

# Accelerating MLOps with Kubernetes, CI/CD & GitOps – *for a Retail Coupon App*

Audrey Reznik  
Sr. Principal Software Engineer  
Red Hat

# Agenda

- **Case Study:** Deliver an intelligent retail (coupon) application
  - Overview
- AI/ML models and the **work of data scientists**
  - MLOps
- **Productizing a ML model** with OpenShift
  - GitOps and Pipelines
- **Demo:** Globex - AR Coupons

# Case study: Retail Coupon App

Globex - AR Coupons

# Globex Retail Coupons



Is this shirt on sale? If so, how much is the discount?

## What's the Discount?

Give a customer the ability to find merchandise discounts, for shirts, as they browse clothing in a department store.

Leadership Team - can you build an intelligent app (PoC) with MLOps?

- *Streamlining the process of taking machine learning models to production, and then maintaining and monitoring them.*

Developed Globex Retail Coupons app.

# People + Collaboration



**Business  
leadership**



**Data  
engineer**



**Data  
scientist**



**App  
developer**



**ML  
engineer**



**IT  
operations**

## Data Science & MLOps applications and Managed services



(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



(Red Hat OpenShift Data Science)

Cloud service that gives Data Scientists and Developers a powerful AI/ML platform for building intelligent applications

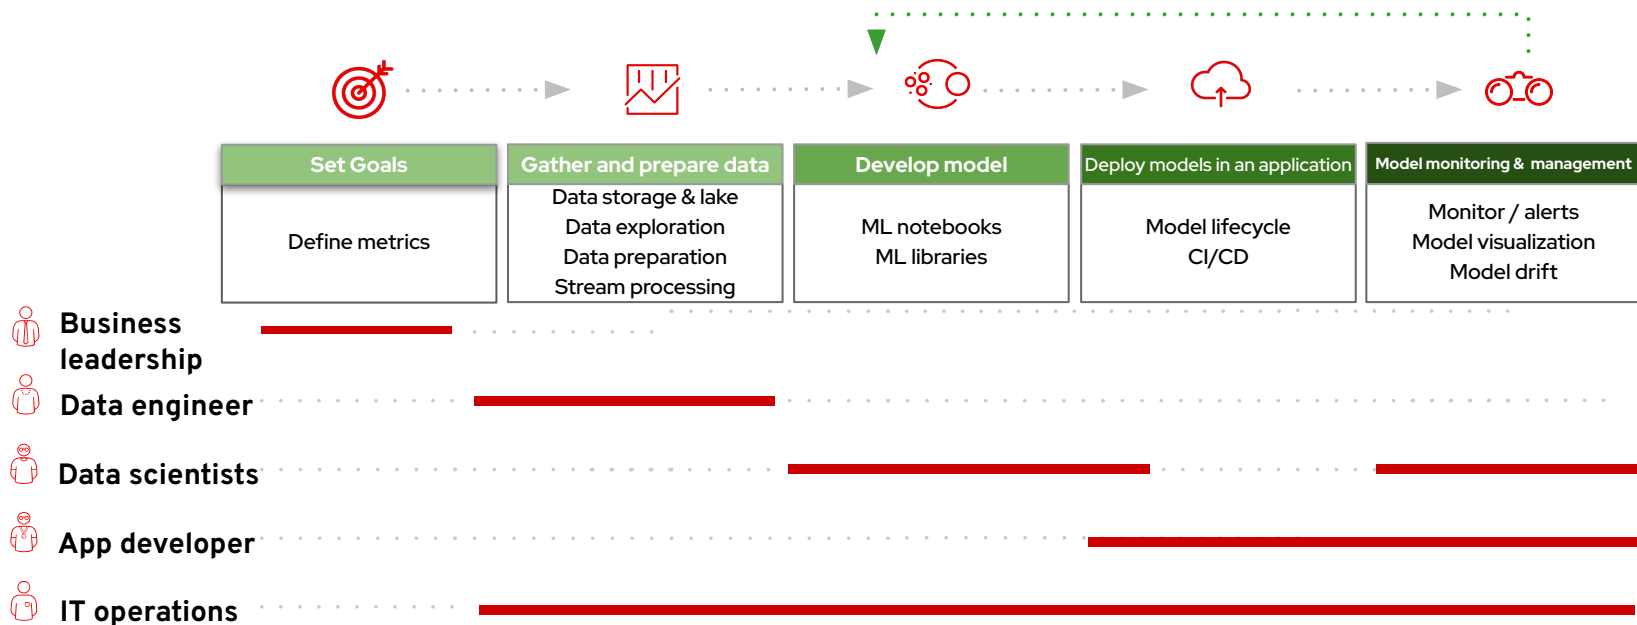


Toolkit and workflow to build reproducible container images from source code

# AI/ML model and the work of a data scientist

Level Setting

# Managed Services (and Personas) for a Model Lifecycle





# Developing a ML Model

```

To begin, we import a range of python packages:

[ ]: # For running inference on the TF-Hub module.
import tensorflow as tf
import numpy as np
from PIL import Image
from PIL import ImageColor
from PIL import ImageDraw
from PIL import ImageFont

import matplotlib.pyplot as plt
import os

