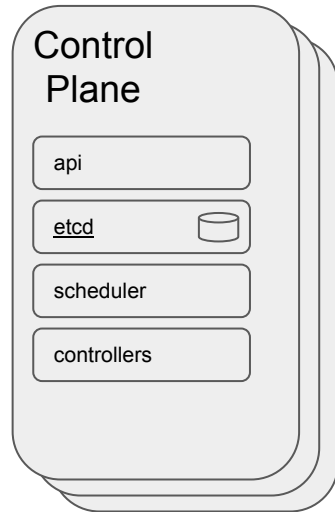
# Kubernetes Cluster - Nodes



Dev



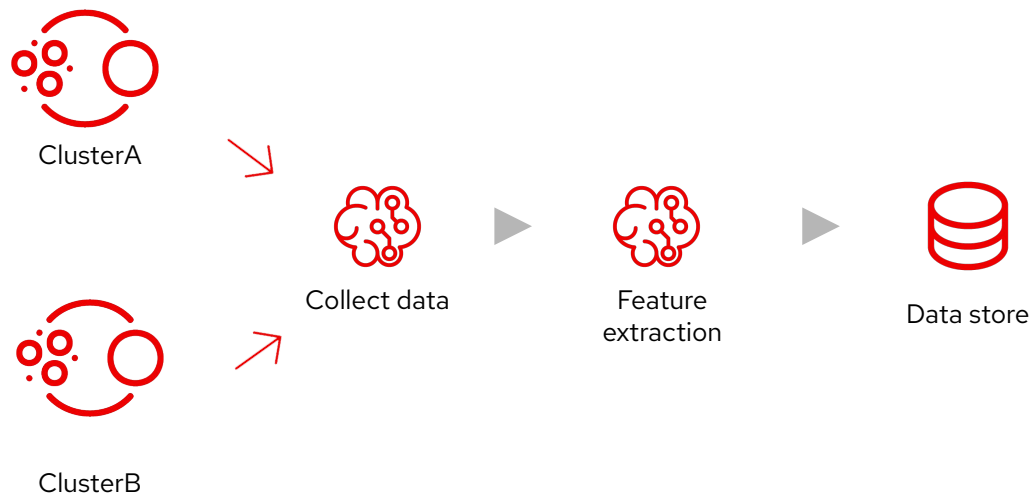
Ops



---

# MLOps: DevOps to ML

# MLOps: Data Pipeline





An open-source project for providing a set of shared and standard components for building Kubernetes-style CI/CD systems

