5 ways Red Hat and AWS build AI value

Navigating the complexities of artificial intelligence (AI) deployments and integrations can be a daunting task. Red Hat® OpenShift® AI combined with Red Hat OpenShift Service on AWS offers a comprehensive, managed application platform in the cloud that can help businesses accelerate their AI strategy. Explore 5 ways this combined solution can help data scientists, IT leaders, and developers operationalize AI in their organization.

1 Operationalize AI in any environment

Red Hat OpenShift Service on AWS provides a comprehensive, managed platform to adopt AI across the hybrid cloud and to the edge of the network.

With built-in developer tools, managed maintenance, and an integrated security approach, teams can focus on building AI-enabled applications rather than managing infrastructure, so they can:

- **Streamline AI adoption.** Take AI models from experimentation to deployment.
- **Simplify workflows.** Train and deploy AI-enabled applications efficiently.
- **Use a unified platform.** Reduce complexity and free up resources for innovation.

2 Build, train, and deploy AI models efficiently

OpenShift AI on Red Hat OpenShift Service on AWS simplifies the lifecycle management of AI models.

With access to the right tools for building, training, serving, and deploying AI models in a managed cloud environment, this solution reduces the barriers to entry for data scientists, making it easier to work with AI models.

A managed cloud environment lowers barriers for data scientists, equipping them with the tools needed to build, train, serve, and deploy AI models efficiently.

Data scientists gain self-service capabilities that include:

- **Pre-configured tools and libraries.** Streamline development with ready-to-use resources.
- **Scalable infrastructure.** Handle resource-intensive tasks like on-demand provisioning and training.
- **Collaboration and version control.** Support version control and team sharing for notebooks and models.
- **Model registry.** Centrally manage and track AI models with a KubeFlow community preview.
- **Model reliability tools.** Detect model drift and bias to boost accuracy and fairness.

This streamlined approach enables quick experimentation while maintaining a security focus and compliance standards.

3 Achieve consistency with OpenShift AI on premise or on AWS

While hybrid cloud environments offer flexibility, scalability, and enhanced security, they can also increase complexity which can threaten consistency.

OpenShift AI on Red Hat OpenShift Service on AWS provides operational consistency across on-premise and cloud environments.

Organizations can train AI models where their data resides—whether on premise or in the cloud–maintaining data privacy while allowing for scalability.

- **Data-local training.** Train AI models where data resides while maintaining privacy.

- **Consistent platform.** Develop and deliver AI-enabled applications alongside other workloads.
- **Optimized costs.** Freely move workloads between environments to maximize efficiency.

With consistent workflows and reduced integration headaches, teams can scale AI solutions without compromise.

4 Achieve tighter integration between DevOps and MLOps

Effective AI and ML implementation depends on collaboration.

OpenShift AI on Red Hat OpenShift Service on AWS bridges the gap between DevOps and MLOps by providing a consistent user experience for data scientists, developers, and IT Operations by providing:

- **Guardrails and approvals.** Track, test, and validate changes to reduce risk.
- **Streamlined continuous integration and continuous delivery (CI/CD) pipelines.** Automate workflows from model development to production inference.
- **Effective collaboration.** Integrate data science pipelines with DevOps workflows to accelerate delivery.

By aligning teams and automating processes, organizations can accelerate deployment and increase reliability across their AI-enabled applications.

5 Scale AI workloads, even in data-intensive scenarios

Data-intensive workloads demand infrastructure that can scale reliably.

OpenShift AI on Red Hat OpenShift Service on AWS connects organizations with the AWS's GPU capabilities and global cloud infrastructure. This enhanced computing power allows businesses to scale their AI applications up or down as needed using only the GPU power they need depending on what their business demands.

By deploying trained models in the cloud, organizations gain access to elastic resources that handle spikes in demand without compromising performance.

This scalability, combined with cost optimization and enhanced efficiency, makes the platform ideal for organizations looking to advance their AI capabilities.

Learn more

Discover how Red Hat OpenShift AI on Red Hat OpenShift Service on AWS can help your organization operationalize AI and accelerate your AI strategy.

Try Red Hat OpenShift Service on AWS. Sign up for a hands-on experience.

Read more about Red Hat OpenShift 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	Europe, Middle East, and Africa	Asia Pacific	Latin America
1 888 REDHAT1 www.redhat.com	00800 7334 2835 europe@redhat.com	+65 6490 4200 apac@redhat.com	+54 11 4329 7300 info-latam@redhat.com

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

redhat.com