os.environ['TF_CPP_MIN_LOG_LEVEL'] = '3'
print('Imported libraries')
print(tf.version.VERSION)

Connect to data in s3

Before we get started with exploring our model, we first need to load some data to test the
model against. This data is stored in an s3 bucket, and we connect to it using the boto3
library. The boto3 library was built into the standard data science notebook image, which we
selected from the spanner page. As such, it is already installed in our environment.

[ ]: import boto3
from botocore import UNSIGNED
from botocore.client import Config
print('Imported s3 libraries')

[ ]: s3 = boto3.client('s3', config=Config(signature_version=UNSIGNED))
s3.download_file('retailshopper', 'retail1.jpg', 'retail1.jpg')

You should be able to see that this file has been added to your file directory on the left hand
side of the screen.
    
```

```

print('Defined image display functions')

[ ]: result = {key: value.numpy() for key, value in result.items()}
image_with_boxes = draw_boxes(
    tf.groceries.numpy(), result["detection_boxes"],
    result["detection_class_entities"], result["detection_scores"])
display_image(image_with_boxes)


So, you can see we've got some predictions for 'bottle', and we can now see what parts of the
image have been identified as such.

However, if we look at the detection scores, some of the predictions are not made with high
confidence:

[ ]: result["detection_scores"]

Let's only draw boxes if the model was over 30% confident that the object was correctly
classified.

[ ]: per = 0.3
image_with_boxes_per = draw_boxes(
    tf.groceries.numpy(), result["detection_boxes"][result["detection_scores"]
    result["detection_class_entities"][result["detection_scores"] > per],
    result["detection_scores"][result["detection_scores"] > per])
display_image(image_with_boxes_per)
    
```



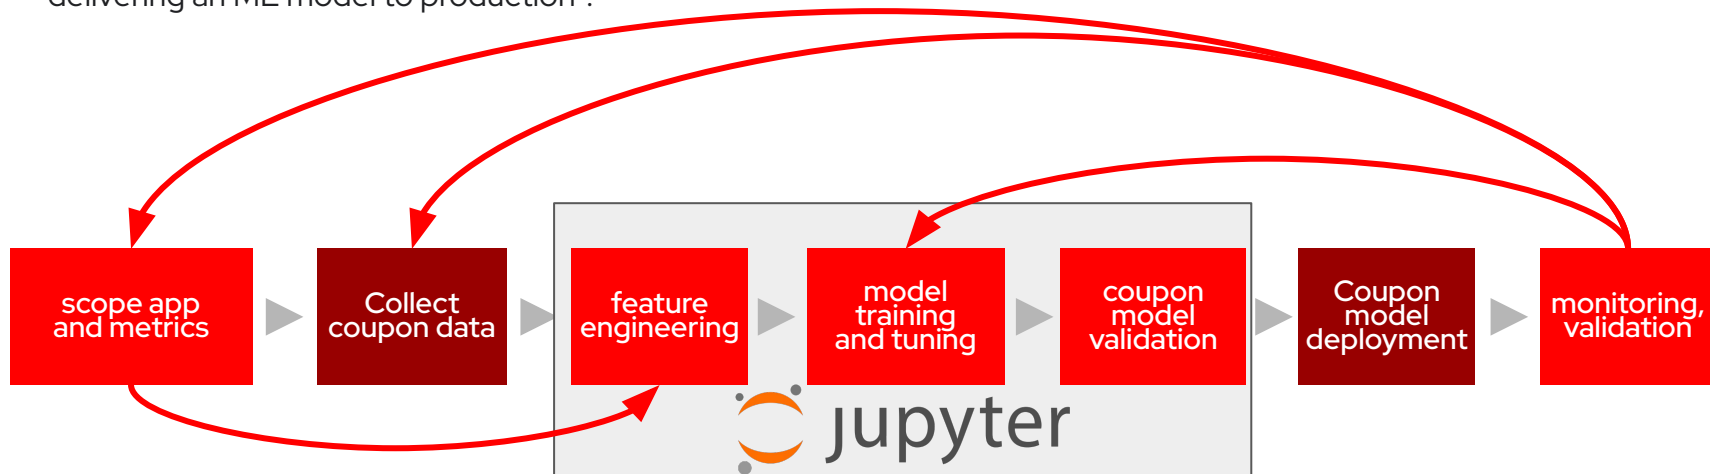
# **DEMO - Red Hat OpenShift Data Science Platform**