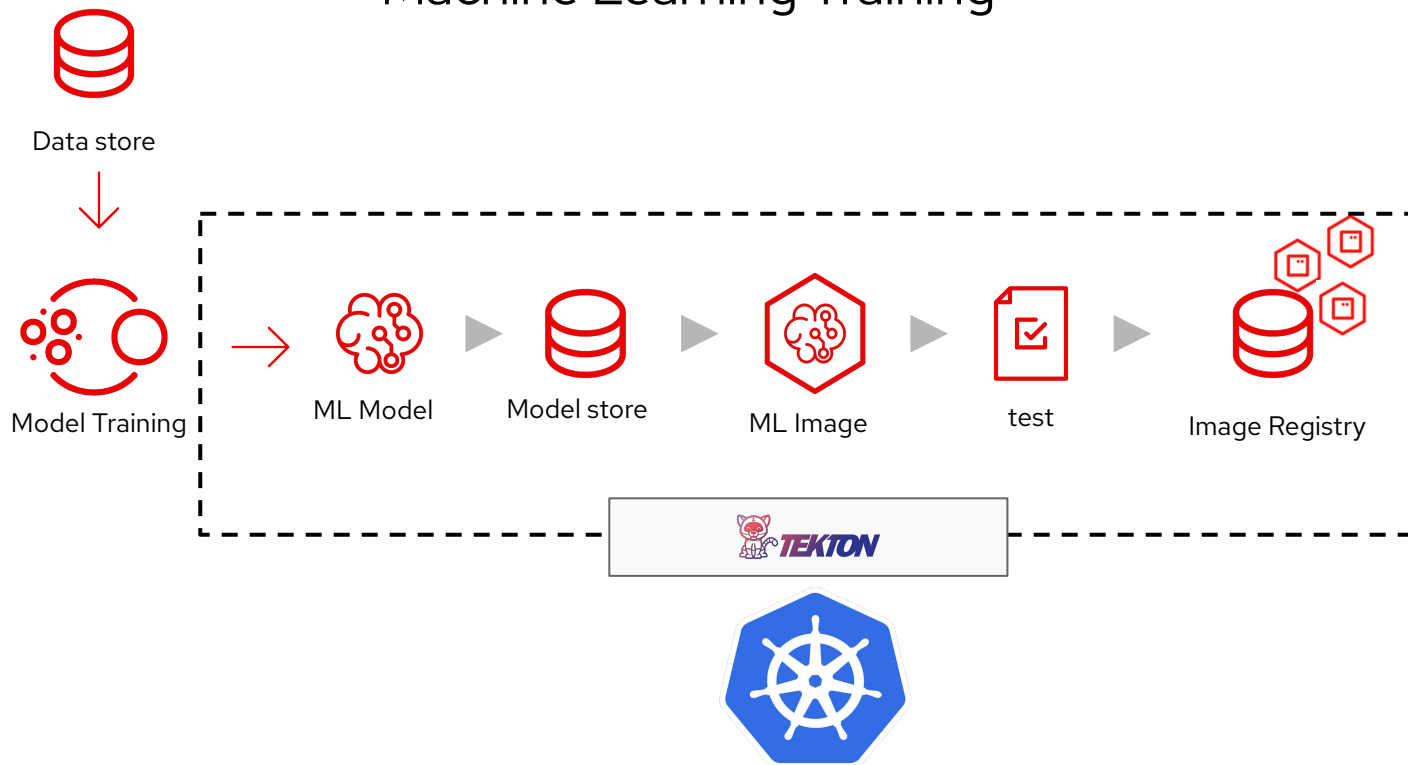


CD.FOUNDATION

Governed by the Continuous Delivery Foundation


Contributions from Google, Red Hat, Cloudbees, IBM, Pivotal and many more

# Machine Learning Training





# MLOps: Release Pipeline



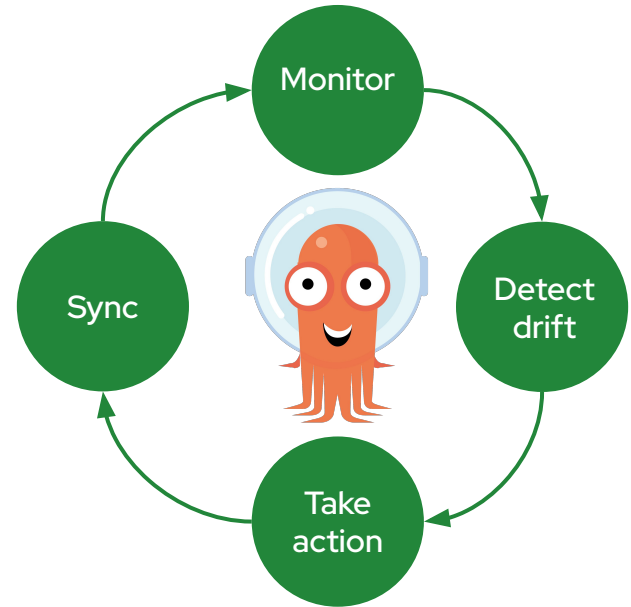
## Argo CD

Argo CD is a declarative, GitOps continuous delivery tool for Kubernetes.

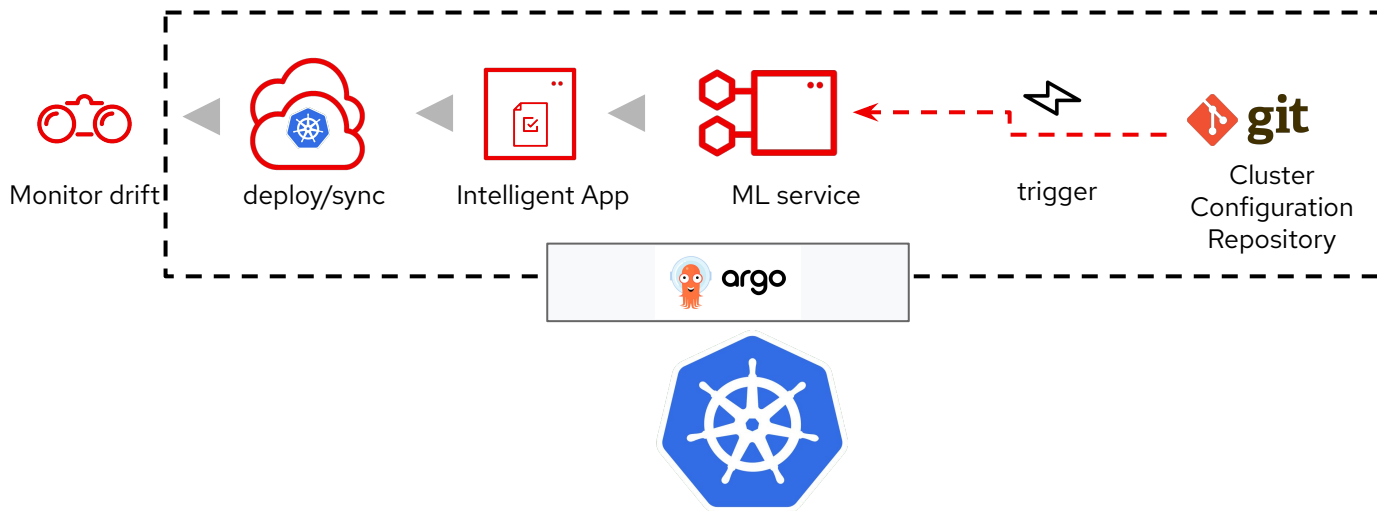
Cluster and application configuration versioned in Git

