



Red Hat Dallas Emerging Tech Summit

December 5, 2019

Blueprint for AI on Red Hat OpenShift with Open Data Hub

Sherard Griffin, Senior Manager AI Center of Excellence

Introductions



Sherard Griffin

**Senior Manager, AI Center of Excellence
Raleigh, NC**

**Responsible for OpenDataHub.io and
Red Hat's AI partner enablement**

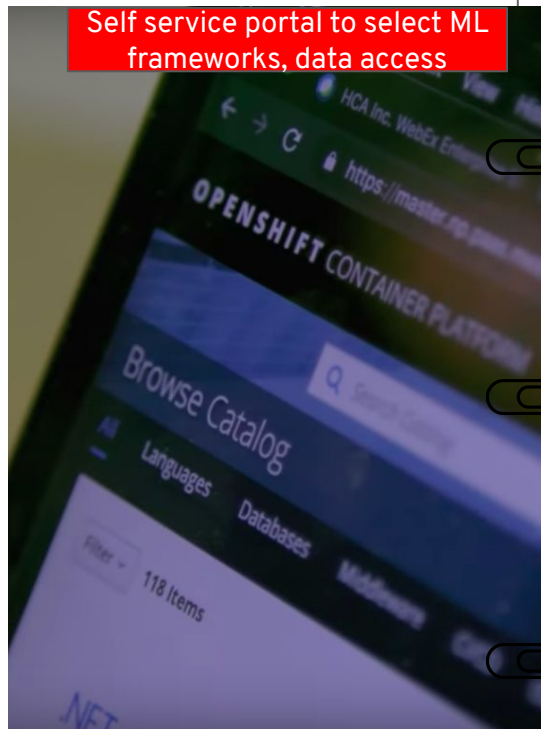
What do Data Scientists Want?

“Self-service cloud like” experience for my machine learning projects

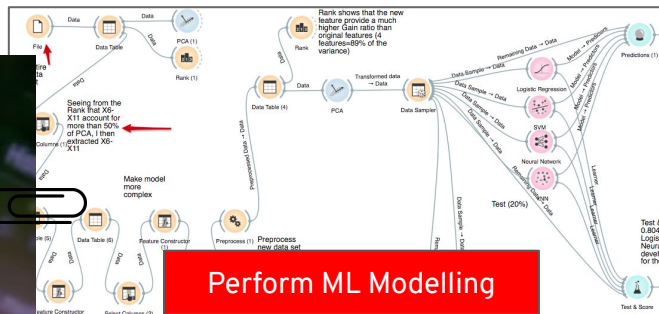
Access to a rich set of modelling frameworks, data, and computational resources

Collaborate with colleagues

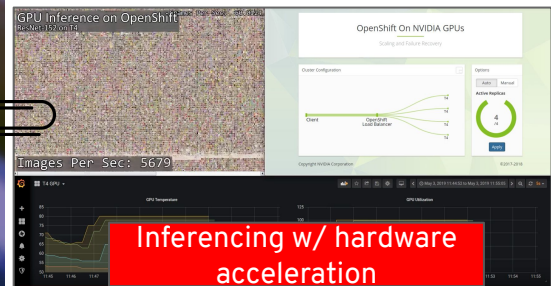
Deliver my work into production with speed, agility and repeatability to drive business value




Self service portal to select ML frameworks, data access



Perform ML Modelling



Inferencing w/ hardware acceleration



ML model deployment in app dev

Today's Challenges with AI Infrastructure

Limited non-shared resources

Modeling takes a long time or can't achieve desired accuracy without pooled resources