**Create an Object Detection model**

**Place code/model in Git repo**

# MLOps

How do we automate this process of delivering an ML model to production ?



# Productizing ML models with OpenShift GitOps and Pipelines

# OpenShift Pipelines

Kubernetes-native CI/CD



Built for Kubernetes



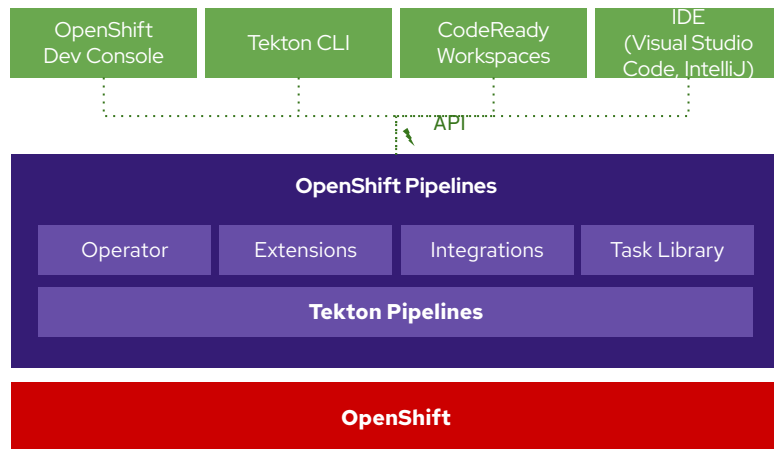
Scale on-demand



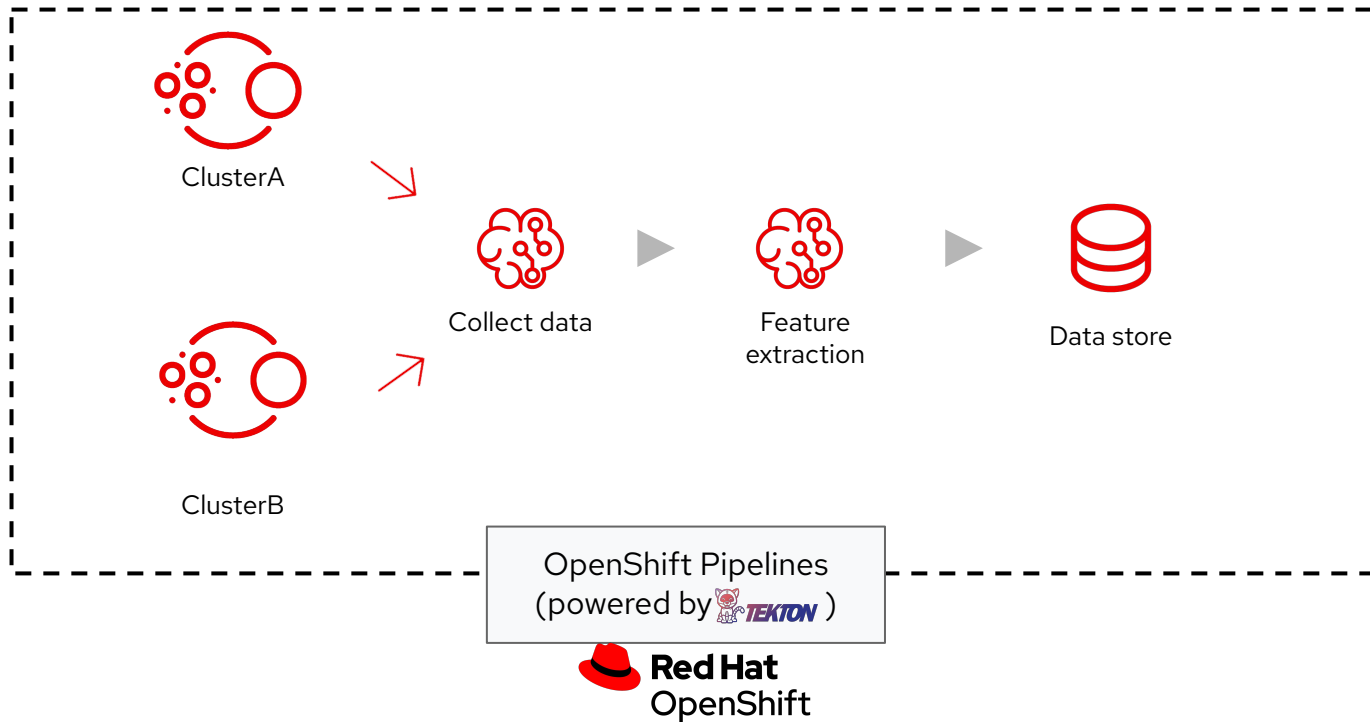
Secure pipeline execution



Flexible and powerful



# MLOps: Data Pipeline


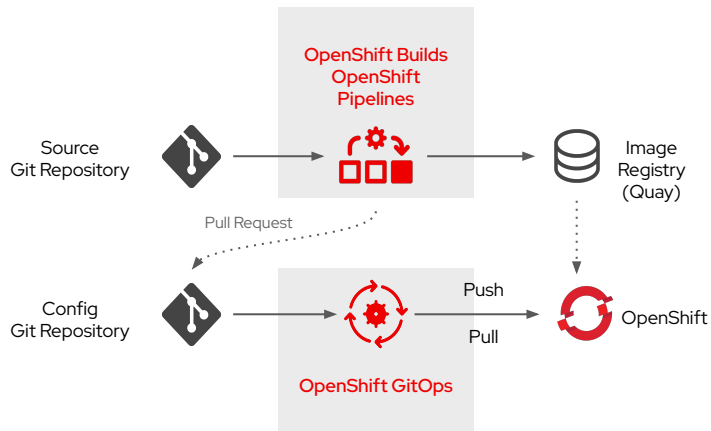


## **DEMO - Apache Kafka**

**Create a service to process data streams**

# OpenShift GitOps

## Kubernetes-native GitOps



Multi-cluster config management



Automated Argo CD install and upgrade



Opinionated GitOps bootstrapping

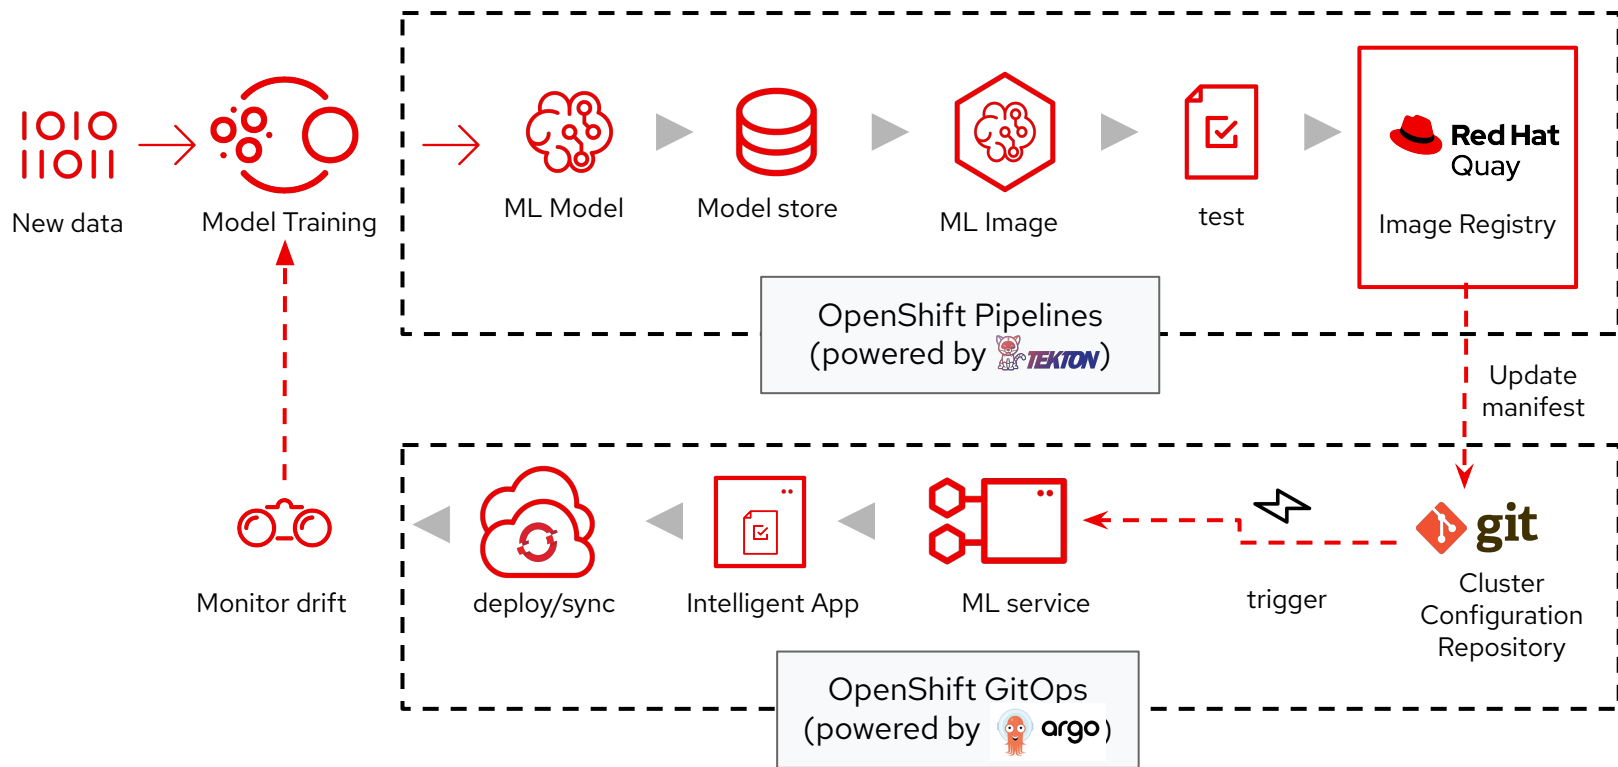