Automatically syncs configuration from Git to clusters

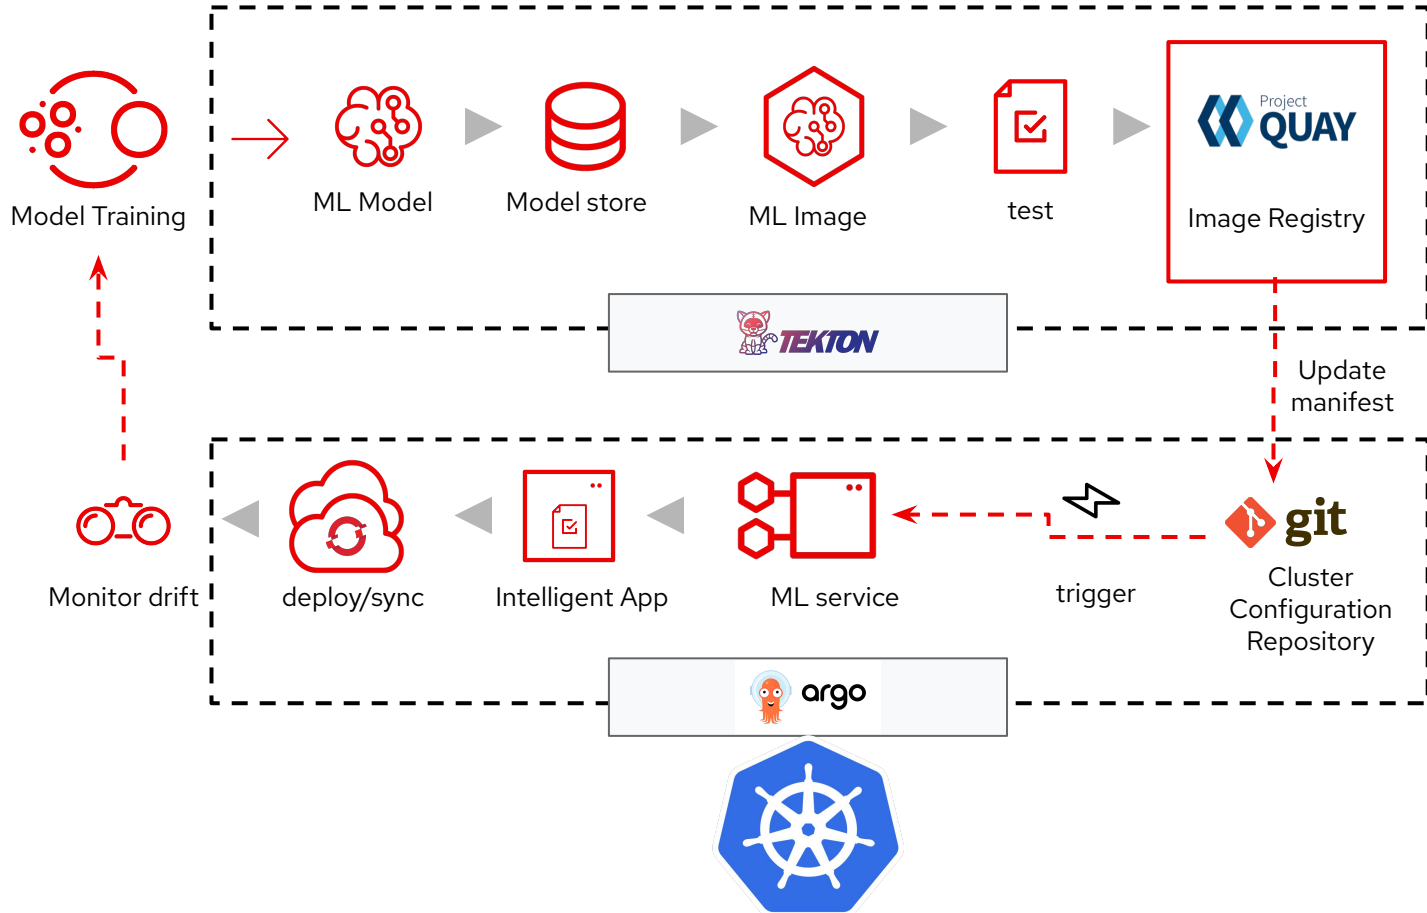
Drift detection, visualization and correction