Difficult to share and collaborate

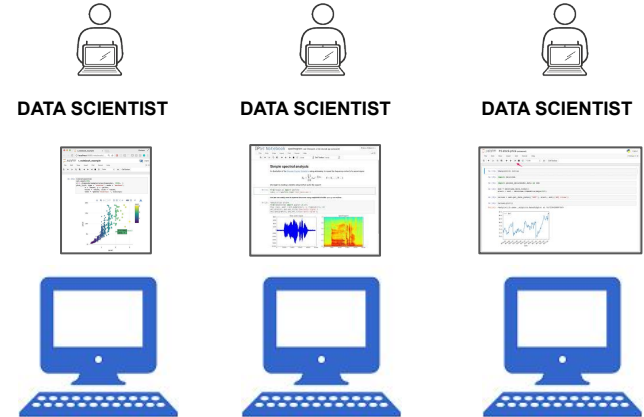
Manual, error-prone methods of sharing data science work with others

Team(s) of data scientists and developers

Limited multi-user environments

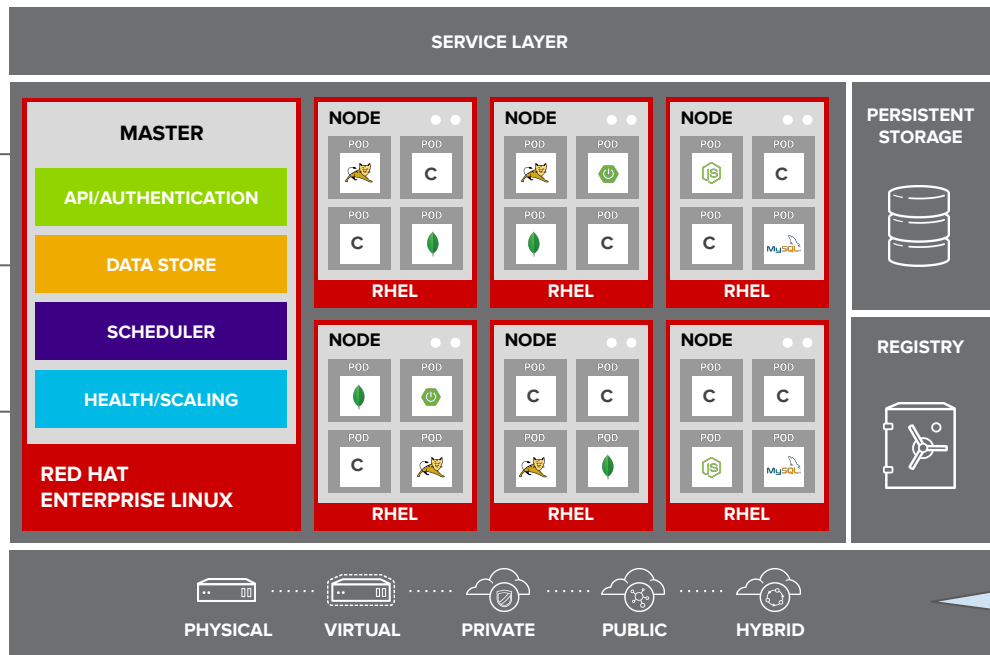


Challenging production deployments

Complicated methods for deploying models into production pipeline



How Does OpenShift Help with AI/ML workloads?

Software Development Lifecycle



ML services, load-balanced, and scaled

ML microservices scheduled and orchestrated on shared resources

ML deployed across clouds, datacenter, and edge

ML in Production

AI/ML on OpenShift Momentum is Strong

Healthcare and public sector



Public cloud alternative
for research



Data-driven diagnosis



Data-driven diagnosis

Ministry of Defence (Israel)

Jupyter notebooks as a service

Automotive



Connected Drive &
Autonomous Driving
Back-end AI Platform

Financial

RBC Bank

Containerized Apache
Spark

Discover Financial Services

Oil and gas



Democratize data science for oil
and gas exploration

Why Red Hat for your AI/ML lifecycle



Proven track record

Expertise helping top brands operationalize containers & Kubernetes powered hybrid cloud for AI/ML.



Comprehensive portfolio

Helps complete the AI/ML architecture.



