

# Red Hat

## AI offering



TD SYNnex AI POSITIONING

Red Hat AI STRATEGY

Red Hat AI SOLUTIONS

EDUCATION - PROFICIENCY

LAB & ASSETS

Red Hat MARKETING INITIATIVES

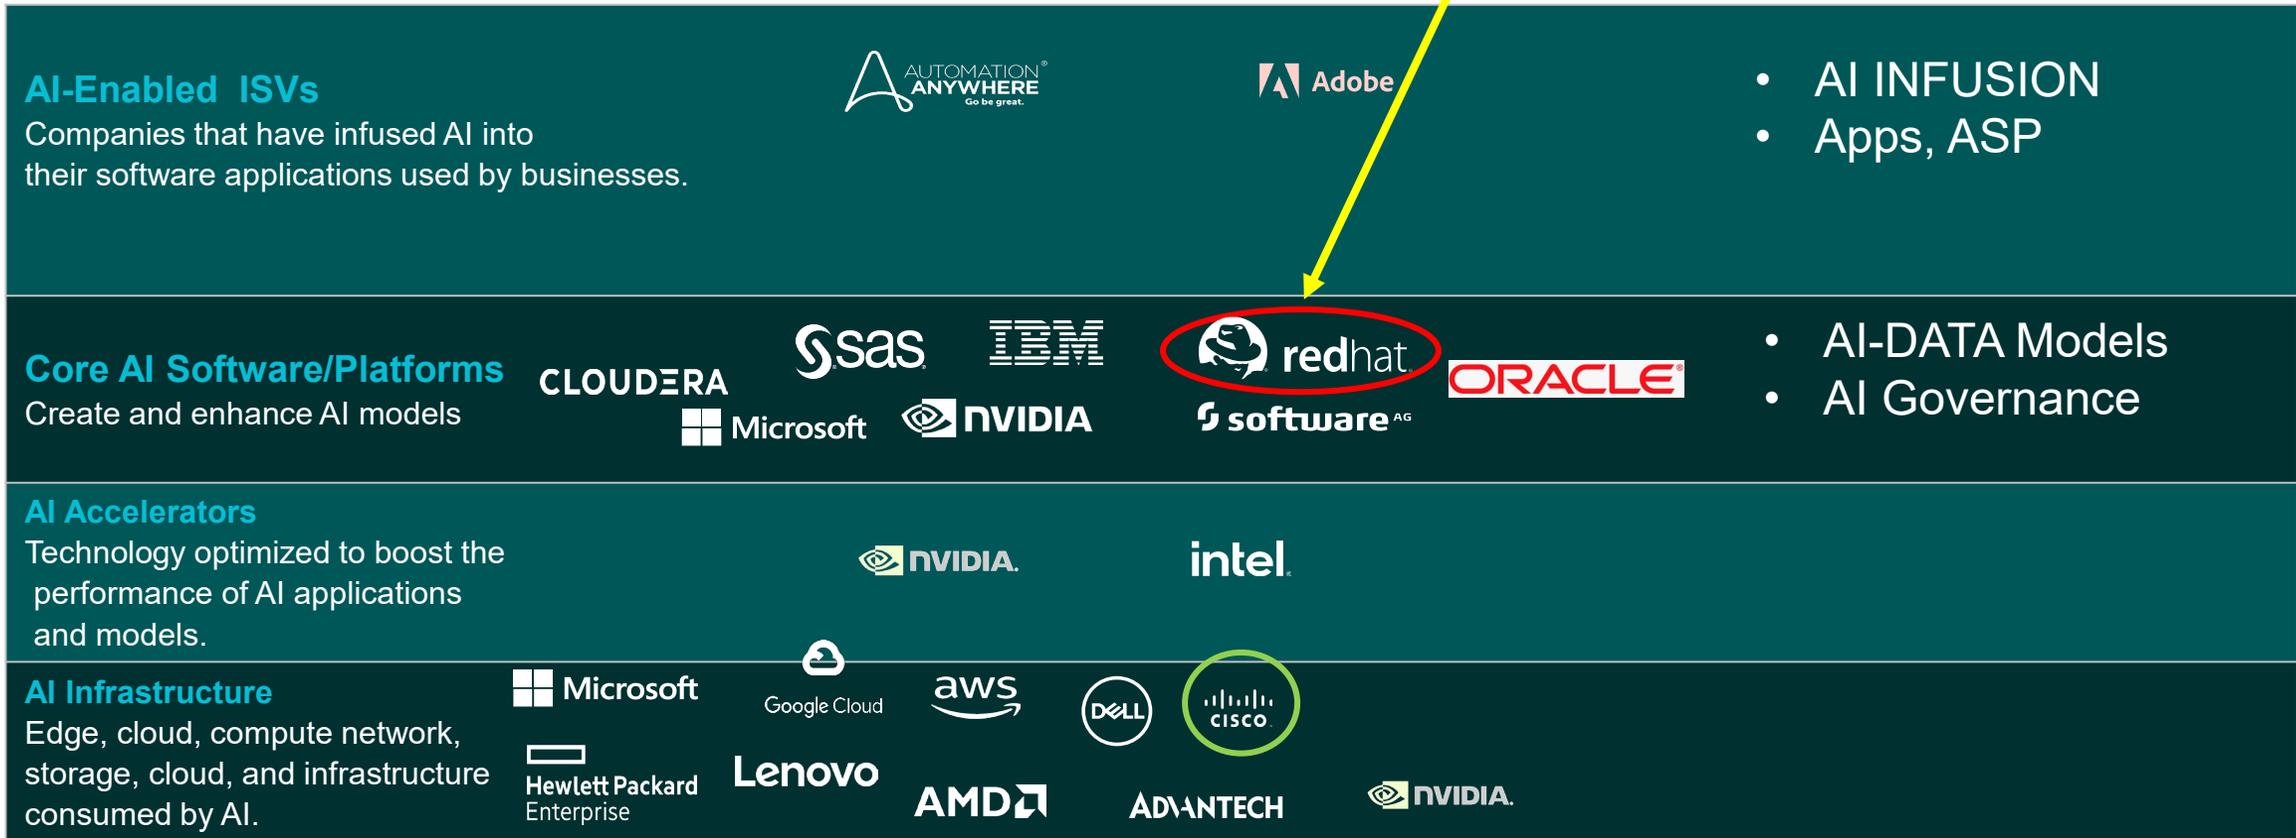
USE CASES & REFERENCES

\* Eventuali brand dei vendor citati e riferimenti a scenari e soluzioni Red Hat, possono essere contenuti pubblici condivisi dal vendor

# The TD SYNnex AI Ecosystem – RED HAT positioning

Vendors that have selected TD SYNnex to *amplify* their position in AI

Current Vendors



\* Some Vendors have solutions across multiple segments. \*\* Not all vendors are available in every region



TD SYNnex

June 11<sup>th</sup> 2025

# Red Hat solutions for AI

*Mauro Lavezzari  
Hybrid Cloud Architect  
TD Synnex Advanced Solutions  
mauro.lavezzari@tdsynnex.com*

# — Red Hat strategy on AI

Open AI solutions on hybrid cloud

# Generative AI customer adoption challenges



## Cost

Generative AI frontier model services are cost prohibitive at scale for most enterprise customer use cases.



## Complexity

Tuning models with private enterprise data for customer use cases is too complex for non-data scientists.



