EMEA AI: Unlocking AI with Red Hat for Partners

The format of an Enablement Journey is a phased and time-boxed approach to enable teams to address a focus topic with a target customer or segment.

The aim is to prepare participants to start a first conversation, continue with an architecture discussion, move on with a technical demonstration based on the same storylines and flow, all the way up to architecting a successful Demo, preparing a proper POC with your customers, and supporting to win the Technical Decision Points (TDP).

This specific Enablement Journey will give Partners and internal sales teams (eg PODs, Territory) an in-depth understanding of Red Hat AI from a Sales and TechSales perspective. Starting from a first meeting, identifying the customer's current AI experience and know-how, to potential opportunities, and pitching the value proposition of Red Hat AI.



EMEA AI: Unlocking AI with Red Hat for Partners

Self-Paced Training

- **Step 1:** Master your first meeting
- **Step 2:** Skill Up on Red Hat AI
- **Step 3:** RHOAI Architecture Workshop
- **Step 4:** AI POC Enablement "Technical Deep Dive - Winning AI POCs"

Virtual Instructor-Led Training

- **Step 5:** Red Hat AI Factory with NVIDIA (vILT) | Check upcoming dates here

EMEA AI: Unlocking AI with Red Hat for Partners

Step 1: Master Your First Meeting | Self-Paced Training

Audience: Sales and Technical Sales (highly recommended for everyone on the team)

Duration: 1 hour

Description: Webinar discussing the value proposition of Red Hat AI, questions to ask to help identify opportunities, and how to secure a second meeting.

- Identify Customers' key challenges
- Recap on terminology: Gen AI, LLM, RAG, etc.
- What is Red Hat's AI solution, Why Red Hat for AI?
- Engaging with the customer
- What's in it for Partners

Step 2: Skill Up on Red Hat AI | Self-Paced Training

Audience: Technical Sales (recommended for technology interested salespeople)

Duration: 3 hours

Description: Training on basic Sales and Tech Sales knowledge of RH AI. Attendees will be able to describe the features and components of RHEL AI and RHOAI and the benefits they provide to customers. They will have a deeper under standing of the different demo assets available and how to deliver value propositions to different customer personas (in episodic fashion). Also looking at real customer wins and the POC opportunity for the coming year.

- AI Landscape
- Red Hat AI products overview
- Positioning Red Hat AI to win
- RHOAI Technical Overview
- RHEL AI Introduction
- Demoing Red Hat AI
- Red Hat AI: Customer Successes
- POC Opportunities

Step 3: RHOAI Architecture Workshop | Self-Paced Training

Audience: Technical Sales and Delivery

Duration: 3-4 hours

Description: Workshop for Tech Sales and Delivery providing knowledge and giving a technical deep dive into the various components and frameworks that make up RHOAI, and get your hands on the product to work through an industry-specific use case. Using Generative & Predictive AI, to increase efficiency and reduce human error in the organisation.

- Hands-on Lab with RHOAI
- Build, train, and serve Gen AI and predictive AI models
- Using Data Science Pipelines
- Integrating models and smart application
- RAG enhancement of LLMs

Step 4: AI POC Enablement "Technical Deep Dive - Winning AI POCs" | Self-Paced Training

Audience: Technical Sales

Duration: 3 hours

Description: Workshop going deeper in the technical aspects of RHAIIS (based on ETX POC training) Provides attendees with fundamental knowledge about identifying and approaching POC opportunities. The various considerations, along with tools needed for evaluating and benchmarking POC performance. Attendees will get a deeper understanding of RHAIIS, the benefits it provides as well as how it works under the hood.
We will look at a number of complex concepts:

- Continuous Batching
- Paged Attention
- Speculative Decoding
- Model Compression

Step 5: Red Hat AI Factory with NVIDIA (vILT) | Virtual Instructor-Led Training

Audience: Sales and Technical Sales

Duration: 2 hours

Description: Webinar discussing the value proposition of Red Hat AI Factory with NVIDIA, this session will cover everything from the problem statement and solution architecture, to How to Buy and Frequently asked questions about the joint offering.

- The shift to production AI
- Red Hat AI Factory with NVIDIA
- Platform Architecture
- Example Use Cases
- FAQs

> Check upcoming dates here



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.

facebook.com/redhatinc
x.com/RedHat
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 © 2026 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.