

Red Hat Consulting: AI Accelerator

Establish a ML platform and adopt new AI capabilities

"Red Hat is very important for the implementation of [natural language processing] NLP at Banco Galicia, since it provides us with the technology and architecture, and allows us to understand everything related to Red Hat OpenShift. We also started designing an architecture directly in the cloud."

—

Matias Loruso

Solutions Architect, Banco Galicia

Red Hat OpenShift provides an on-premise platform with a highly available and reliable foundation, and AI workloads are containerized to optimize inference processing.

[Learn more](#) about how Banco Galicia worked with Red Hat Consulting to design an intelligent AI-based NLP platform.

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

Leaders in various industries are turning to the power of artificial intelligence and machine learning (AI/ML) to achieve desired business outcomes. Red Hat has the leading open hybrid cloud platform to help organizations train, deploy, and monitor AI/ML workloads and models on-premise and in a public cloud environment.

Deploying these technologies effectively can seem daunting, which is why it is important for organizations to use a platform best suited to their needs. Not using the right platform can lead to obstacles—including insufficient access to compute, underused resources, and a lack of knowledge sharing and various versions of libraries—that make it difficult to collaborate and produce quality models at an accelerated pace.

To help address these challenges, Red Hat created [Red Hat® OpenShift® AI](#) to provide a unified platform of open source tooling for building, deploying, and managing ML models. Following the release of Red Hat OpenShift AI, many customers aimed to expand their data science capabilities but needed help to do so. As a result, Red Hat Consulting introduced the AI Accelerator engagement to help customers identify the tools, integrations, and customizations they need in an AI/ML platform.

What is Red Hat Consulting: AI Accelerator?

Overview

Red Hat Consulting developed the AI Accelerator engagement to focus on helping customers start up their Red Hat OpenShift AI journey and integrate it with their existing enterprise services. With a central platform in place, users will have access to standardized libraries and tools, increased compute, ability to collaborate more effectively, and improved onboarding of data scientists and other users.

Through this engagement, Red Hat experts will join an organization's teams to evaluate the current environment and approach and identify future requirements. They will help use a GitOps approach to deploy and manage Red Hat OpenShift AI and integrate it with other data science tools in customers' environments to get the most out of the technology.

Red Hat's tailored approach

Architecture sessions

Red Hat consultants will provide hands-on experience with a focus on deploying to multiple clusters—such as a development, test, and production cluster—and can extend to end users and multiple data science teams.

First, we will get started with architecture sessions to help identify any critical integration points in existing systems and additional tooling needed to meet each customer's unique requirements. Customers can expect Red Hat Consulting to aid in the selection of tools required for their AI/ML platform. Red Hat Consulting will also perform an architecture review to understand their environment and determine how Red Hat OpenShift AI will be used and adopted.

Team collaboration

Following these architecture sessions, Red Hat consultants will collaborate with teams to assist in performing the deployment and configuration of third-party tooling components defined in the architecture design sessions. In this phase, Red Hat Consulting may engage in the following:

- ▶ Red Hat consultants work with you to appropriately size configurations, integrate with data systems, and configure hardware accelerators, such as GPUs.
- ▶ Assistance in the deployment, configuration, and integration of third-party tooling.
- ▶ If necessary, integrating existing OpenShift functionality (e.g., serverless, GitOps, pipelines) to support platform requirements.

Customizations

Red Hat acknowledges every customer has different business objectives, so we customize components of the engagement, based on the architecture design sessions, with a focus on:

- ▶ Developing custom pipelines for model training and data integration.
- ▶ Creating pipelines to deliver custom notebooks.
- ▶ Configuring model mesh with custom runtimes.

Red Hat prioritizes helping to bridge the cluster administration, infrastructure teams and data scientists as part of a core foundation of this engagement. Red Hat will work directly with both organizations to help support teams for success to maintain the current state and lay the groundwork for future requirements.

After establishing the AI Accelerator, Red Hat will work directly with data scientists to help onboard them and provide customizations to help improve their user experience.

Considerations

The engagement is predicated on having an appropriately sized OpenShift cluster. Red Hat Services can both help scope and deliver such a cluster as a prerequisite. The AI Accelerator does not require you to have any functioning ML models for this engagement, and Red Hat is able to meet you wherever your team is on your data science journey.

For customers looking to move beyond model experimentation and develop strategies for deploying models to production, [Red Hat Consulting: MLOps Foundation](#) is a great follow-up engagement to the AI Accelerator engagement.

Getting started with Red Hat

Red Hat strives to help our customers to not only deliver their first ML model, but the foundation for all of their ML systems. By creating a reusable pattern to train and deploy an ML model as a production-ready solution, Red Hat seeks to help our customers succeed in delivering future ML projects.

Using a GitOps approach, Red Hat Consulting can help to provision data science projects, accelerate the building, training, testing, deploying and monitoring of ML models and consistently promote models to production across hybrid cloud and edge environments—from cloud based clusters to fully disconnected on-premise environments.

The Red Hat difference

Red Hat Services consist of Red Hat Consulting, Red Hat Training and Certification, and Red Hat Technical Account Management (TAM) and provide customers with comprehensive support and guidance for building modern cloud applications.

Consulting: With hands-on mentoring, Red Hat consultants build skills and encourage independence while also streamlining processes, aligning teams, and ensuring systems and applications work together.

Training: This helps to close skills gaps and hone teams' Red Hat product expertise by developing role-based, hands-on knowledge in emerging and foundational open source technologies. Check out our robust [OpenShift AI curriculum](#).

Red Hat TAMs: TAMs partner with organizations to resolve potential problems before they occur, minimizing disruption and alleviating time to focus on key business challenges.

Interested in AI/ML services offered by Red Hat? Speak to your Red Hat Account team, or talk to a [consultant](#).



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.

 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