## Flexibility

Enterprise AI use cases span data center, cloud & edge and can't be constrained to a single public cloud service.

# Guiding principles: Open Source

“Meta’s LeCun Says Open-Source Is The Key”<sup>1</sup>

*"InstructLab lets anyone contribute skills to models, making AI truly collaborative. It's how open source won in software -- now we're doing it for AI."*

Chris Wright, Red Hat CTO

Alibaba makes AI model for video, image generation publicly available<sup>2</sup>

*“Every AI application requires a database, a **Kubernetes** cluster, and a model on an AI accelerator,”*  
Satya Nadella, Microsoft

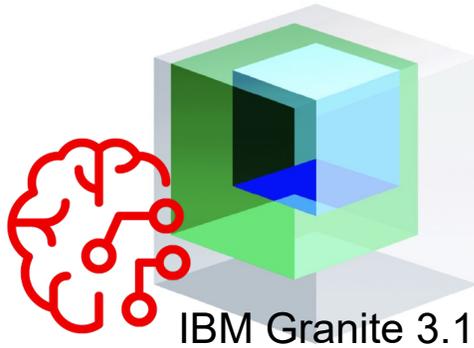
*“A shocking Chinese AI advancement called DeepSeek is sending US stocks plunging”<sup>3</sup>*

<sup>1</sup> <https://www.forbes.com/sites/luisromero/2025/01/27/chatgpt-deepseek-or-llama-metas-lecun-says-open-source-is-the-key/>

<sup>2</sup> <https://www.reuters.com/technology/artificial-intelligence/alibaba-release-open-source-version-video-generating-ai-model-2025-02-25/>

<sup>3</sup> <https://edition.cnn.com/2025/01/27/tech/deepseek-stocks-ai-china/index.html>