Deployments and environments insights

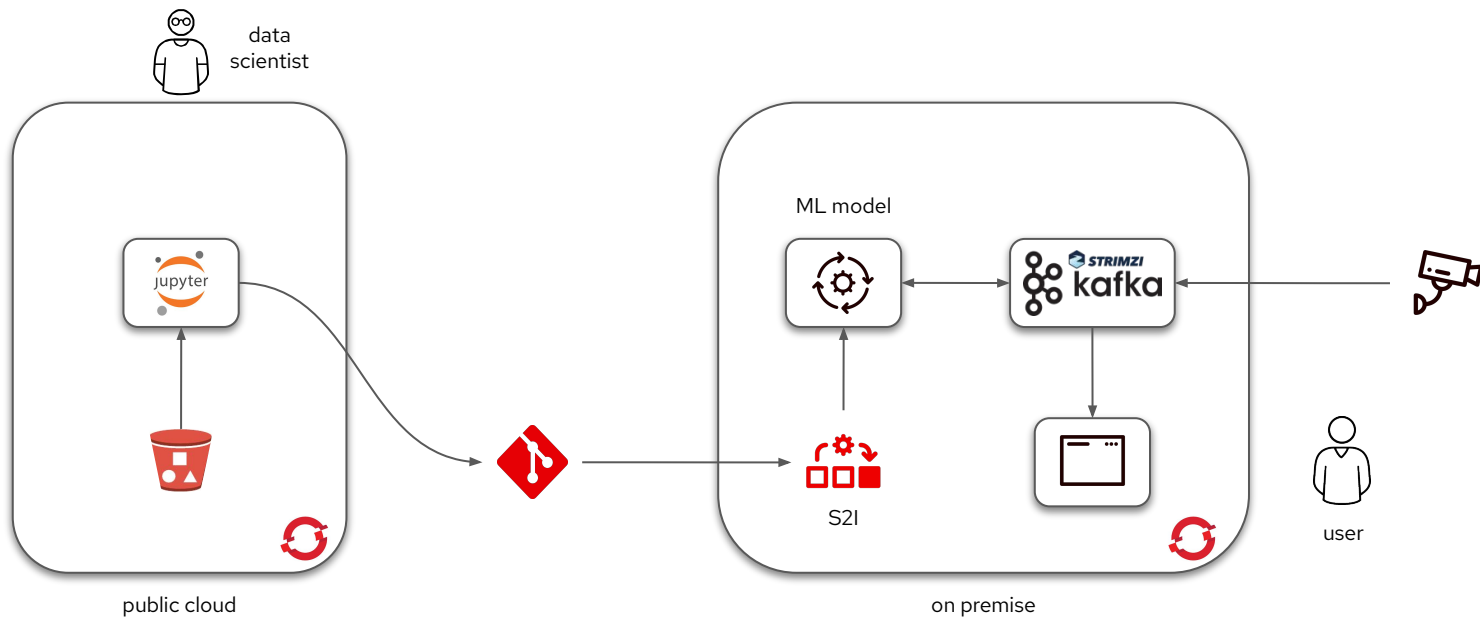
*Consistency when deployed to different clusters in different environments, such as: development, staging, & production.*



# MLOps with Red Hat OpenShift



# Deploying an intelligent application



# **DEMO - OpenShift Platform**

## **Creating a containerized Image for Deployment**

# Globex - AR Coupons

Example use case for MLOps with OpenShift  
GitOps and Pipelines


## Demo - Globex AR Coupons

1. Scan this QR Code with your phone



2. Use the app to take a picture of these items





# Summary



# Accelerate MLOps with OpenShift

Red Hat OpenShift is for every innovator



**Red Hat**  
**OpenShift**



## Ready for IT operators

- ▶ Automate processes. Reduce complexity.
- ▶ Operate more securely from end to end.



