

# Red Hat Consulting: AI Platform Foundation

**Establish a ML  
platform and adopt  
new AI capabilities**

"Red Hat Consulting's technical expertise and communication skills were evident and instrumental in our success," said Yavuz. "The team provided mentorship and worked with us on a one-to-one basis, helping us adopt standardization and best practices, and make full use of OpenShift AI's capabilities."

---

**Emre Yavuz**  
Head of Data and AI,  
Turkish Airlines

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

[Learn more](#) about how Turkish Airlines worked with Red Hat Consulting to implement OpenShift AI and develop a culture of experimentation.

 [facebook.com/redhatinc](https://facebook.com/redhatinc)  
 [@RedHat](https://twitter.com/RedHat)  
 [linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

## Accelerate a path to production for gen AI with Red Hat Consulting

Leaders across industries are turning to generative AI (gen AI) to reinvent processes and create new customer experiences. Red Hat has one of the leading open hybrid cloud platforms, helping organizations develop, deploy, and manage cutting-edge AI applications from the datacenter to a cloud environment.

However, bringing AI applications from pilot to production presents a unique set of challenges, including restrictive graphics processing unit (GPU) costs, complex model serving requirements, and inconsistent development environments that hinder repeatable fine-tuning and deployment of large language models (LLMs).

To help address these challenges, Red Hat created Red Hat® AI, which includes [Red Hat AI Inference Server](#), [Red Hat OpenShift® AI](#) and [Red Hat Enterprise Linux® AI](#) to provide a flexible, end-to-end platform of open source tooling for the entire LLMOps lifecycle—from data preparation and model fine-tuning to scalable inference and monitoring. To advance this journey, Red Hat Consulting: AI Platform Foundation helps customers establish LLMOps strategies and the AI platform needed to rapidly operationalize generative AI applications and achieve tangible business results.

## What is Red Hat Consulting: AI Platform Foundation?

Red Hat Consulting developed the AI Platform Foundation engagement to help customers rapidly deploy and adopt their Red Hat AI investment while advancing their AI practices. 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, AI engineers, and other users.

Through this engagement, Red Hat experts will join an organization's teams to evaluate the current environment, approach, and identify future requirements. A GitOps approach is used to deploy and manage this Red Hat AI platform and integrate it with other tools in customers' environment to implement a robust end-to-end platform.

## 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 such as data scientists, AI engineers, and application developers.

We will get started with architecture sessions to 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 assist in the selection of tools and models required for their AI platform. Red Hat Consulting will also perform an architecture review to understand their environment and determine how Red Hat 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, a Red Hat Consultant may engage in the following:

- ▶ Collaborate to appropriately size configurations, integrate with data systems, and configure hardware accelerators, such as GPUs.
- ▶ Assist in the deployment, configuration, and integration of third-party tooling.
- ▶ Integrate existing platform functionality (e.g., serverless, GitOps, pipelines) to support platform requirements.

## Customizations

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

- ▶ Developing custom pipelines for model training, model deployment, and data integration.
- ▶ Creating pipelines to deliver custom notebooks.
- ▶ Establishing Models-as-a-Service (MaaS).
- ▶ Creating pipelines to benchmark and evaluate a model's accuracy.

Red Hat helps to bridge the cluster administration, infrastructure teams, and data scientists as a priority and core foundation of this engagement. Red Hat works directly with both organizations to provide team support to maintain the current state and lay the groundwork for future requirements.

After establishing the AI Platform Foundation, Red Hat will work directly with data scientists and AI engineers to get them onboarded and provide customizations to help improve their user experience.

## Considerations

Red Hat is able to meet you wherever your team is on your gen AI journey. For customers looking to operate beyond model experimentation and develop strategies for deploying models to production, [Red Hat Consulting: Operationalizing AI](#) is a great follow-up engagement to the AI Platform Foundation.

## 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 curriculum](#) around Red Hat AI.

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

## Getting started with Red Hat

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

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

Interested in AI 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](https://facebook.com/redhatinc)  
 @RedHat  
 [linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

**North America**  
1888 REDHAT1  
[www.redhat.com](http://www.redhat.com)

**Europe, Middle East,  
and Africa**  
00800 7334 2835  
[europe@redhat.com](mailto:europe@redhat.com)

**Asia Pacific**  
+65 6490 4200  
[apac@redhat.com](mailto:apac@redhat.com)

**Latin America**  
+54 11 4329 7300  
[info-latam@redhat.com](mailto:info-latam@redhat.com)