Powerful partnerships


Established strategic partnerships with key AI/ML ISVs.



Open source advocate


A leader and trusted provider in the open source space.

Red Hat portfolio helps complete the ML architecture



ML Software tools  Red Hat Middleware Portfolio

Data Services - Databases (SQL, NoSQL, etc.), Data Lake, etc.  Red Hat Middleware Portfolio

Hybrid, Multi Cloud Platform with self service capabilities  Red Hat OpenShift

Compute Acceleration (GPU, FPGA, TPU)

Software-defined Infrastructure  Red Hat Enterprise Linux  Red Hat OpenShift Container Storage  Red Hat OpenStack Platform  Red Hat Virtualization  Red Hat Ceph Storage

Red Hat is the design partner for AI/ML solutions across industries

Healthcare and public sector



Data driven diagnosis



Data driven diagnosis

Ministry of Defence (Israel)

Jupyter notebooks as a service

Automotive



Connected Drive & Autonomous Driving

Financial

RBC Bank

Containerized Apache Spark

Discover Financial Services

Oil and gas



Democratize data science for oil and gas exploration

What is Open Data Hub?



OpenDataHub.io

Reference Architecture Community Project

Blueprint for building an AI-as-a-service platform on OpenShift

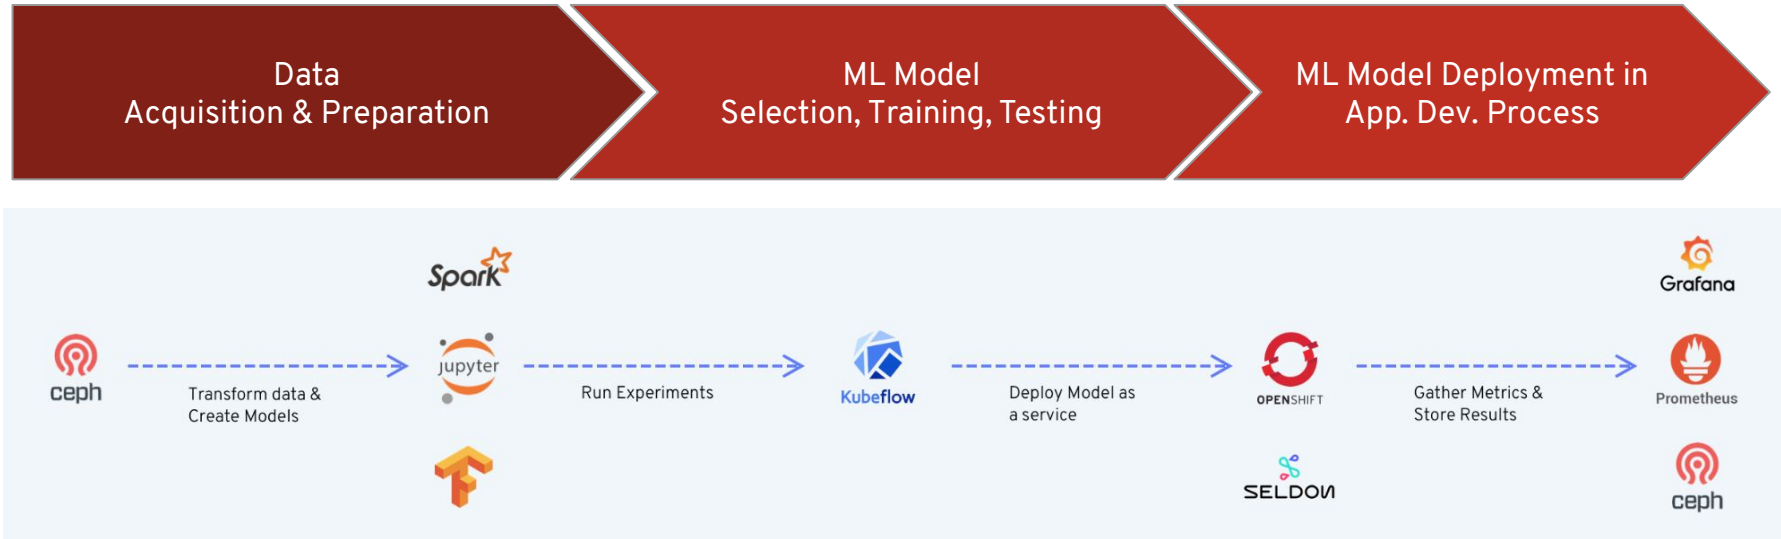
Meta-Operator

Integrates the best open source AI/ML and data engineering community projects as an operator

Production-ready

Running in Red Hat's datacenter and used by teams for over a year

End-to-End AI/ML Process on OpenShift





Open Data Hub v0.5 Operators



Object Storage



Interactive Notebooks



Model Lifecycle



Data Ingestion



Metadata Management Hive Metastore

Workflows



Data Exploration and Engineering



Service Monitoring



Upstream Community Projects in the AI/ML Space



NVIDIA NGC
GPU optimized
and curated



Kubeflow



ML toolkit and lifecycle
Kube



Tensorflow



Others

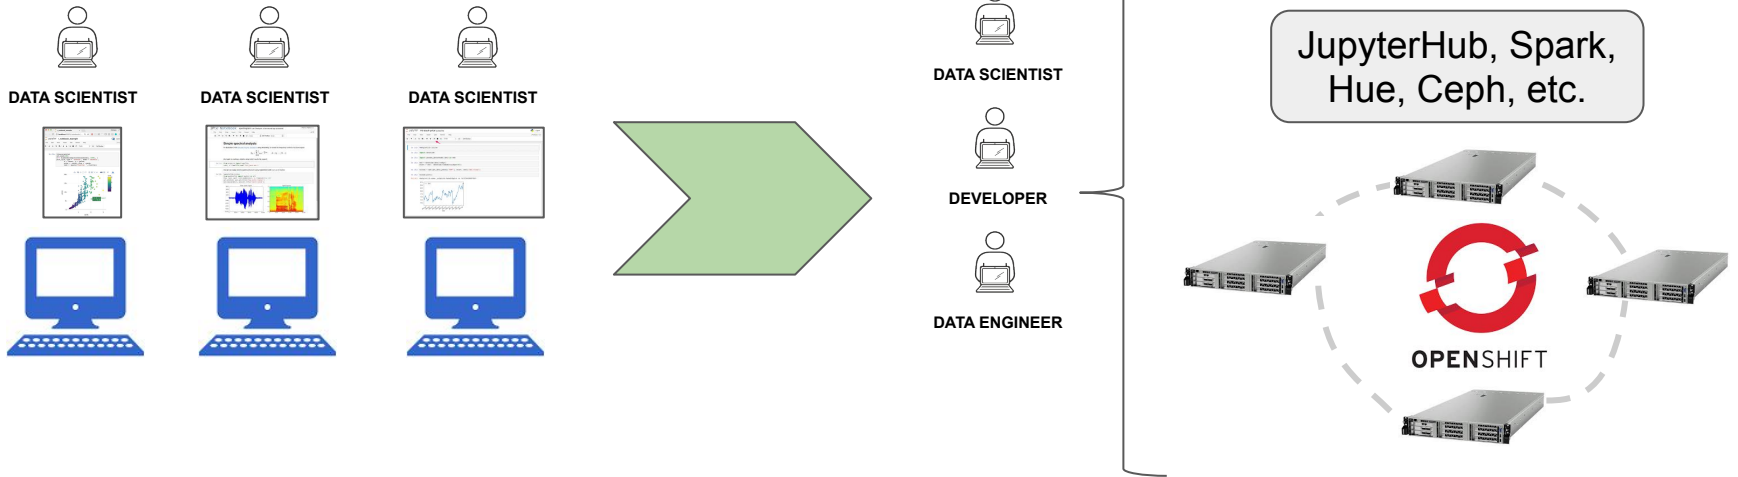


AI-as-a-service platform based on
OpenShift, Ceph, Kafka, JupyterHub,
Spark, and more

