

Accelerate Al using OpenShift Data Science

Erica Langhi Solution Architect - Red Hat



Al is becoming a part of our everyday lives





Project abandonment rate

On average, organisations abandon over one third of their AI/ML projects before they enter production





Q. To the best of your knowledge, what percent of AI/ML projects in proof-of-concept are abandoned before production? Base: ML in production (n=287). Source: Voice of the Enterprise: AI and Machine Learning, Infrastructure 2023.



Market Intelligence

0

Infrastructure is the main driver of project abandonment



Q. In which of these areas did challenges lead to project abandonment at your organization within the past 12 months? Pleases dect as many as apply. Base: All respondents (n=686). Source: Voice of the Enterprise: Al and Machine Learning, Infrastructure 2023.



© 2023 S&P Global.

Complexity of operationalising models





Red Hat's AI/ML engineering is 100% open source





Provide your data science team with the tools they need







8

Hybrid MLOps platform

Collaborate within a common platform to bring IT, data science, and app dev teams together

Available as fully managed cloud service or traditional software product on-prem or in the cloud!



Model development

Conduct exploratory data science in JupyterLab with access to core AI / ML libraries and frameworks including TensorFlow and PyTorch using our notebook images or your own.



Model serving & monitoring

Deploy models across any cloud, fully managed, and self-managed OpenShift footprint and centrally monitor their performance.



Lifecycle Management

Create repeatable **data science pipelines** for model training and validation and integrate them with devops pipelines for delivery of models across your enterprise.



Increased capabilities / collaboration

Create projects and share them across teams. Combine Red Hat components, open source software, and ISV certified software.



Everything put together is MLOps





The model operationalization life cycle

Red Hat OpenShift Streams for Apache Kafka Pachyderm Starburst Galaxy Extract and transform data IRM ANACONDA **Red Hat** jupyte OpenShift Data Science Run experiment and create models **OpenVINO**[®] 🕺 NVIDIA Deploy models as services **Red Hat** Red Hat OpenShift OpenShift API Management Data Science OpenVINO[®] Monitor models and track performance IBM **OpenVINO**

10

Starburst Galaxy

Unlock the value of your data by making it fast and easy to access data across hybrid cloud.

Pachyderm

Brings data versioning and governance to your most precious asset

Anaconda Professional

Curated access to an extensive set of data science packages to be used in your Jupyter projects.

IBM Watson Studio

Build, run, and manage AI models at scale with Watson Machine Learning and Watson OpenScale.

NVIDIA

GPU-enabled hardware makes it easier for customers to stand up resource-intensive environments and accelerate their data science experiments.

Intel OpenVINO Notebook Images

Toolkit of pre-trained models optimized for intel processors and GPUs

Intel OpenVINO Model Server

Scalable, High-Performance serving engine



Cloud service and self-managed components





Dashboard user interface

Red Hat OpenShift Data Science					III 🌲 🕢 🔻 jdemo 👻		
Applications Enabled Explore	Explore Add optional applications to your Red Hat OpenSh	lift Data Science instance.					
Data Science Projects Model Serving	Partner managed Anaconda Professional by Anaconda	Self-managed	Self-managed	Red Hat managed	٩		
Resources	Anaconda Professional is a popular open source package distribution and Red Hat OpenShift Data Science	IBM Watson Studio is a platform for	by Intel®	A multi-user version of the notebook		III ♣ Ø ▾ jdemo ▾	
	Applications Enabled Explore Explore						
	Model Serving Intel Resources Setting	oneAPI AI Analytics Toolkit aliner (1 managed) Al Analytics Toolkit	pyter	NVIDIA GPU Add-on Self-managed NVIDIA GPU Add-on prepares NVIDIA GPU Add-on prepares	OpenVINO Self-managed OpenVINO is an open source tool beln optimize deen learning	1 kit to	
12	The A accel and a archit	AI Kit is a set of AI software tools to lerate end-to-end data science analytics pipelines on Intel® tectures.	aunch application	workloads on a single node or on multiple nodes.	performance and deploy using an inference engine onto Intel® hardw	vare.	