# Guiding principles: high quality models in the client context



IBM Granite 3.1

## Winning with Models

open source,  
small language models



## Customizing models

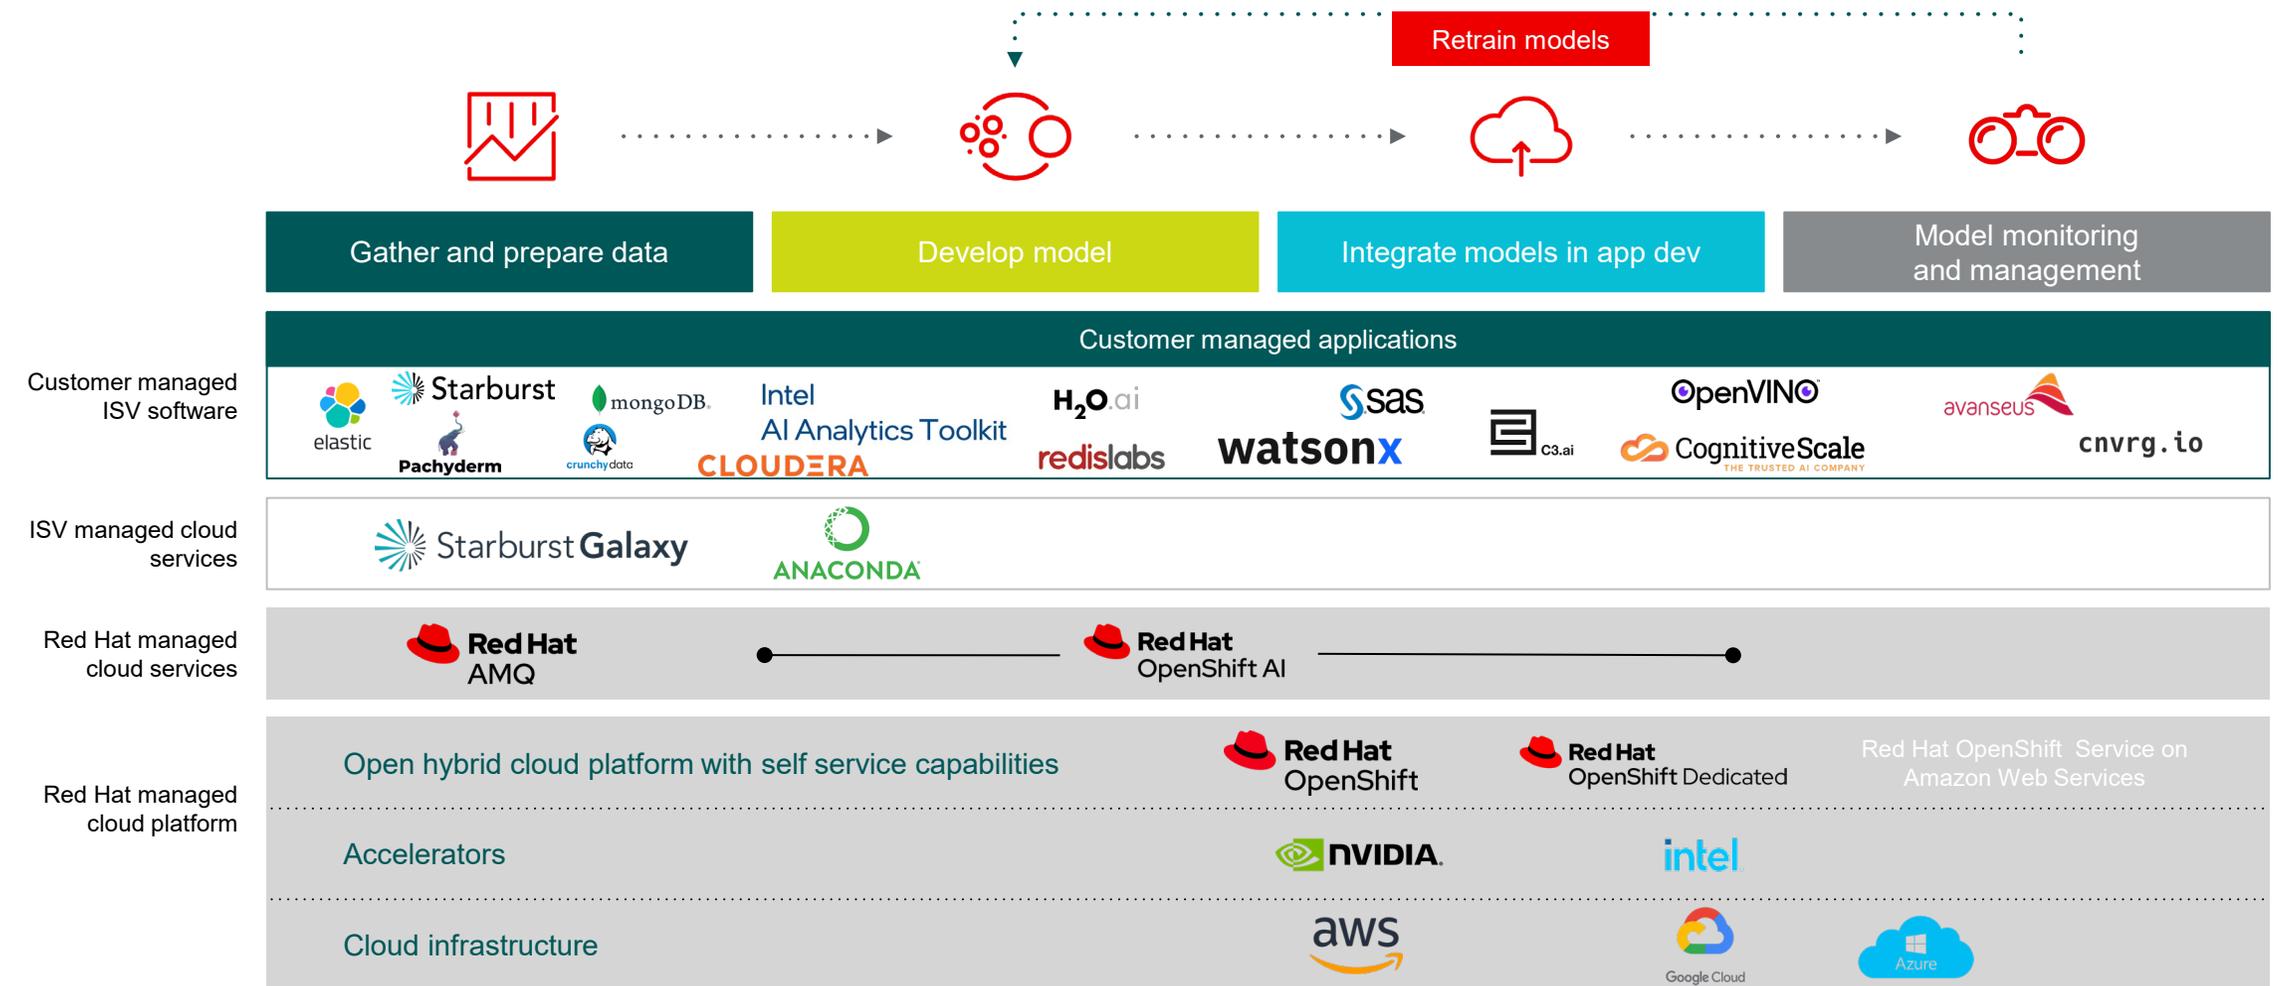
with enterprise,  
private data



## Deploying across the hybrid cloud

models, code and  
intelligent applications

# Guiding principles: Partner Ecosystem



# Red Hat AI / ML Strategy



## AI workload support

Become the default choice for customers' AI deployments by supporting the **requirements of AI workloads on Red Hat infrastructure.**

*e.g., hardware acceleration*



## Platform for intelligent apps

Provide a consistent, purpose-built application development **platform** for customers **to build and deploy AI-enabled applications.**

*e.g., Red Hat OpenShift AI, Red Hat OpenShift Data Science*



## AI-enabled platforms

Leverage **AI-enabled tools and services to drive adoption of existing Red Hat products and services.**

*e.g., Ansible Lightspeed, Developer Hub*

---

## What are the challenges impeding progress?

- 1. How do I privately include my own domain expertise in an AI app?*
- 2. How do I manage the entire lifecycle of my AI App?*
- 3. How do I maximise the efficiency of scarce resources?*

# How do I privately include my own domain expertise in an AI app?

## RAG

*Retrieval Augmented Generation*

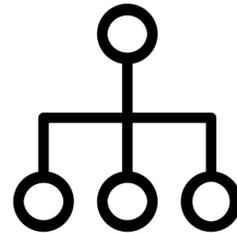


Enhance Gen AI model-generated text by retrieving relevant information from internal and external sources.