# MLOps: Release Pipeline



# MLOps example



# Kubernetes and cloud-native supporting ML workloads



(OpenShift Pipelines)

Framework to create CI/CD pipelines to build, test, and deploy across environments



(Red Hat Quay)

Builds, analyzes, and distributes container images



(Red Hat AMQ Streams)

Distributed event streaming platform for data pipelines, analytics, and integration



(OpenShift GitOps)

Argo CD is a declarative, GitOps continuous delivery tool for Kubernetes



Open Data Hub

(Red Hat OpenShift Data Science)

Blueprint for AI-as-a-service on OpenShift Container Platform and Ceph Object Storage



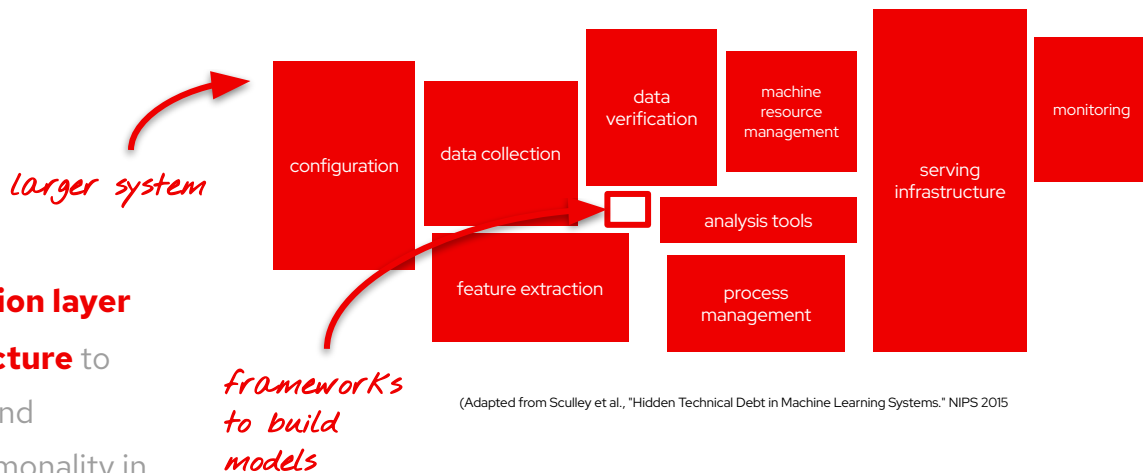
Toolkit and workflow to build reproducible container images from source code