Dashboard resources

■ Red Hat OpenShift Data Science					🏭 🌲 🛛 🕶 jdemo 🕶			
Applications Chabled Explore Date Colored Brainete	Resources Access all learning resources for Red Hat OpenShift Data Science and supported applications. To learn more about Red Hat OpenShift Data Science, view the documentation.							
Model Serving	Data analysis Data cleaning	Q Search	Sort by name 👻 🕽 🛓	•	19 of 50 items			
Resources Settings	Data management Data preprocessing Data visualization Getting started Model development Model monitoring Model optimization Model serving Model serving Notebook environments Package management	★ <p< td=""><td>VIDIA Creating a Machine Learning Model using the NVIDIA GPU Add-on. by NVIDIA GPU Add-on. Quick start S minutes Install GPU add-on Open</td><td>Image: Second Secon</td><td></td></p<>	VIDIA Creating a Machine Learning Model using the NVIDIA GPU Add-on. by NVIDIA GPU Add-on. Quick start S minutes Install GPU add-on Open	Image: Second Secon				
	Enabled state Enabled (19) Not enabled (31) Resource type Documentation (10) HowTo (14)	Getting started with Pachyderm concepts by Pachyderm Tutorial 15 minutes	How to install Python packages on your notebook server by Jupyter How-to 15 minutes	©penVIN⊙ How to serve a model using OpenVINO Model Server by OpenVINO How-to 10 minutes				



Ecosystem support



OpenShift Data Science Ecosystem

- Tier 0 (supported)
- Tier 1 (staging)
- Tier 2 (unsupported)
- Certified Integrations (mixed)



How we differentiate versus alternatives



Hybrid cloud

Deploy models in containerized format for intelligent apps on-premise or in cloud



Easy to manage

Simple configurations on a secure and proven platform, that you can scale up or down with low effort



Collaborate

Collaborate on a common platform to bring IT, data science and application development teams together



Open Source

Red Hat tracks changes and fixes to open source AI/ML tooling and enables customer access to upstream innovation



Customers

HCA*	Royal Bank of Canada	BMW GROUP	E≫onMobil
Data driven diagnosis	Containerized Apache Spark and Al-powered intelligent apps	Connected drive and autonomous driving	Diagnostics and surveillance of diseases and outbreaks
Boston Children's Hospital	Public Health England	†) Galicia	W TURKCELL
Data driven diagnosis	Diagnostics and surveillance of diseases and outbreaks	NLP solution to onboard corporate banking clients faster	Al-powered apps to transform customer experience and drive revenue
משרד הביטחון MINISTRY OF DEFENCE	NTT EAST	KPMG	NEC
Jupyter notebooks as a service	Image and video analytics at the edge to detect congestion, improve retail customer experience, and reduce shoplifting	"Ignite"-modular AI-platform	Speed traveler experience at Narita International Airport in Japan



17

Red Hat approach strategy for AI



Mission : To be the recognized industry leading platform for AI workloads and a trusted advisor to customers navigating the open source AI landscape.

Use the foundational value of containers, Kubernetes, DevOps with Red Hat **OpenShift AI** to accelerate AI/ML projects from pilot to production.

Provide Red Hat OpenShift Data Science as a fully supported environment for data scientists to develop, train, test, deploy, and monitor AI/ML models.

Promote and show the depth and breadth of our AI/ML ecosystem (ISVs and hardware acceleration).

Provide Red Hat Consulting Services as a modular approach for customer implementations.





Red Hat OpenShift Al

An AI-focused portfolio that provides tools across the full lifecycle of AI/ML experiments and models and includes Red Hat OpenShift Data Science.

Red Hat OpenShift Data Science is a platform for building, training, testing, and deploying models optimized for hybrid cloud environments



Want to learn more?





https://www.redhat.com/en/technologies/cloud-computing/opens hift/openshift-data-science/trial

https://ai-on-openshift.io/





Thank you



linkedin.com/company/red-hat



youtube.com/user/RedHatVideos



facebook.com/redhatinc



twitter.com/RedHat

