

New release: Unlocking AI with Red Hat for Partners

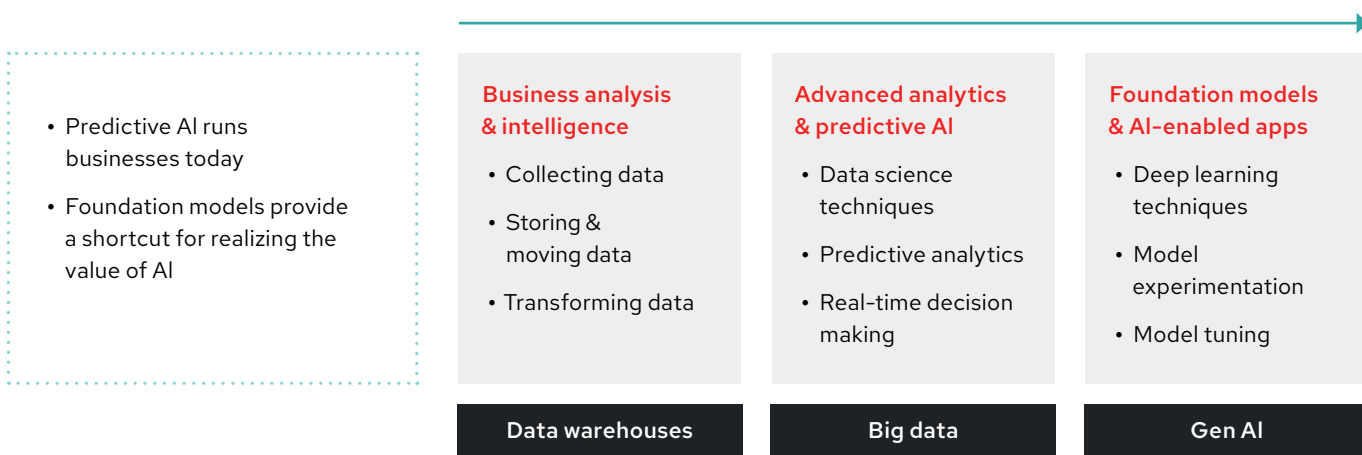
Red Hat is offering EMEA partners and distributors a structured, time-limited initiative to help teams focus on a specific topic with a target customer or segment. The goal is to guide participants from initiating conversations and architecture discussions to technical demos, preparing a PoC, and winning technical decision points (TDP).

This initiative is designed to provide our partner ecosystem with a deep understanding of Red Hat® AI – from both sales and tech sales perspectives – covering the first meeting, assessing the customer's AI maturity, identifying opportunities, and pitching Red Hat AI's value proposition, including its powerful **AI model inference capabilities** powered by **Red Hat Inference Server** and **OpenShift AI (RHOAI)**.



AI has undergone significant evolution

The evolution of AI: from business intelligence to generative AI



Growing demand for AI solutions and services

The worldwide AI software market will grow to nearly \$790 billion by 2026 (5 yr CAGR 18%)¹

51%

of organizations cite 'lack of MLOps tools' as a challenge?²

65%

of organizations are currently investing in generative AI³

¹ [IDC - Worldwide AI Software Forecast, 2023-2027](#)

² [IDC Market Analysis Perspective: Worldwide Artificial Intelligence Software, 2023](#)

³ [IDC's AI Strategies View 2022, CY22](#)

The new initiative: Unlocking AI with Red Hat for Partners

To help you guide customers through one of the most transformative periods in the history of technology, we've designed a new initiative that equips your sales and technical teams with the confidence and skills to position Red Hat's AI portfolio – including Red Hat's **production-ready inference server solution** – effectively and competitively.



Step 1: Master your first meeting – the Red Hat AI opportunity

Audience: Sales and Technical Sales

Duration: 1 hour

Description: Learn how to have a constructive, non-technical conversation with customers about Red Hat's AI capabilities, including how Red Hat enables reliable, scalable model inference and deployment in production.

Topics include:

- What problems are we solving for our customers?
- **The power of inference:** Why Red Hat?
- **Business use cases:** Hear from 3 Red Hat partners how they leverage Red Hat technology to address customer AI business challenges
- Questions to ask, objection handling, and customer case studies

Leave this session confident in qualifying AI opportunities with a focus on inference, scalability, and enterprise-readiness.

[> Discover the available dates and sign up](#)



Step 2: Sales & Technical Training – skill up on Red Hat AI

Audience: Sales and Technical Sales

Duration: 3 hours

Description: This session covers:

- AI fundamentals & market trends
- Red Hat's AI product portfolio
- Positioning Red Hat AI to win
- Red Hat OpenShift® AI (RHOAI) technical overview
- **Red Hat Inference Server** – efficient AI inferencing at scale
- Live demo strategy and customer success examples
- Handling technical objections
- Structuring PoCs

Learn how to position Red Hat's AI stack, including **Red Hat Inference Server**, as a secure, scalable foundation for running LLMs and other AI models at the edge, in data centers, or in hybrid cloud environments.

[> Discover the available dates and sign up](#)



Step 3: Architecture workshop – Red Hat Inference Server

- Audience:** Technical
- Duration:** 4 hours
- Description:** A technical deep dive into the components and architecture of Red Hat Inference Server & Red Hat OpenShift® AI, with hands-on labs including:
- Working with RHOAI pipelines and model lifecycle
 - Integrating predictive and generative AI into intelligent apps
 - **Deploying AI models using Red Hat's inference server**
 - Enhancing LLMs with Retrieval-Augmented Generation (RAG)

Understand how Red Hat enables **cost-effective inference at scale** for production-grade AI deployments.

[> Discover the available dates and sign up](#)



Step 4: Red Hat AI partner-led workshop

This flexible format can be tailored:

One-to-many approach

- For hosting joint technical events for multiple customers.

One-to-one approach

- Focused on advancing or closing a customer opportunity.
 - Use the workshop to **demonstrate real-world AI inferencing** with Red Hat tools in a lab environment.



Step 5: AI PoC

Coming soon!

Pilot Red Hat AI technologies in customer environments, focusing on:

- Integrating with customer data pipelines
- Deploying LLMs using Red Hat's AI inference stack
- Demonstrating performance, observability, and support benefits

Leverage Red Hat's enterprise-grade support and NVIDIA collaboration to deliver **secure, optimized inference** with certified model containers.

Why Red Hat for AI?

Red Hat's AI platform gives customers:

- **Enterprise-ready inference with Red Hat Inference Server**, built on **vLLM technology**
- A full stack from data science to deployment with OpenShift AI
- Flexibility to run AI workloads anywhere: cloud, edge, or hybrid
- Support for **open source model serving** frameworks like **ONNX Runtime, Triton, and Hugging Face Text Generation Inference (TGI)**



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 x.com/RedHat
in linkedin.com/company/red-hat

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

redhat.com

Copyright © 2025 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Red Hat logo, and JBoss are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.