NEW

## InstructLab

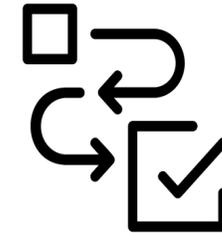
*Large-scale Alignment for chatBots*



Allow non-experts to contribute their domain expertise to a model

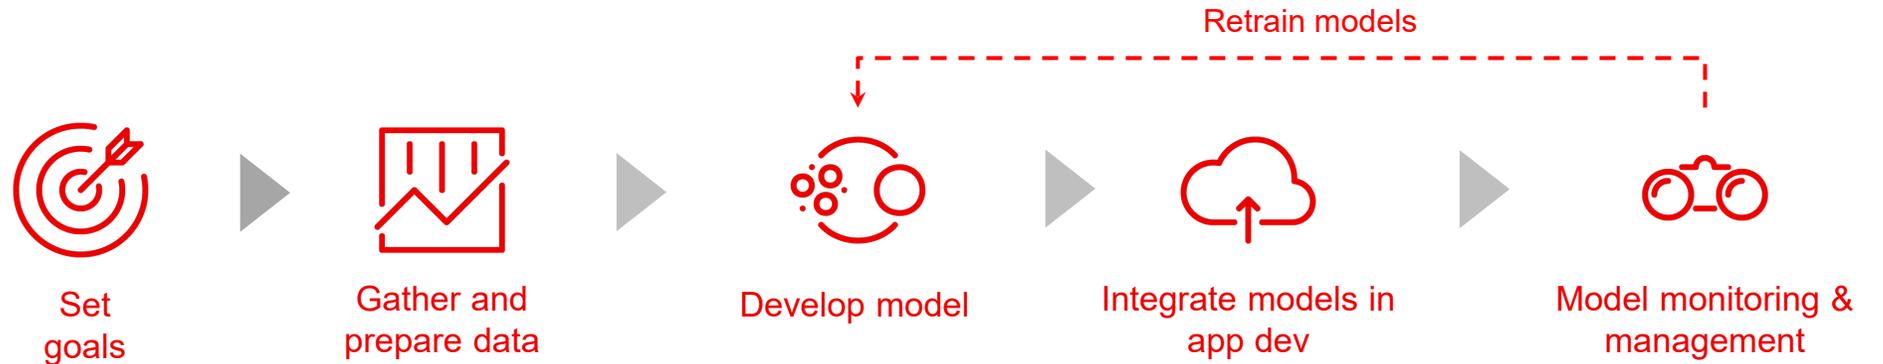
## Fine tuning

*Fine Tuning*



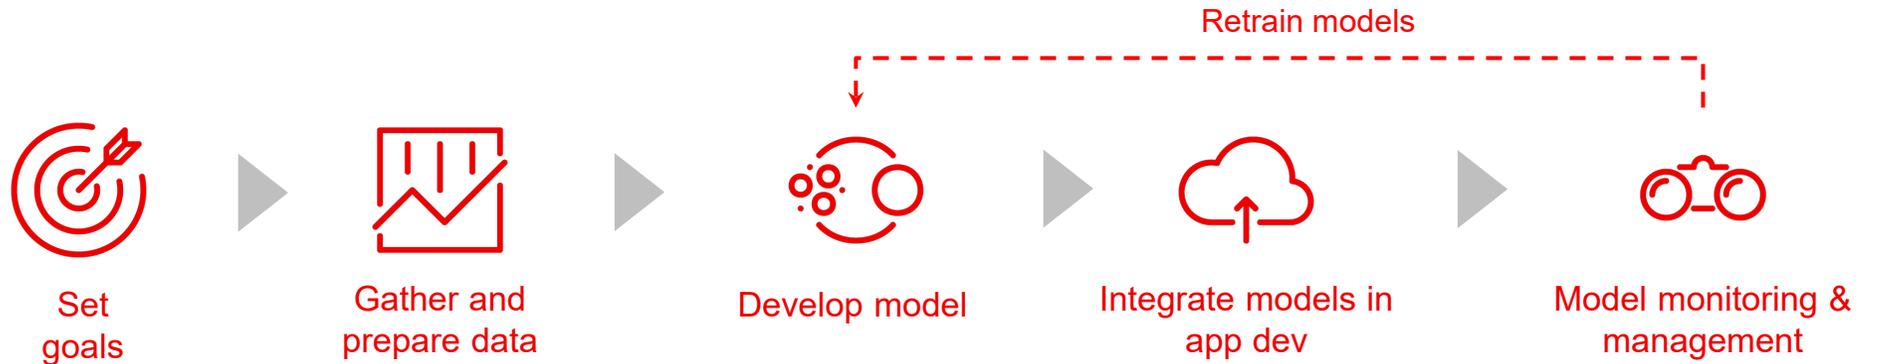
Adjust a pre-trained model on specific tasks or data, improving its performance and accuracy for specialized applications without full retraining.