## Complexities of operationalizing models

### "a consistent application

**platform** for the management of existing, modernized, and cloud-native applications that runs on any cloud."

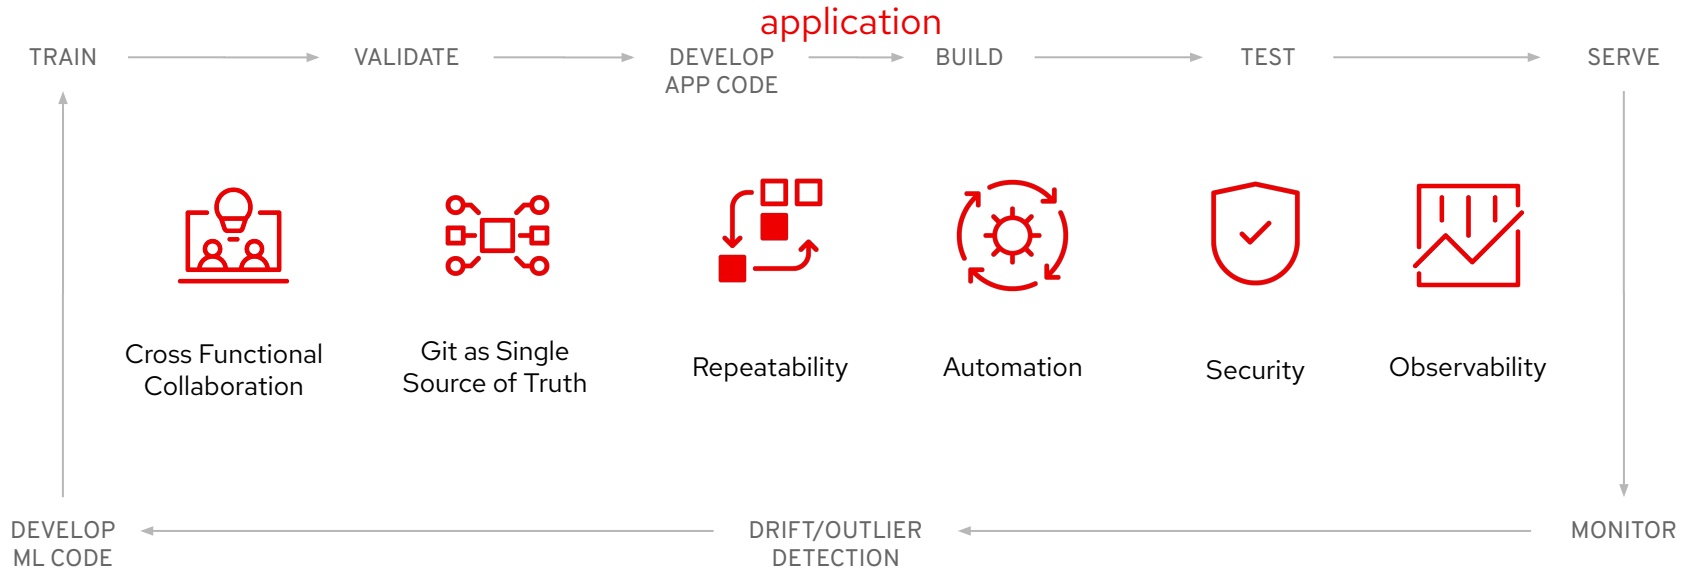
"a common abstraction layer across any infrastructure to give both developers and operations teams commonality in how applications are packaged, deployed, and managed."



(Adapted from Sculley et al., "Hidden Technical Debt in Machine Learning Systems." NIPS 2015)

# Enter MLOps

MLOps incorporates DevOps and GitOps to improve the lifecycle management of the ML



29

[1] Reference and things to read: <https://cloud.redhat.com/blog/enterprise-mlops-reference-design>



---

# OpenDataHub

# Open Data Hub

An open source MLOps suite



- Multi-tenant data science platform
- Self-service workspaces



- Preinstalled machine learning libraries
- Custom stack can be integrated



**Elyra**

- AI pipeline editor
- Define workflows through Jupyter



**Kubeflow Pipelines**

- Machine learning workflow orchestration
- Experiment tracking



**Kserve ModelMesh**

- Deploying machine learning models as micro-services
- Pre-built inference servers



