

Red Hat OpenShift Data Science

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

Key capabilities

- Rapidly develop, train, test, and deploy ML models in the cloud, without designing and deploying 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 using the Source-to-Image (S2I) tool for integration with intelligent apps. Rebuild and deploy based on changes to the source notebook.
- Streamline model training and workflows.

Accelerate artificial intelligence and machine learning adoption

Artificial intelligence (AI), machine learning (ML), and deep learning (DL) have rapidly become critical for businesses and organizations as they seek to convert their data to business value. However, as data scientists strive to build their models, their efforts are often complicated by a lack of alignment between rapidly evolving tools, affecting productivity and collaboration among their teams, software developers, and IT operations. On-premise resource limitations can limit scalability, such as quickly provisioning hardware. Popular cloud platforms offer the desired scale and attractive tools, but 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. Delivering integrated capabilities from upstream efforts—such as Apache Kafka, Strimzi, and Kubeflow—Open Data Hub provides an architecture for building an AI-as-a-Service platform on Red Hat OpenShift and Ceph® object storage.

With this solution, data scientists and developers can rapidly develop, train, test, and iterate ML and DL models in a fully supported sandbox environment—without waiting for infrastructure provisioning. Available as an add-on to [Red Hat OpenShift Dedicated](#) and [Red Hat OpenShift Service on AWS](#), OpenShift Data Science combines Red Hat components, open source software, and certified partner technology from [Red Hat Marketplace](#) with the public cloud hyperscalability of Amazon Web Services (AWS).

Develop, model, experiment, and deploy more rapidly

Red Hat OpenShift Data Science lets organizations quickly deploy an integrated set of common open source and partner tools to perform AI/ML modeling. The platform makes it simpler to use hardware infrastructure without requiring time-consuming Kubernetes provisioning and management tasks. OpenShift Data Science supports rapid experimentation with user-supplied data where the model outputs are hosted on OpenShift Dedicated for integration into a customer-defined, intelligent application or exported to self-maintained environments.

This solution represents a flexible, customizable alternative to prescriptive AI/ML suites from individual cloud providers. It provides open source tools and platform technology for collaborative creation of experimental models without infrastructure concerns or cloud-specific vendor lock-in. Teams can extend the base platform with partner tools to further enhance and expand their modeling capabilities. Models can be exported in a container-ready format for consistency deploying across hybrid cloud and edge environments.

Customize your environments with popular open source and commercial tools

Red Hat OpenShift Data Science provides a subset of tools found in the upstream Open Data Hub project (Table 1). Red Hat provides regular updates to the included open source tools through the managed cloud service, removing the burden of integration and testing. Additional commercial technology partner offerings can also be added from Red Hat Marketplace.



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



Table 1. Tools included in Red Hat OpenShift Data Science

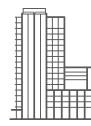
AI/ML modeling and visualization	Jupyter Hub with pre-defined notebook images, TensorFlow, PyTorch, Anaconda Commercial Edition (optional), IBM Watson Studio (optional)
Data engineering	Starburst Galaxy (optional)
Data streaming	Red Hat OpenShift Streams for Apache Kafka (optional)
Hardware acceleration	NVIDIA (with GPU operator)*, Intel OpenVINO and AI Toolkit (optional)
Model serving	OpenShift Source-to-Image tool, Red Hat OpenShift API Management (optional), Seldon Deploy (optional)

*coming soon

Support your AI/ML adoption with expert services

Meet challenges like the adoption of DevOps and ML best practices with Red Hat Consulting. See how an AI/ML residency with Red Hat consultants and technology experts through Red Hat Open Innovation Labs can help you succeed with Red Hat OpenShift Data Science. [Read the brief](#).

About Red Hat



Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.



facebook.com/redhatinc

@RedHat

linkedin.com/company/red-hat

North America

1888 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