# How do I manage the entire lifecycle of my AI App?

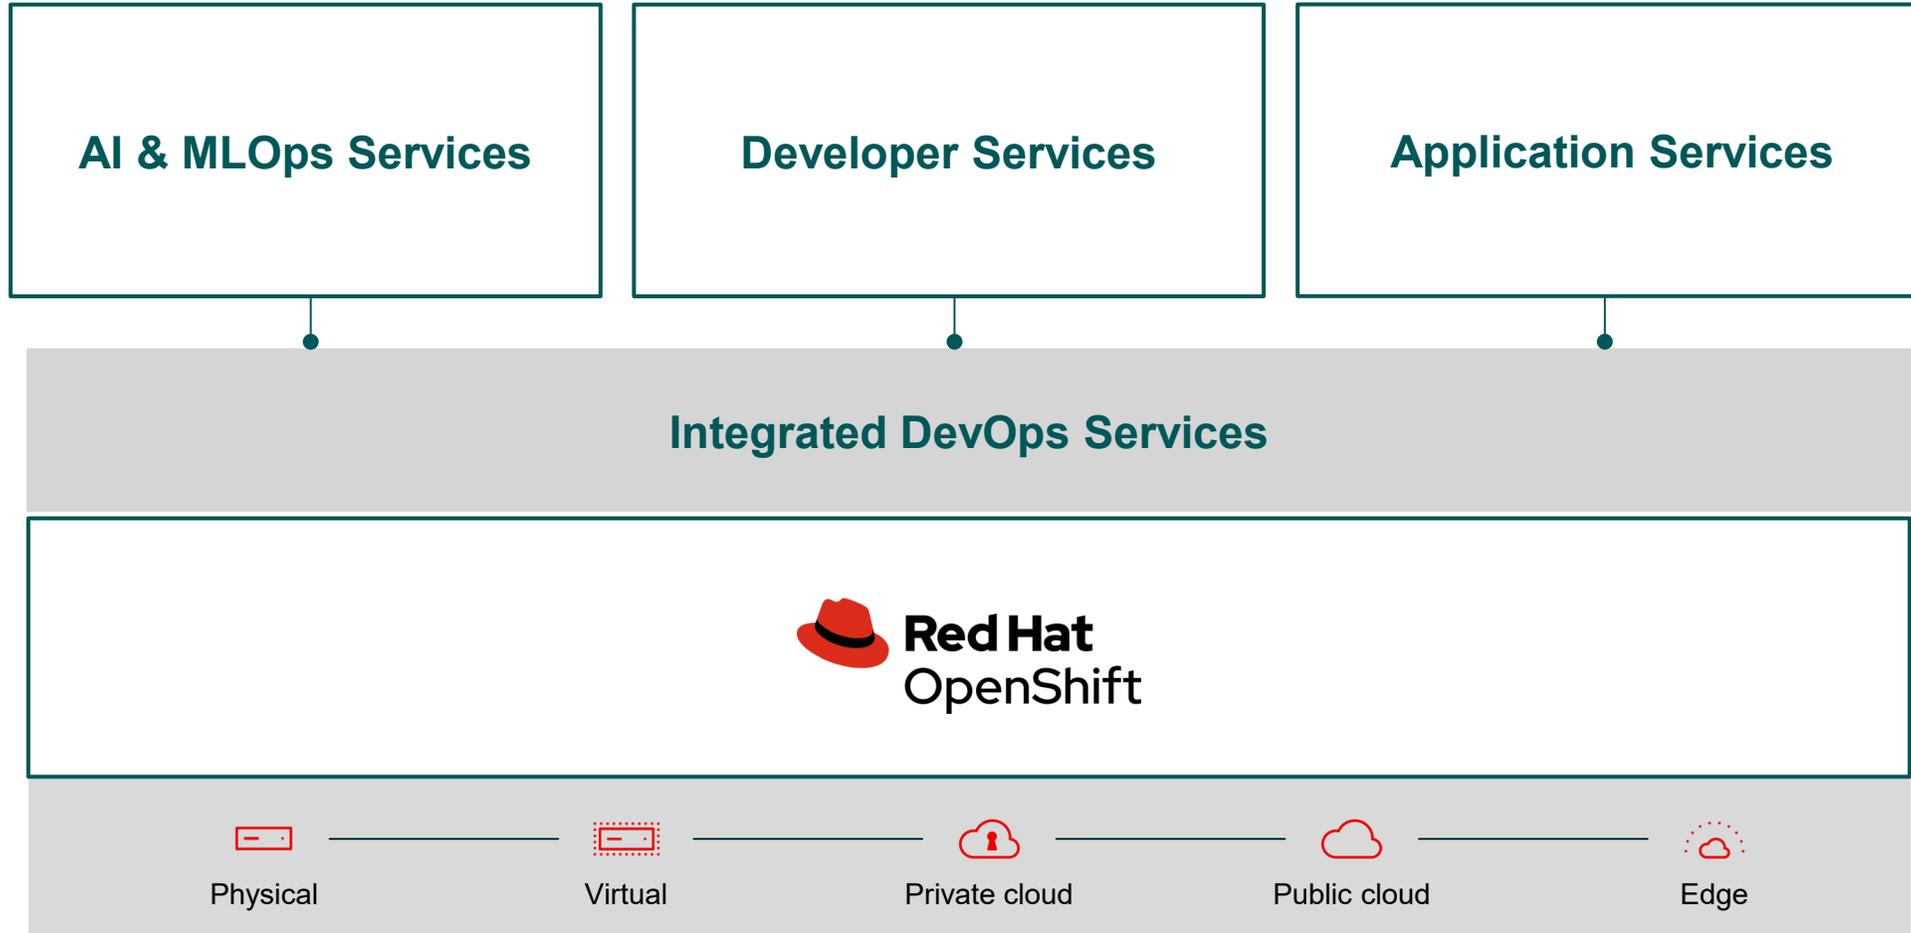


-  Business leadership
-  Data engineer
-  Data scientists
-  ML engineer
-  App developer
-  IT operations

# How do I manage the entire lifecycle of my AI App?



# How do I manage the entire lifecycle of my AI App?



*“Every AI application requires a database, a **Kubernetes** cluster, and a model on an AI accelerator,”*

*Satya Nadella, Microsoft*

---

## Summary: Why Red Hat?

*1. How do I privately include my own domain expertise in an AI app?*

InstructLab allows non -experts to contribute their domain expertise

*1. How do I manage the entire lifecycle of my AI App?*

OpenShift's DevOps and MLOps capabilities allows all stakeholders to effectively contribute to the lifecycle of an application

*1. How do I maximise the efficiency of scarce resources?*

Model-as-a-Service simplifies Development, Management and the efficiency of AI Applications

# Red Hat AI portfolio

## Generative AI & Predictive AI

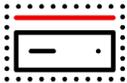


Trusted, Consistent and Comprehensive foundation

NVIDIA AMD intel Hardware Acceleration aws Google IBM



Physical



Virtual



Private  
Cloud



Public  
Cloud



Edge

\* NVIDIA GPUs fully supported in Red Hat AI. AMD Instruct \$ Intel Gaudi in Preview, AWS Inferentia/Neuron, Google TPU, IBM AIU are on our roadmap

# Red Hat AI portfolio (II)



## Foundation model platform for developing, testing and running Granite family LLMs

- Provides a simplified approach to get started with generative AI that includes open-source models
- Makes AI accessible to developers and domain experts with little data science expertise
- Provides the ability to do training & inference on individual production server deployments



## Integrated MLOps platform for model lifecycle management at scale anywhere

- Provides support for both generative and predictive AI models with *bring-your-own-model* (BYOM) approach
- Includes distributed compute, collaborative workflows, model serving and monitoring
- Offers enterprise MLOps capabilities and the ability to scale across hybrid-clouds
- Includes Red Hat Enterprise Linux AI, including the Granite family models

# From RHEL AI to OpenShift AI



**InstructLab**

## Open Source

Learn and experiment via limited desktop-scale training method (qlora) on small datasets.



**Laptop/Desktop**



**Red Hat**  
Enterprise Linux AI

## Small Scale

Production-grade model training using full synthetic data generation, teacher and citric models. CLI tooling with building blocks.



**Server/VM**



**Red Hat**  
OpenShift AI

## Large Scale

Production-grade model training, serving and monitoring using full power of hybrid cloud app platform for scaling, automation and MLOps services.