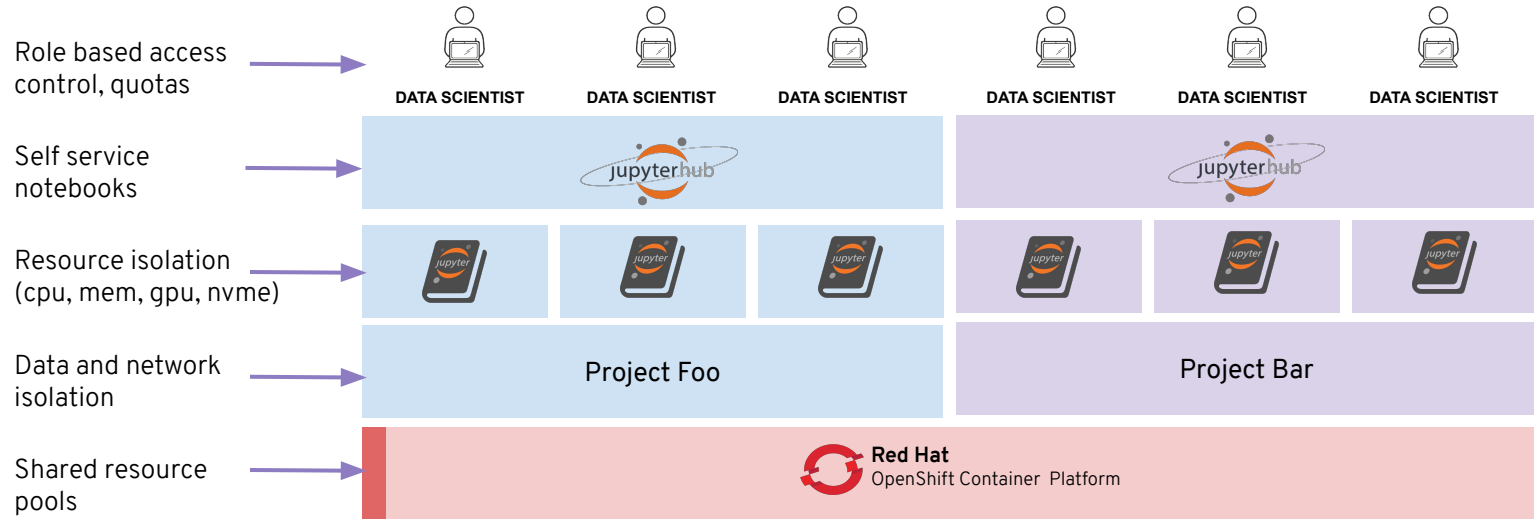


Home for K8's community to share
operators for various apps/tools

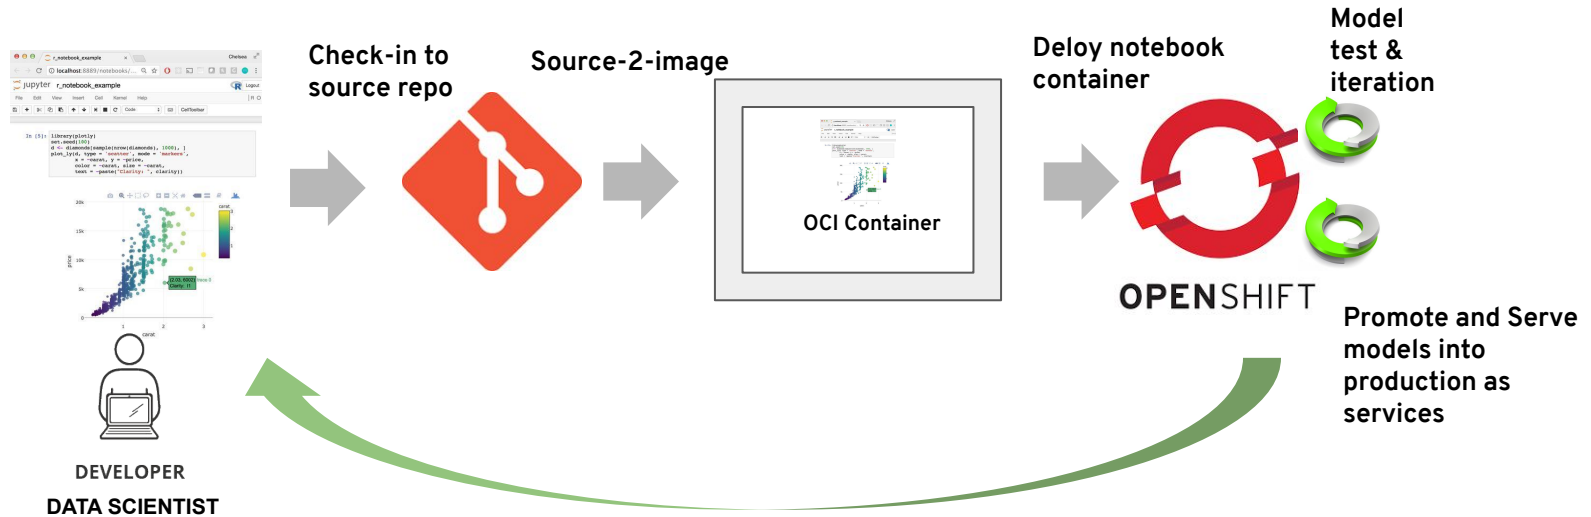
Data Scientists and Engineers Can Share Resources on OpenShift