Prometheus

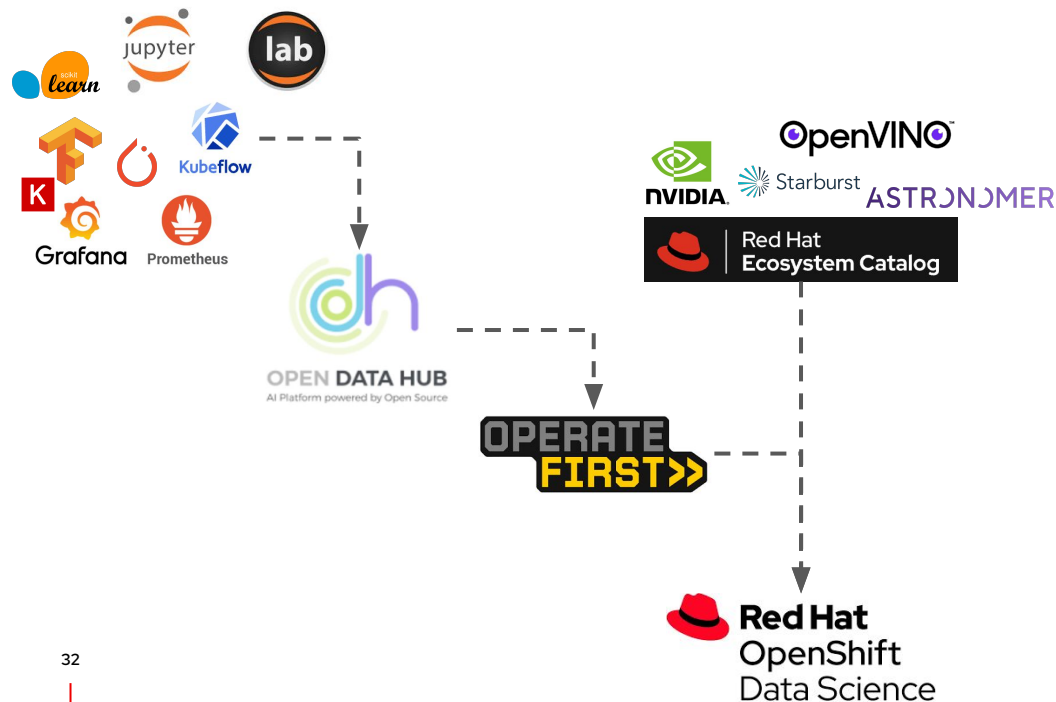


Grafana

- Monitoring stack
- Native integration with OpenShift

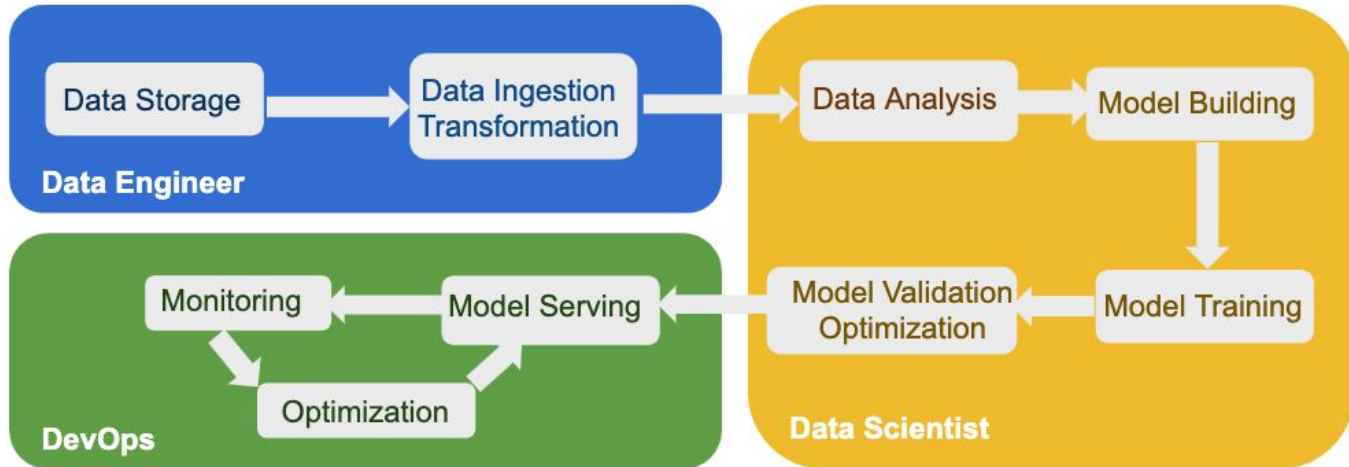
# Red Hat's investments in Data Science

Community projects curated into supported Offerings

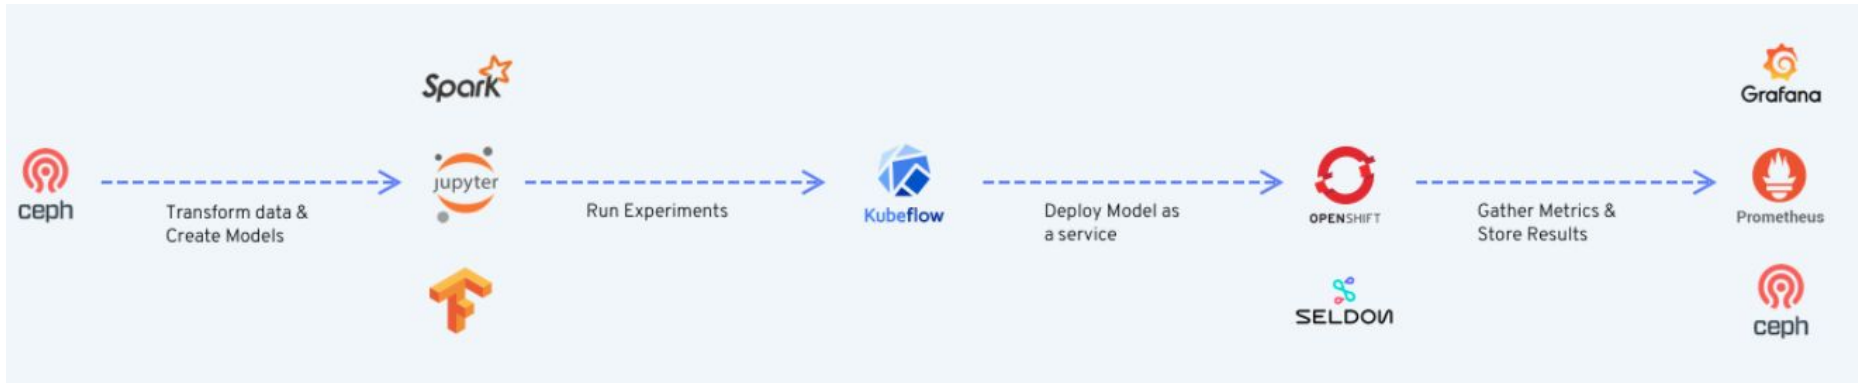


- **Kubeflow** is the open source ML toolkit for Kubernetes "K8s" and community project.
- **Open Data Hub** "ODH" is the open source ML toolkit for Red Hat OpenShift "OCP"; It is an OCP Operator; It is the community project.
- **Operate-First** is an all-Open Source community cloud to learn and make things, like ODH.
- **Ecosystem Catalog** lists certified to run on OCP (like Starburst, Kafka, etc.) as well as integration solutions.
- **Red Hat OpenShift Data Science** "RHODS" is a managed cloud or self-managed service offering based on the open source Open Data Hub project.
- **Red Hat Services** provides a paved and time-bound road to operationalize data science Platform or MLOps automation.

- **Collaborate** on a Data & AI platform for the Hybrid Cloud - [opendatahub.io](https://opendatahub.io)
