

Red Hat Vision for AI and Open Source



Chris Wright

Chief Technology Officer and Senior Vice President,
Global Engineering
Red Hat



Where do we go next?



A photograph of two people, a woman on the left and a man on the right, sitting at a desk in a dimly lit room. They are looking at a laptop screen which displays code. In the background, there are several other computer monitors, some showing code and others showing data or charts. The overall atmosphere is that of a modern, tech-oriented workspace.

**What's challenging your
IT operations?**


A person in a dark blue uniform is standing in a server room aisle, leaning over a blue rolling cart. The room is filled with rows of server racks on both sides, with many small lights visible. The floor has a grid pattern. The overall lighting is dim and blue-toned.

**What's driving
your innovation?**


Red Hat enables open hybrid cloud



Any cloud



Secure
automated
infrastructure



Any
application



Team
collaboration

Open hybrid cloud is evolving

Mission-critical
automation

Developer
productivity

Trusted
security

Edge

AI

Open Hybrid Cloud



open source and AI






**Generative
AI era**

You are here

Every previous era has grown and expanded with "open"



We see a future where open source plays a huge role in AI



Open source models






Red Hat

IBM Research

open sourcing the Granite family
of language and code models





Big results from **smaller models**




Big results using **data** provided by **experts across your organization**



And big results for much **lower cost**

Code assistance model performance






InstructLab

contribute to and train
Large Language Models

The background is a solid blue color with a subtle gradient. There are several 3D-rendered spheres: a large teal sphere in the top right, a small red sphere to its right, a large red sphere in the bottom left, and a small light pink sphere partially visible next to it.

Red Hat allows your
business to **own your AI.**

The background is a dark blue gradient. In the top right, there is a large teal sphere and a smaller red sphere. In the bottom left, there is a large red sphere and a smaller pinkish sphere.

Red Hat is enabling your
business to run **your AI**
in your way.

Red Hat's AI portfolio

Trust

Choice

Consistency

AI models
RHEL AI

AI platforms
OpenShift AI

AI-enabled portfolio
Red Hat Lightspeed

AI workloads

Open Hybrid Cloud platforms

Partner ecosystem

Hardware | Accelerators | Models | Delivery

AI platforms



Red Hat Enterprise Linux AI



Granite language and
code models

InstructLab

open source AI tooling

AI-optimized Linux



AI platforms



Red Hat OpenShift AI



Can be used for both
GenAI and Predictive AI
use cases

Helps accelerate getting
AI apps and models
into production

Provides a trusted platform, choice
of AI models and hardware, and
consistent operations



AI platforms



Red Hat Ansible Automation Platform



Automation is a key enabler for AI-driven workflows.

The content and structure of Ansible playbooks are a fertile and safe AI training ground.

Automation delivers the savings needed to support your AI transformation.



Red Hat Lightspeed



Neural Magic + Red Hat



Neural Magic's expertise in software and algorithms that accelerate generative AI (gen AI) inference workloads.



nm-vllm



LLM Compressor

DeepSparse



Red Hat

Red Hat's vision of high-performing, functional AI workloads that directly map to customer-specific use cases and data, anywhere and everywhere across the hybrid cloud.



Red Hat Enterprise Linux AI



InstructLab



Red Hat OpenShift AI



The potential to supercharge LLM deployments anywhere and everywhere across the hybrid cloud by providing a ready-made, highly-optimized, open inference stack.



Open source meets AI.

AI meets **Open Hybrid Cloud.**