



Advance your open source AI/ML initiatives with Red Hat

Migrate to Red Hat OpenShift AI to gain production-grade consistency, flexibility, and support across your organization


Contents

1 Advance your AI/ML initiatives with Red Hat OpenShift AI

2 Plan your migration

3 Is Red Hat OpenShift AI right for your organization?

4 Ready to begin your migration?



Advance your AI/ML initiatives with Red Hat OpenShift AI

Artificial intelligence (AI) technologies power innovative applications that speed processes, personalize experiences, and unlock insights from vast data sets. AI is the science and engineering of applications that can perform tasks that typically require human intelligence, like problem solving, learning, perception, and reasoning. A subset of AI, machine learning (ML) uses algorithms and statistical models trained on massive data sets to make predictions or decisions without being explicitly programmed.

AI and ML are key tools for enterprise technology leaders who want to generate business benefits. Many organizations start by creating their own AI platform using open source projects like Jupyter, PyTorch, and Kubeflow. While this approach lets teams work close to the innovation happening in communities like [Open Data Hub](#), it also requires larger teams and more effort to test, modify, and integrate these projects together.

Red Hat delivers foundational technology, proven expertise, and strategic partnerships to help you meet your AI and ML goals.

A natural evolution of the initiatives that created Open Data Hub, [Red Hat® OpenShift® AI](#) is an AI platform that provides tools for training, tuning, serving, monitoring, and managing AI/ML experiments and models on [Red Hat OpenShift](#). OpenShift AI gives data scientists and developers a powerful technology platform for gathering insights and building intelligent applications. Teams can quickly move from experiment to production in a collaborative, consistent environment.

With [multiple versions](#) available, OpenShift AI includes a core set of development and deployment features—like AI/ML libraries and frameworks, graphics processing unit (GPU) accelerator support, data science pipelines, and distributed workload capabilities—integrated with an [ecosystem of trusted AI tools](#). Data scientists can start with their choice of tools, create self-service development environments, and collaborate in real time, while developers can integrate container-ready models into AI-enabled applications with less effort. At the same time, both of these teams can deploy containerized models and applications on a unified, security-focused platform and quickly scale workloads to handle demands—including volume of data, duration of training run, size of model, and required acceleration—on-site, in the cloud, or at the edge.

Experience benefits across teams

Red Hat OpenShift AI delivers benefits across your organization. Here are some of the many reasons to migrate your open source AI/ML technologies, tools, and applications to OpenShift AI.



Production-grade support

Your OpenShift AI subscription provides access to online and phone support, so you can work with an expert to rapidly resolve issues. Speak with a product specialist from the start or find product documentation, life cycle and support information, troubleshooting tools, Red Hat Knowledgebase articles, account and subscription management tools, and security updates online via the [Red Hat Customer Portal](#).



Simplified upgrades

Red Hat provides a simplified upgrade path for OpenShift AI, so you don't need to keep track of the latest versions of open source projects or curate bug fixes in upstream communities. Take advantage of our expertise in creating enterprise open source solutions to save time and shift focus back to your business priorities. Managed cloud service users benefit from automatic upgrades with a 99.95% uptime service-level agreement (SLA), while self-managed users can better plan upgrades with [transparent life cycle policies](#).



Certified integrations

Customize your OpenShift AI deployment with your preferred open source or commercially available tools. OpenShift AI directly integrates a selection of open source tools. You can also incorporate certified partner products for data engineering, model development, hardware and software acceleration, and model management. These optional products and technologies can be accessed directly from the OpenShift AI dashboard and many also offer learning modules and tutorials. Or substitute your preferred upstream technologies.



Software supply chain security

Red Hat delivers enterprise-ready solutions—including OpenShift AI—via a documented, security-focused [software supply chain](#). A dedicated team of security engineers analyzes threats and vulnerabilities, offers proactive advice, and develops security patches for Red Hat products.

Plan your migration

Use these questions to assess your existing open source AI environment and needs, as well as future requirements, before migrating to Red Hat OpenShift AI.

Hardware requirements

- ▶ Do you know how to right-size and configure your hardware, including graphics processing units (GPUs) and other accelerators?

Management policies

- ▶ Will users be able to self-provision projects?
- ▶ Which existing use cases require a project?
- ▶ How will access to projects be managed?
- ▶ Which resources and tools will be pre-deployed for teams?

Processes and pipelines

- ▶ Do you have or need a custom pipeline for model training and data integration?
- ▶ Do your data scientists need custom notebooks with unique packages?
- ▶ Do you need to configure your ModelMesh with custom runtimes?

Integrated components and tools


- ▶ How will you integrate other data science tools into your new environment?
- ▶ Can you integrate your data storage systems with OpenShift AI?
- ▶ Do you have existing Red Hat OpenShift functionality that needs to be integrated?
- ▶ Are there critical integration points in existing systems that require additional tooling?

Solution deployment

- ▶ How will you deploy and manage OpenShift AI on multiple development, test, and production clusters?
- ▶ How will you deploy and configure other third-party tools in your environment?
- ▶ How will you manage the promotion of models across environments and clusters?

User education

- ▶ How will teams adopt and use OpenShift AI?
- ▶ Who will train data scientists and provide customizations to improve user experiences?



Get help from Red Hat experts

Red Hat Consulting offers several services to help you plan your migration:

- ▶ **Red Hat Consulting: Red Hat OpenShift AI Pilot** is an engagement that helps you identify the tools, integrations, and customizations needed in your AI/ML platform and deploy them in your environment.
- ▶ **Red Hat Consulting: MLOps Foundation** is a service that helps you advance your ML strategies, create reusable patterns for production-ready inference services, and automate the full ML model life cycle using cloud-native tooling and architectures.

[Learn more](#) about Red Hat Consulting.

Common migration steps

While no two migrations to Red Hat OpenShift AI are the same, this section provides an overview of the migration process with some of the challenges that you may encounter along the way.



NOTE: The easiest and safest way to install Red Hat OpenShift AI is on a new, unused cluster. If you need to reuse an existing cluster, make sure to fully remove any previously installed components before installing Red Hat OpenShift AI.

Tips for Open Data Hub users

Open Data Hub is a community project that provides a blueprint for integrating common open source AI/ML tools into an OpenShift environment. Because the Open Data Hub architecture is based on OpenShift, it is a popular choice for organizations that want to experiment with AI/ML before moving to a production-grade platform like Red Hat OpenShift AI.

There are several versions of Open Data Hub, and we call out additional version-specific tasks throughout this section for organizations that are migrating their Open Data Hub environments to Red Hat OpenShift AI.

STEP 1**Survey your current AI platform installation.**

- ▶ Collect all manifests and configuration files.
- ▶ Gather all relevant settings, including confidential data like Kubernetes Secrets and non-confidential data like ConfigMaps.
- ▶ Record the versions of all installed components like Grafana and Strimzi.
- ▶ Locate all notebook image versions, including your custom notebooks.
- ▶ Find all persistent volume claims (PVC) used by your AI platform, or other applications or users.
- ▶ Identify any other applications or tools installed in the same namespace as your AI platform.

Recommended for Open Data Hub releases prior to 1.4

- ▶ Plan migrations to alternative, community supported options for any deprecated tier 2 components.
- ▶ Prepare to migrate legacy JupyterHub notebooks to Kubeflow Notebook controller.
- ▶ Create plans consolidating multi-instance deployments to a centrally managed instance.

Recommended for Open Data Hub 1.x releases

- ▶ Evaluate all components deployed in the KfDef for migration to Red Hat OpenShift Dev Spaces.

Recommended for Open Data Hub 2.x releases

- ▶ Identify which components can remain in place and which must be removed and recreated post-migration.

STEP 2**Back up all data and configurations.**

Perform a full backup, including raw data, user data, configurations, and secrets prior to any change.

Actual back up procedures will vary depending on the type and volume of data, the allowed downtime, and your current backup and restore capabilities.

STEP 3 Prepare your current AI platform.

Upgrade your current AI platform to the latest stable version to ensure compatibility.

Recommended for Open Data Hub 1.4 and earlier

Migrate Kubernetes resources used by individual users and notebooks—including Secrets and ConfigMaps—to individual projects.

STEP 4 Uninstall existing AI platforms.

- ▶ Use automatic processes to remove as many components as possible.
- ▶ Manually remove any remaining artifacts and orphaned objects.
- ▶ Clean up old components as necessary.

STEP 5 Install Red Hat OpenShift AI.

Follow the [instructions](#) included with Red Hat OpenShift AI to complete the installation.

STEP 6 Restore all data and configurations.

Perform a full restore of previously saved raw data, user data, configuration, and secrets.

Is Red Hat OpenShift AI right for your organization?

Red Hat OpenShift AI could be the right choice for your project or organization. Here are some criteria shared by organizations that successfully use OpenShift AI for their AI/ML workloads.

- ▶ We want a data science platform that runs on Red Hat OpenShift.
- ▶ We want an open platform that provides foundational components we can build upon.
- ▶ We want to integrate other commercial and open source products into our environment simply and easily.
- ▶ We have compliance, data sovereignty, and privacy concerns that require us to deploy certain workloads on-site.
- ▶ We want to deploy workloads in hybrid, multicloud, or disconnected environments.
- ▶ We want to centrally manage users, data access, and resources for multiple teams.
- ▶ We support open source technologies.
- ▶ We want to simplify our workload deployment processes.
- ▶ We want to run AI workloads like large language models (LLMs) across multiple cluster resources with automatic scaling.
- ▶ We lack the time or expertise to integrate, maintain, and deliver a reliable platform based on multiple upstream open source projects.
- ▶ We want support for our AI platform from a dedicated team of experts.
- ▶ We want our data science teams to use the same practices and tools used by our other development teams.
- ▶ We already have Red Hat OpenShift skills, processes, and tools and want to use them for our AI initiatives and workloads.

If any of these apply to your organization, Red Hat OpenShift AI would make a great addition to your AI modernization plans.

Ready to begin your migration?

Red Hat OpenShift AI can help your organization rapidly build and advance AI-powered applications from development to production.

Migrating to a new AI platform can be challenging. **Red Hat Services** offers a variety of **consulting**, **training**, and **certification** services to help you build optimized environments to speed deployment of your AI/ML models and intelligent applications. In fact, organizations that engaged Red Hat to deploy Red Hat OpenShift improved infrastructure utilization and reduced their virtual machine footprint by 50%.¹ Red Hat's services and support have also helped these organizations realize a return on investment of 703%.¹



Red Hat Consulting

Red Hat Consulting offers hands-on mentoring to help you build skills, streamline processes, align teams, and ensure systems work together. Whether you're just starting or ready to develop strategies for deploying models to production, Red Hat Consulting can help.



Red Hat Training and Certification

Red Hat Training and Certification offers **targeted skills paths** for Red Hat OpenShift administrators and developers. Topics for hands-on training and certifications include OpenShift AI foundations and creating ML models.



Red Hat Technical Account Management

Red Hat Technical Account Managers (TAMs) are technical advisors that can help plan and deploy Red Hat software to support your AI/ML initiatives. In fact, including TAMs in key projects resulted in a 67% reduction in outages after 3 years.²

¹ Forrester Consulting study, commissioned by Red Hat. "The Total Economic Impact™ Of Red Hat Services And Support For OpenShift," May 2022. Results are for a composite organization representative of interviewed customers.

² Forrester Consulting study, commissioned by Red Hat. "The Total Economic Impact™ Of Red Hat Technical Account Managers," March 2023. Results are for a composite organization representative of interviewed customers.