- **Open Source** community Meta-Project to integrate open source projects into an **end-to-end AI/ML platform** on OpenShift
- **Install** today from the Openshift [OperatorHub.io](https://operatorhub.io) Community Operators



# OpenDataHub.io



# OpenDataHub: Kubeflow

- Kubeflow - A new open source project dedicated to making deployments of machine learning (ML) workflows on Kubernetes **simple, portable and scalable** - [kubeflow.org](https://kubeflow.org)
- It is now integrated into Open Data Hub and runs on **OpenShift**.
- Kubeflow brings multiple new AI/ML capabilities and features
  - For **model training** Tensorflow and Pytorch
  - For **model serving** Seldon and KFServing
  - For **pipelines** Kubeflow Pipelines based on Argo CD



Kubeflow

## Documentation

- About
- Getting Started
- Components of Kubeflow
- Jupyter Notebooks
- Pipelines
- Fairing
- Kubeflow on AWS
- Kubeflow on Azure
- Kubeflow on GCP
- Kubeflow on IBM Cloud
- Kubeflow on OpenShift

---

# OpenShift AI

# Red Hat OpenShift AI

Upstream code enhanced with operational excellence

## Open Data Hub

Community driven upstream meta-project demonstrating AI/ML platform on Red Hat OpenShift comprised of open source projects

