

Accelerating AI into production with Red Hat Services

Artificial intelligence (AI) is powering a new era of productivity across industries, and companies need to find the right use cases to give them a competitive edge. With Red Hat® OpenShift® AI and Red Hat Enterprise Linux® AI, companies can more easily develop, test, and run AI-enabled applications across hybrid cloud environments. However, navigating this emerging technology in the most cost-effective and impactful way can be challenging. Here are 4 lessons from customers who all worked with Red Hat Services to accelerate their AI strategies.

1 Automate to accelerate AI adoption

A common pain point experienced by customers when it comes to deploying AI is engaging and empowering the right people to collaborate and innovate at scale. Some customers have more than 100 data scientists developing and testing AI and machine learning (ML) models before passing them to IT teams for deployment. The user experience is often impacted by a lack of easy access to compute resources and—on top of this—workstations that are complex to use. This can be frustrating, unproductive, and slow for everyone.

To improve this, customers are adopting a self-service approach for data scientists and developers, using existing technologies and skills as much as possible. Red Hat Enterprise Linux AI offers a low barrier of entry for AI experimentation while OpenShift AI allows self-service deployment to automate data science pipelines and ease resource consumption. This empowers data scientists to take greater ownership over their AI projects and increases efficiency, leaving teams more time to focus on their main responsibilities.

How Red Hat helps: Red Hat Consulting embeds experts within customer IT teams to provide leadership, mentorship, and advice on how to maximize the effectiveness of existing skills and architecture. For one customer, Red Hat consultants recommended that developers use the Java Quarkus programming language they were already familiar with, rather than learning Python.

2 Eliminate technical barriers to reduce time to market

Organizations need the right AI foundation in place to reduce time to market and scale up development and deployment. Customers might invest in the latest technologies but often need support and advice on how to best use them for their specific needs.

Red Hat Enterprise Linux AI provides a stable operating system with all the functionality companies need to experiment with AI. It also makes AI more cost-effective and more accessible with preconfigured AI packages, smooth integration with DevOps tools, and the high performance levels that AI workloads require. OpenShift AI, with features like smart scaling to save time and costs, can be integrated with Red Hat Ansible® Automation Platform to speed up model development and testing.

With these technologies, the time to push images to production can be reduced—in the case of one particular customer, from 2 months to a few days. OpenShift AI can be combined with Red Hat Developer Hub to support collaboration throughout the development process.

How Red Hat helps: Red Hat Services simplifies installations with expert advice. They build strong relationships to understand a customer's architecture and how best to optimize it, guiding future rollouts through workshops and proofs of concept.

[InstructLab](#)—an open source project included with Red Hat Enterprise Linux AI that improves the alignment of large language models (LLMs)—is also empowering organizations with minimal ML experience to start working with AI.

3 Unlock the full value of AI faster

Knowing the right processes to automate with AI can be the difference between making real productivity gains and merely following the AI trend.

For example, Red Hat AI has been used to augment an app to help staff search knowledge articles from Slack to find fast resolutions to issues without raising a support ticket. This was achieved by deploying retrieval augmented generation (RAG) and generative artificial intelligence (gen AI) in Red Hat OpenShift, whereby an LLM generates synthetic data based on the information it is given from data storage solutions like Red Hat Data Services.

Another use case that was invaluable to one company was improving the reliability of predictive analytics with AI to make numerous departments more productive while delivering better employee and customer experiences.

How Red Hat helps: By working closely with development teams, Red Hat Services has combined insights into customers’ goals and their environments with AI expertise to offer trusted advice on areas that will deliver results.

Through proofs of concept, Red Hat Services has empowered teams to continue exploring OpenShift AI, and tight feedback loops between Red Hat and its customers have resulted in the sharing of mutually beneficial insights, which help to continuously improve the use of AI.

4 Collaborate with trusted experts

A lack of shared expertise and teams working in isolation can hinder AI projects. This is a common challenge in every customer project and, once addressed, a culture of collaboration and knowledge-sharing can speed up development. It is beneficial for customers to hear different perspectives, learn best practices, and challenge their preconceptions about AI.

How Red Hat helps: Red Hat Services has the flexibility to tailor support for different needs.

- ▶ Helping teams to define clear roles and responsibilities through support and mentorship
- ▶ Running workshops to define project goals and outcomes, mapping a pathway to achieving them with the aid of a prepackaged Open AI/ML platform framework to guide a customer’s transformation
- ▶ Helping customers build their own AI centers of excellence (CoE)

Red Hat Services can also provide one-to-one mentorship over a shared screen with tailored training and enablement around topics such as distributed workloads, Kubernetes containers, GitHub, and virtual LLMs.

Red Hat’s ongoing support, including daily standups, documentation, and diagrams, has helped customers learn sustainable skills and feel confident leading their in-house AI projects in the future.

Learn more about how [Red Hat AI](#) can work for you

See how [Red Hat Services](#) can help you make AI work better



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.

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

redhat.com

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