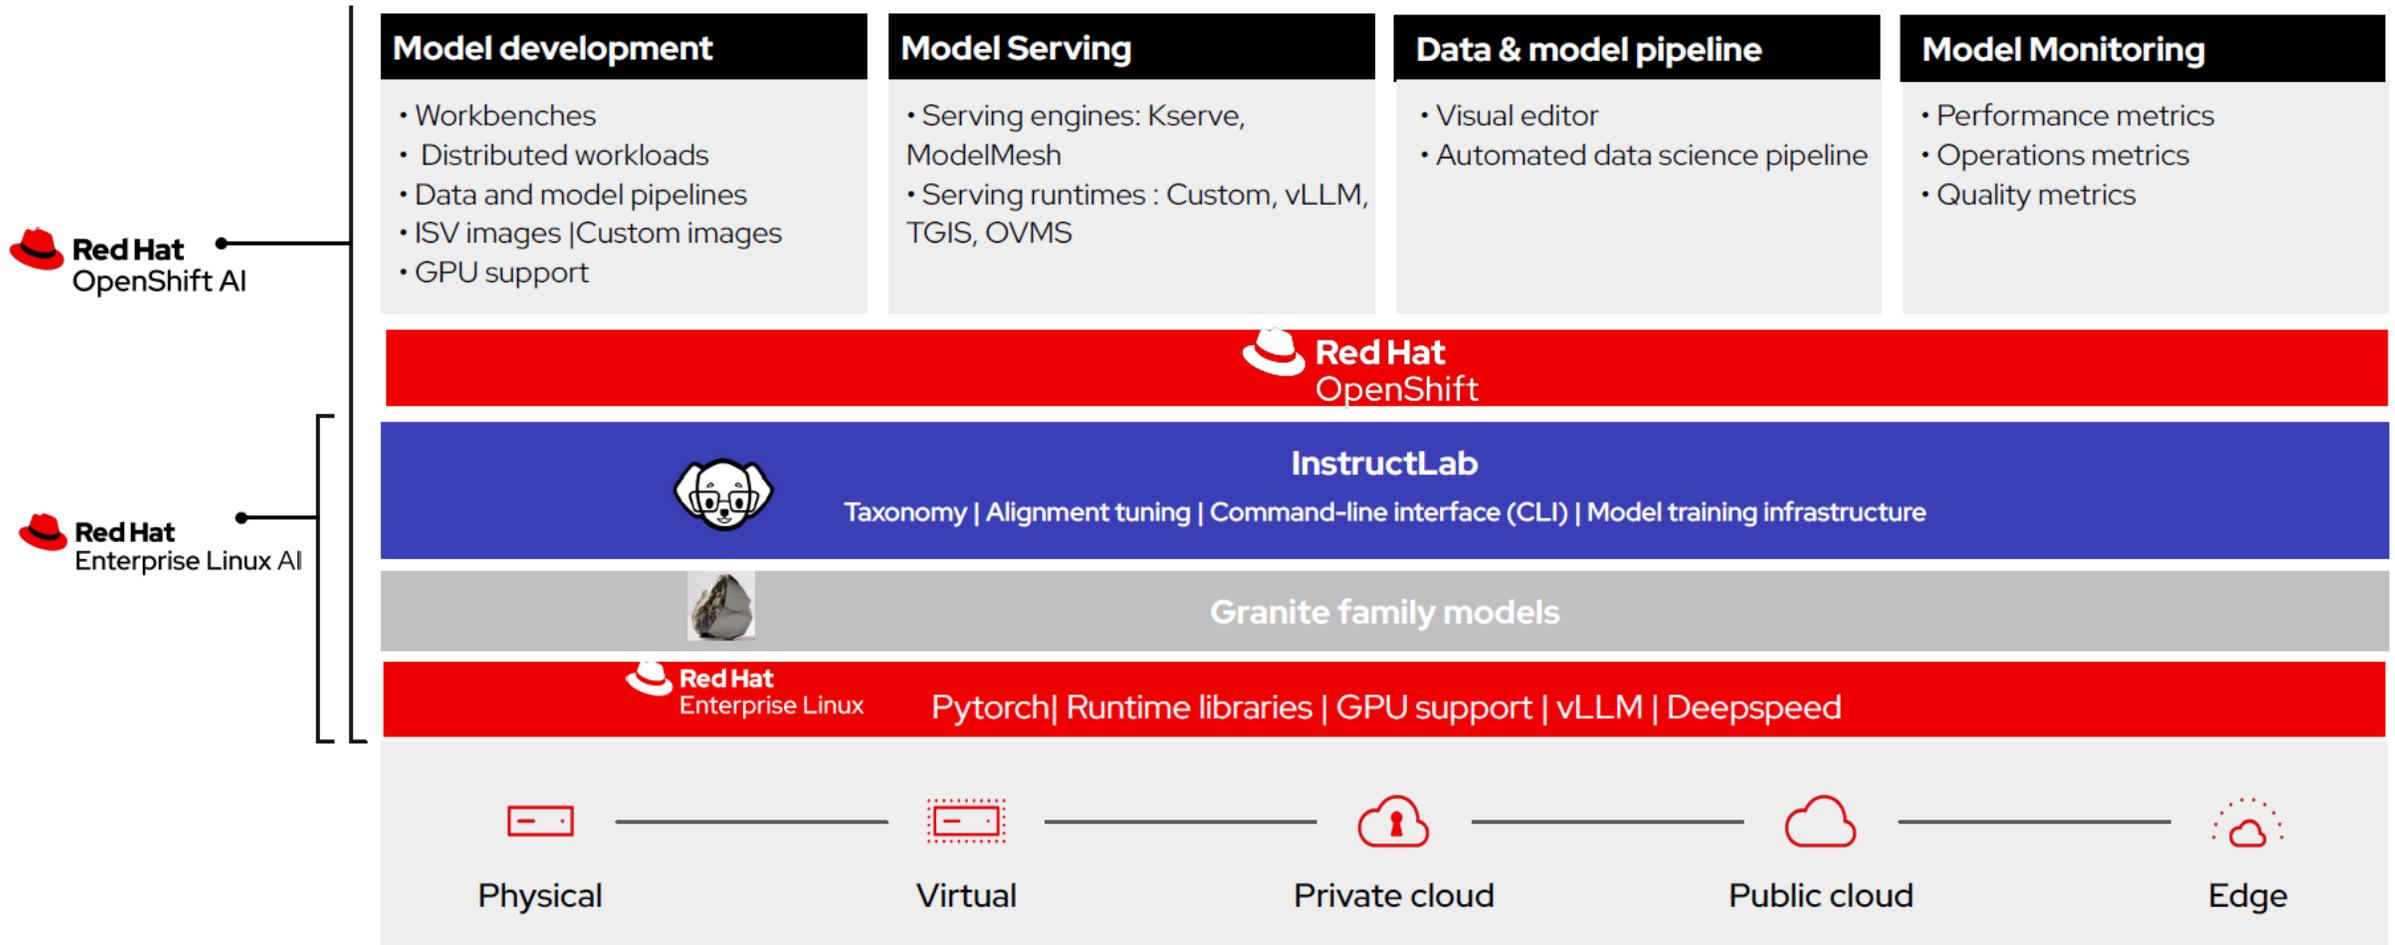
**Cluster**

# AI/ML 100% open source

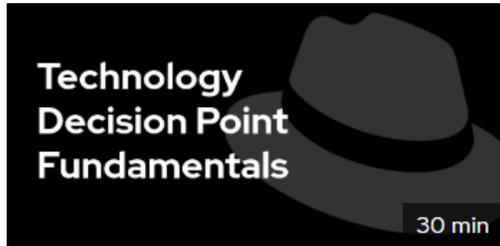


# Red Hat OpenShift AI platform

## Generative AI and MLOps capabilities for building flexible, trusted AI solutions at scale



# AI Enablement & Trainings



Technology Decision Point Fundamentals

30 min



5 Ore

Course

Technology Decision Point Fundamentals



Giocare



Deploying RHEL AI on AWS

5 Ore



30 Minuti

Course

Red Hat AI: Solving Technical Challenges



Giocare



Red Hat AI: Solving Technical Challenges

30 Minuti



4 Ore

Course

InstructLab: Training Merlinite for Parasol Insurance



Giocare



InstructLab: Training Merlinite for Parasol Insurance

4 Ore



Red Hat Solutions for Generative AI and LLMs

15 Minuti

Course

Red Hat Solutions for Generative AI and LLMs



Giocare



AI Fundamentals

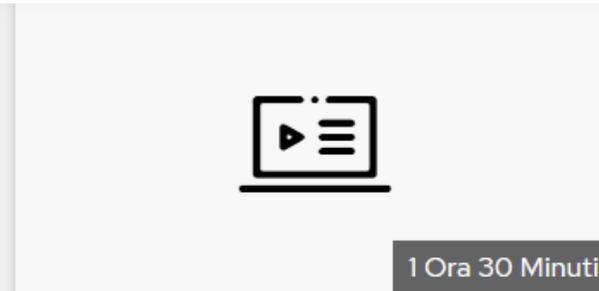
1 Ora 30 Minuti

Course

AI Fundamentals



Giocare



Serving an LLM using OpenShift AI

1 Ora 30 Minuti

Course

Serving an LLM using OpenShift AI



Giocare

# OpenShift AI Certification



Authorized Distributor

EX267

## Red Hat Certified Specialist in OpenShift AI (EX267)

Overview

Objectives

What you need to know

Recommended Training Tab

### Related offer

This exam is part of the [Developing and Deploying AI/ML Applications on Red Hat OpenShift AI with Exam \(AI268\)](#) bundle. This bundle is available at 15% off local MSRP, now through September 30, 2024. Click [get started](#) to checkout from the bundle page and select a course plus the exam to receive this offer. Discount will be applied to the local MSRP of eligible offerings at checkout. This offer may not be combined with other offers or discounts.

---

# Assets from Red Hat



Authorized Distributor

## **["Understanding AI" page](#) in Red Hat public site**

- *Root access to AI products documentation*
- *Documents, videos, blogs and e-books on AI concepts*
- *Client references*

## **[Red Hat Developer site](#) provides interesting material on AI**

- Documentation
- Articles
- Demos
- Labs

---

# Assets from TD SYNEX



Authorized Distributor

## ***TD Synnex Europe Center of Excellence (CoE) on Red Hat provides skill and assets***

Technical sessions dedicated to Red Hat AI products and technologies. 2 examples:

- Red Hat Enterprise Linux AI general overview
- Red Hat Enterprise Linux AI installation and usage for enhancing an existing llm with local knowledge

Technical workshops on general AI techniques. For example:

- