

Accelerate your AI journey with Red Hat AI

Discover 4 ways organizations are using Red Hat solutions to achieve their AI goals faster.

At Red Hat, we offer solutions that help customers innovate so they can achieve their business goals. Red Hat® AI provides a portfolio of solutions that reduce costs and accelerate the development and deployment of artificial intelligence (AI) solutions across hybrid cloud environments. The portfolio includes Red Hat OpenShift® AI, a flexible, scalable AI and machine learning (ML) platform that enables enterprises to create and deliver AI-enabled applications at scale across hybrid cloud environments.

Here are 4 ways customers are embracing Red Hat AI:

1 Implementing scalability and flexibility to support generative AI

Delivering an excellent customer experience means providing responsive, highly personalized digital interactions. AI enables organizations to achieve this at scale. However, generative AI is a rapidly evolving technology with substantial data requirements. To keep up, organizations need flexible and scalable options to serve, run, monitor, and fine-tune their generative AI applications on premise, in the cloud, or within a hybrid environment.

Red Hat customers are using OpenShift AI to create knowledge-based applications and chatbot experiences for a wide range of use cases, including online AI-powered text-based sales assistants, customer service chatbots, and code generation to not only automate and supplement code development but also assist developers in code quality assurance.

How Red Hat helps: OpenShift AI offers a single combined AI platform for generative AI, predictive AI, and more traditional application development. It allows organizations to develop, deploy, and manage generative AI applications with greater flexibility and scalability.

With data augmentation through SDG (synthetic data generation), organizations can create additional high quality data that can be used to provide a more accurate model. OpenShift AI also integrates seamlessly with other leading solutions, providing organizations with access to large language models (LLMs) and accelerating the delivery of generative AI applications.

2 Enabling on-demand infrastructure with a click

To get the best results, data scientists and AI engineers need the freedom to experiment within the constraints of security and compliance requirements. They also need to be able to operate independently, rather than being reliant on other teams to provide the infrastructure and resources they require.

How Red Hat helps: Red Hat customers use OpenShift AI to create self-service environments for data scientists and AI researchers, so they can quickly build experimental models with Jupyter notebooks, TensorFlow, PyTorch, and NVIDIA Graphic Processing Unit (GPU) support without worrying about the underlying infrastructure. The secure platform simplifies the building, testing, serving, and monitoring of AI models, allowing organizations to standardize AI technologies and processes across teams.

OpenShift AI simplifies the reuse and adaptation of models, enhancing productivity and agility. By empowering employees with controlled autonomy, organizations can accelerate innovation and bring new ideas to market swiftly while maintaining security and compliance.

3 Increasing efficiency in developing predictive AI solutions

AI has the power to handle complex analytics problems and enable an automated outcome, without the need for time-consuming manual intervention and reducing the risk of human error. To support data-driven or automated decision-making, organizations need to be able to efficiently create predictive AI solutions.

For example, Red Hat customers are using OpenShift AI to rapidly develop AI-powered solutions that:

- ▶ Automatically classify and route service desk queries to the right team
- ▶ Quantify financial risk with complex analytical models
- ▶ Qualify and prioritize sales leads
- ▶ Identify healthcare patients at risk of certain conditions
- ▶ Generate next-best action suggestions for sales teams.

How Red Hat helps: With OpenShift AI, organizations can develop predictive AI models with greater efficiency. For example, the solution offers common open source tooling enabling teams to standardize, while providing the flexibility for users to add their own custom model images and model serving runtimes to meet organizational needs.

By using OpenShift AI to develop predictive AI applications, organizations can also boost their competitive advantage and deliver significant business benefits, including enhanced employee productivity, customer service quality, and reduced cost and risk.

[Learn more about Red Hat OpenShift AI](#)

4 Establishing an MLOps approach

To develop AI applications effectively, organizations need well-defined processes and practices. Many are adopting a Machine Learning Operations (MLOps) approach - a set of workflow practices inspired by DevOps designed to streamline the development and deployment of machine learning models. MLOps processes help teams work collaboratively to ensure that machine learning models remain accurate and up to date and that model development environments are standardized and automated. To successfully implement MLOps, organizations need the right platform.

How Red Hat helps: Red Hat OpenShift AI is used by organizations to integrate AI into their existing DevOps workflows. It is optimized for data scientists and developers of AI-enabled applications and provides a fully supported environment to establish MLOps best practices, where data scientists and developers can rapidly train, deploy, and monitor ML workloads and models on premise and in public clouds.

The AI models built or tuned in OpenShift AI are easily portable, allowing teams to deploy them in on premise, public cloud, or edge environments. It offers a unified platform for managing the entire lifecycle of code, applications, and AI models, from planning and training to tuning and deploying.

[Discover how Red Hat can help you harness the power of AI](#)



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.

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

f facebook.com/redhatinc
X twitter.com/RedHat
in linkedin.com/company/red-hat