Notebooks as a Service in Open Data Hub



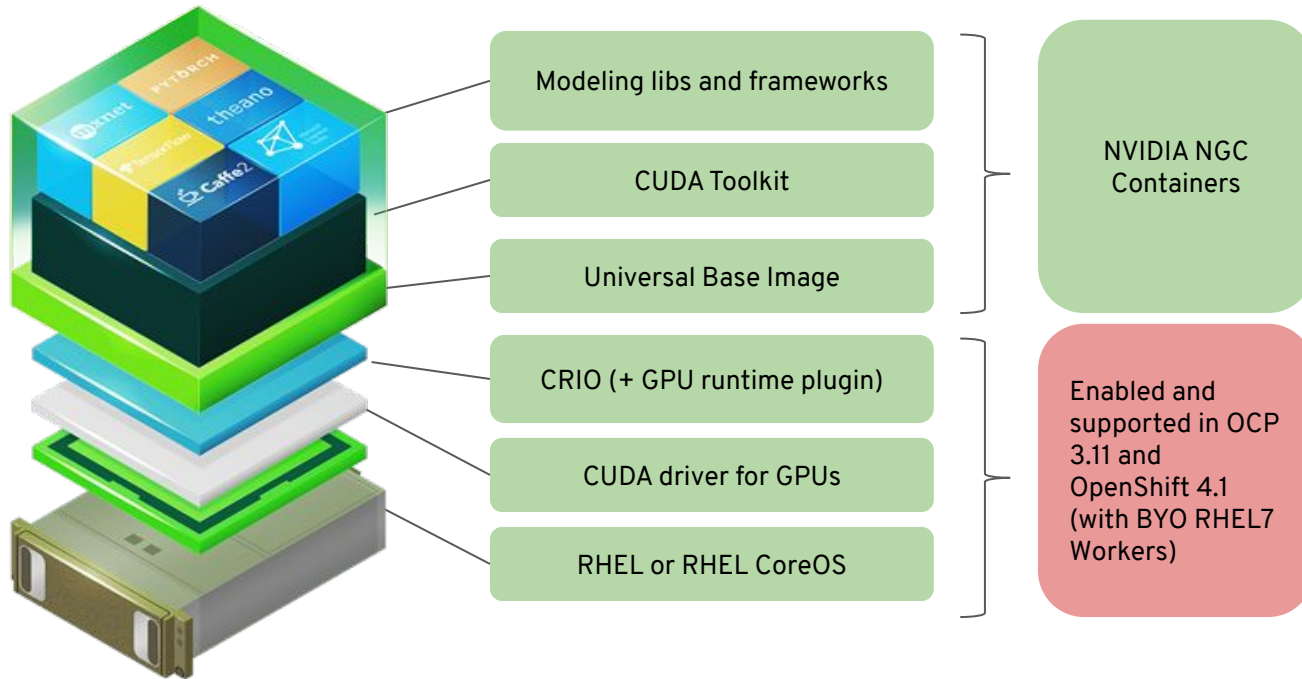
CI/CD to Production for Data Scientists



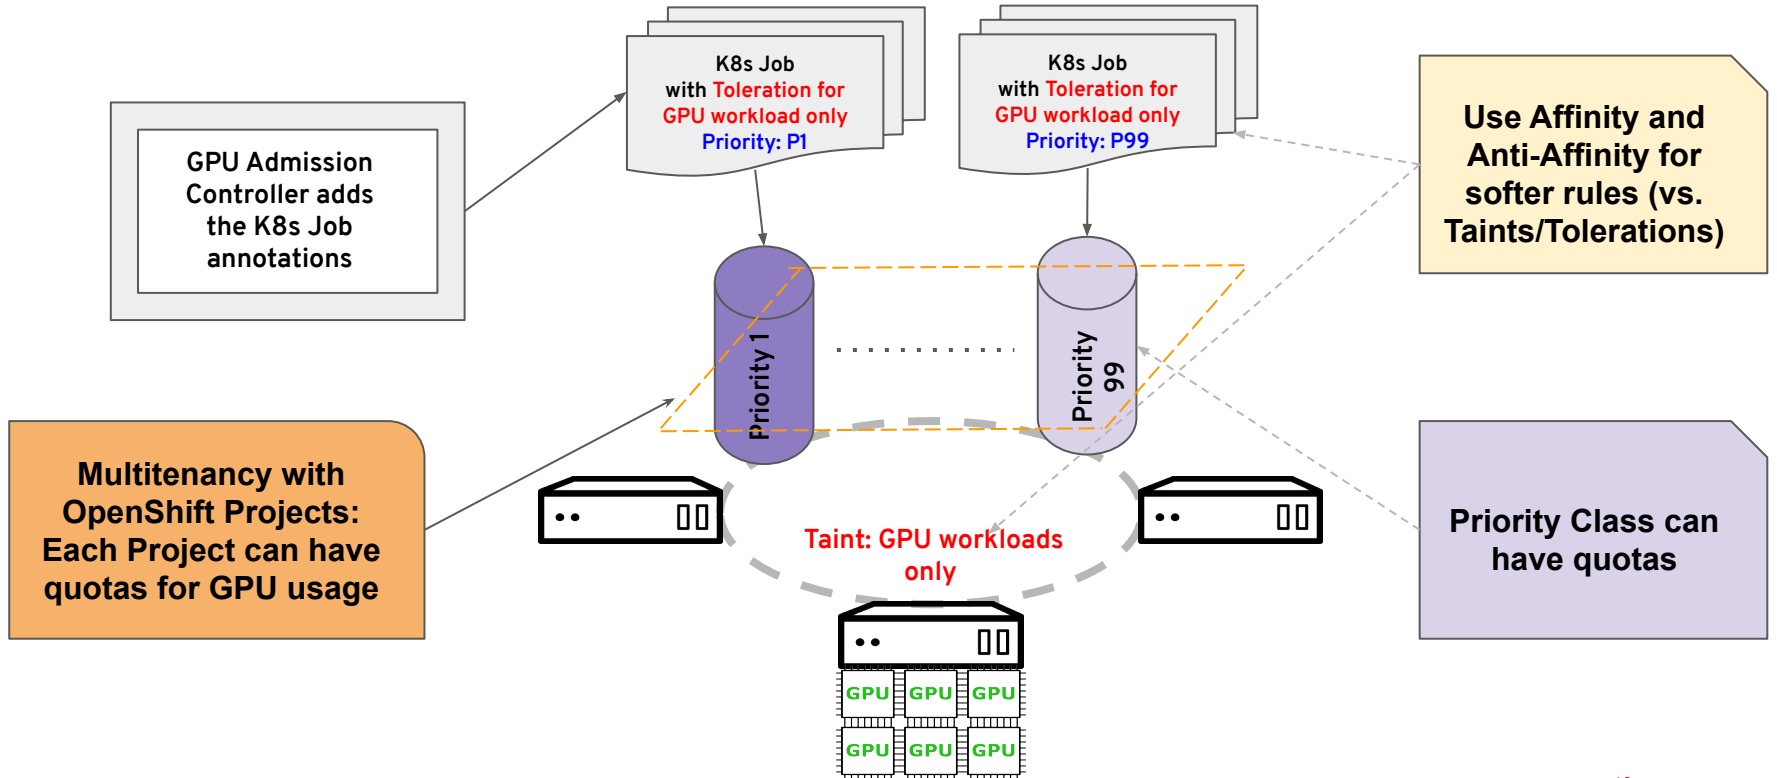
References

Source 2 Image for Jupyter Notebooks: <https://blog.openshift.com/jupyter-on-openshift-part-3-creating-a-s2i-builder-image/>


NVIDIA GPU Enablement for Containers



Sharing GPUs in OpenShift for Model Training



Using GPUs in Open Data Hub



Spawner Options

JupyterHub Server Image

Select desired notebook image

Deployment sizes

Select desired container size


GPU:

Environment Variables

AWS_ACCESS_KEY_ID:

AWS_SECRET_ACCESS_KEY:

Select number of GPUs to use via UI





Demo

AI Partner Program

Mission: With partners, make best-in-class AI/ML hybrid/multi-cloud ready for mainstream and enterprise markets