## Red Hat OpenShift Data Science: cloud service

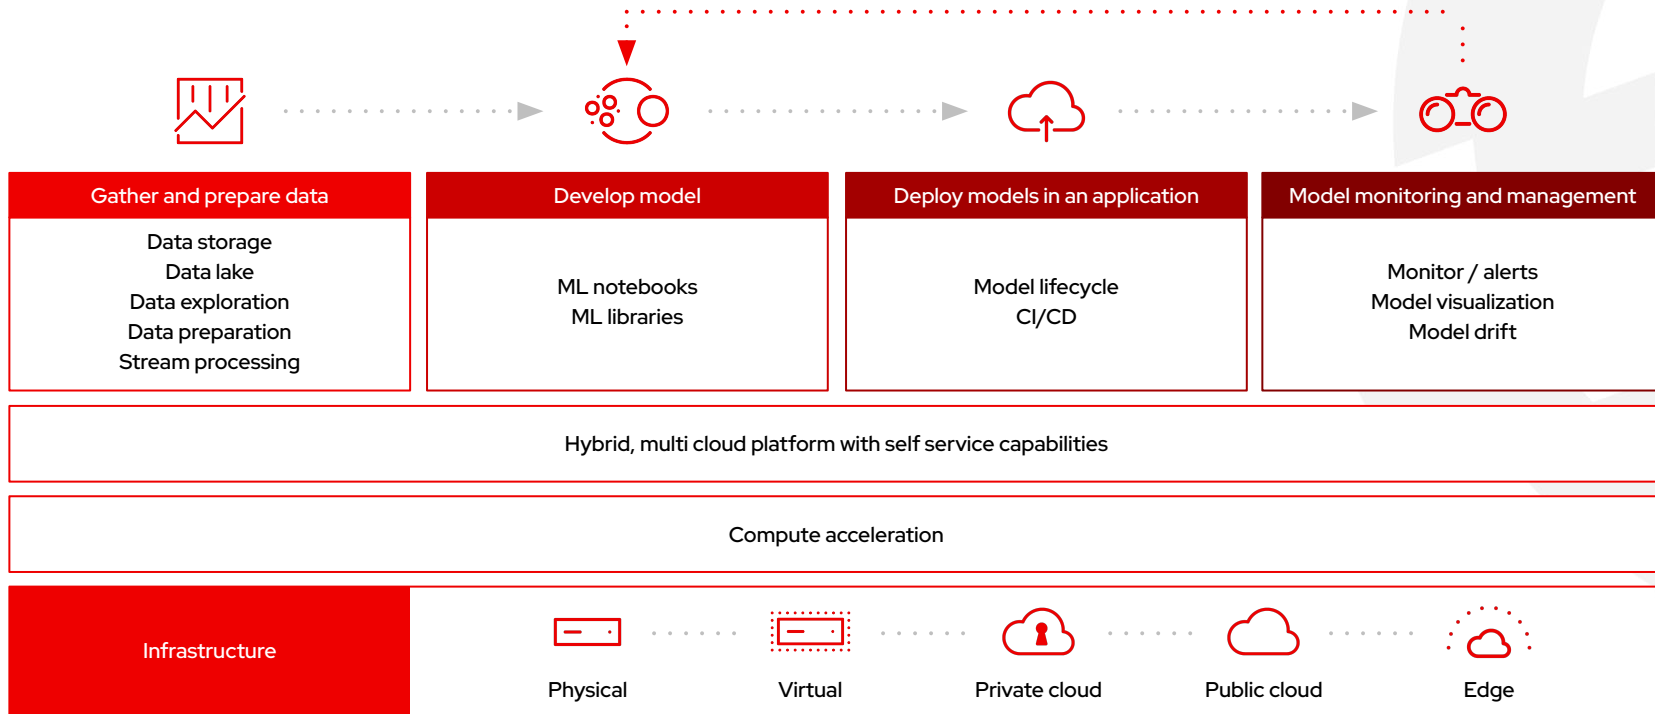
Subset of Open Data Hub delivered as a cloud service on Red Hat OpenShift Managed services with optional ISV offerings

## Red Hat OpenShift Data Science: self managed

Fast moving software stream mirrors the release frequency and capabilities of the cloud service delivered in a self-managed offering for on-premise OpenShift customers



# Conceptual machine learning architecture



DEMO!



# Red Hat OpenShift Data Science

A cloud service that gives data scientists and developers a powerful AI/ML platform for building intelligent applications.


Try OpenShift Data Science in the sandbox

[Object Detection Workshop](#)

<https://red.ht/rhods-sandbox>



# Resources



Learn more about [OpenShift Data Science](#)



Sign up at [developers.redhat.com](https://developers.redhat.com)

Find out more about Red Hat's products and what we do for developers



## Modernizing Enterprise Java


Find out more on Java and Kubernetes with [this free book](#) from Red Hat Developer portal!

## [Developer Sandbox for OpenShift](#)

Start your OpenShift experience for free in four simple steps

## [OpenShift Data Science Sandbox](#)

Start your OpenShift Data Science experience for free



# Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

 [linkedin.com/showcase/red-hat-developer](https://www.linkedin.com/showcase/red-hat-developer)

 [facebook.com/RedHatDeveloper](https://www.facebook.com/RedHatDeveloper)

 [youtube.com/RedHatDevelopers](https://www.youtube.com/RedHatDevelopers)

 [twitter.com/rhdevelopers](https://twitter.com/rhdevelopers)