## Empowering developers and data scientists

- ▶ Self service, develop fast with familiar tools and integrated compute acceleration.
- ▶ Rapidly deliver without roadblocks.



## Proven and consistent for business leaders

- ▶ Choose a platform to power business today.
- ▶ Create a cloud strategy for the future.

# Red Hat is a strategic partner for AI/ML solutions across industries



Data driven diagnosis

**Royal Bank of  
Canada**

Containerized Apache Spark &  
AI-powered Intelligent Apps

**BMW  
GROUP**

Connected Drive &  
Autonomous Driving

**ExxonMobil.**

Democratize data science for oil  
and gas exploration



Data driven diagnosis



Public Health  
England

Diagnostics and surveillance  
of diseases and outbreaks



An Roinn Talmhaíochta,  
Bia agus Mara  
Department of Agriculture,  
Food and the Marine

NLP text analytics to speed up  
farmers' grant applications



AI-powered apps to transform  
customer experience, drive revenue

**Ministry of Defence (Israel)**

Jupyter notebooks as a service

**Discover  
Financial  
Services**



"Ignite" - Modular AI-platform

Orchestrating a brighter world

**NEC**

Speed traveler experience at Narita  
International Airport, Japan



# 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.

**Try RHODS for free:**

<https://developers.redhat.com/products/red-hat-openshift-data-science>

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

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

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

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