- Primary Focus on OpenShift - kubernetes-powered multi/hybrid-cloud AI/ML platform
- Hardware Ecosystem for Performance - GPUs
- SW Ecosystem Enablement
 - Projects w/ Vendor Support - Certified vendors w/ commercial support
 - Balance Platforms & Components - Provide choice
- Cloud-Like Experience - Operators, operatorhub.io, embedded operatorhub
- Certifications with Red Hat Partner Connect (for Container and Operator certifications)

Via:

- AI Reference Architecture - Open Data Hub - RH Ref Arch for AI for upstream/products/partners
- AI Use Case Definitions - define core use cases - RH + Partners Enable these use cases

Who Are We Working With?



Just to name a few...



Find Out More!

Deploy in Minutes


<https://opendatahub.io/docs.html>


Join the Community

<https://opendatahub.io/community.html>

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/company/red-hat](https://www.linkedin.com/company/red-hat)

 [youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)

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

 twitter.com/RedHat

STAY ENGAGED

Developers.redhat.com

Your access point for no-cost developer tools and product subscriptions, how-tos, and demos

Red Hat User Groups

Meetups for networking and tech deep dives
www.meetup.com/Dallas-Red-Hat-Users-Group/

DevNation

Virtual and live events
Catch replays at
<https://developers.redhat.com/devnation/>

Next.redhat.com

Stay in touch with the Office of the CTO