

Red Hat OpenShift Data Science

A cloud service for rapid AI/ML development, deployment, and scaling

Highlights

- Rapidly develop, train, test and deploy containerized machine learning models in the cloud, without having to design and deploy Kubernetes infrastructure.
- Conduct exploratory data science in Jupyter notebooks with access to core AI/ML libraries and frameworks, including TensorFlow and Pytorch.
- Publish models as end-points via source-to-image (S2I) for integration into intelligent applications; rebuild and deploy based on changes to the source notebook.
- Streamline model training and workflows on a supported platform with a full complement of essential tools.

[Open Data Hub](#) is a blueprint for building an AI-as-a-Service (AIaaS) platform on [Red Hat OpenShift](#) and [Red Hat Ceph® Storage](#). It inherits from upstream efforts such as [Kafka Strimzi](#) and [Kubeflow](#), and is the foundation for Red Hat's internal data science and AI platform.

 facebook.com/redhatinc
 [@RedHat](https://twitter.com/RedHat)
 linkedin.com/company/red-hat

Accelerating artificial intelligence and machine learning deployments

Artificial intelligence (AI), machine learning (ML), and deep learning (DL) have rapidly become critical for businesses and organizations. According to IDC, "AI is profound and is impacting businesses and organizations across industries. AI is everywhere across the technology stack."¹ Deploying these technologies, however, can be complicated. As data scientists strive to build their models, they often encounter a lack of alignment between rapidly evolving tools, impacting productivity and collaboration among themselves, software developers, and IT operations. Scaling AI/ML deployments can be resource-limited and administratively complex while requiring expensive GPU resources for hardware acceleration. Popular cloud platforms offer scalability and attractive toolsets, but those same tools often lock users in, limiting architectural and deployment choices.

Red Hat® OpenShift® Data Science is a managed cloud service offering based on the open source [Open Data Hub](#) project. Data scientists and developers can rapidly develop, train, test, and iterate ML/DL models with full support, allowing them to focus on their modeling and application development without waiting for infrastructure provisioning. Available as an add-on to [Red Hat OpenShift Dedicated](#) and [Red Hat OpenShift Service on AWS](#), the service combines Red Hat components, open source software, and technology partner offerings with the hyper-scalability of the [Amazon Web Services](#) (AWS) public cloud platform.

Red Hat OpenShift Data Science

Red Hat OpenShift Data Science offers organizations a way to rapidly deploy an integrated set of common open source and third-party tools to perform AI/ML modeling. The platform makes it simpler to exploit NVIDIA GPU-enabled hardware infrastructure without the need to stand up and perform daily management of Kubernetes on your own.

Red Hat OpenShift Data Science represents an alternative to prescriptive and opinionated AI/ML suites available from individual cloud providers. Adopters gain a collaborative open source toolset and a platform for building experimental models without worrying about the infrastructure or lock-in from public cloud-specific tools. They can then extend that base platform with partner tools to gain increased capability. Models can be exported to production environments in a container-ready format, consistently, across hybrid cloud and edge environments.

Red Hat OpenShift Data Science supports rapid experimentation with user-supplied data where the model outputs are either:

- ▶ Hosted on OpenShift Dedicated or OpenShift Service on AWS for integration into a customer-defined intelligent application.
- ▶ Exported for hosting in your environment.

¹ IDC Market Forecast. "Worldwide Artificial Intelligence Software Forecast, 2022-2026," Doc # US49571222, August 2022.



Red Hat Consulting offers several AI/ML consulting engagements for challenges like implementing MLOps and intelligent application development.

Upstream open source and commercial technology partner tools

Red Hat OpenShift Data Science provides a subset (Table 1) of the 20+ tools found in the upstream Open Data Hub project. Offering this service on Red Hat OpenShift means that organizations don't have to worry about deploying Kubernetes themselves or purchasing and installing NVIDIA graphics processing units (GPUs) for on-premise deployments. Red Hat provides regular updates to open source tools (e.g., Jupyter, Pytorch, and Tensorflow) through the managed cloud service, removing integration and testing burden. The offering also integrates several AI/ML technology partner offerings (Table 1). Additional commercial technology partner offerings can also be added from [Red Hat Marketplace](#).

Table 1 Initial Red Hat OpenShift Data Science ecosystem

AI/ML modeling and visualization tools	Jupyter Hub with predefined notebook images; TensorFlow ; PyTorch ; Anaconda (Commercial Edition is optional); Intel AI Analytics Toolkit (AI Kit) IBM Watson Studio (optional)
Data engineering	Starburst (Galaxy is optional); Pachyderm (optional)
Data ingestion and storage	Red Hat OpenShift Streams for Apache Kafka (OpenShift Dedicated optional add-on); Amazon Simple Storage Service (S3)
GPU support	NVIDIA (with GPU operator)
Model serving	Source-to-Image (OpenShift), Intel Distribution of the OpenVINO toolkit

Conclusion

Red Hat OpenShift Data Science delivers common data science tools as the foundation of an AlaaS platform integrated with partner offerings. The platform simplifies the development, training, and testing of AI/ML models, complete with a shared user interface for navigation, onboarding, and exploring partner options. Organizations can rapidly develop their AI/ML models, expanding further by adding open source tools and Red Hat technology partner solutions.



About Red Hat

Red Hat helps customers standardize across environments, develop cloud-native applications, and integrate, automate, secure, and manage complex environments with [award-winning](#) support, training, and consulting services.

[facebook.com/redhatinc](#)
[@RedHat](#)
[linkedin.com/company/red-hat](#)

North America
1 888 REDHAT1
[www.redhat.com](#)

Europe, Middle East, and Africa
00800 7334 2835
[europe@redhat.com](#)

Asia Pacific
+65 6490 4200
[apac@redhat.com](#)

Latin America
+54 11 4329 7300
[info-latam@redhat.com](#)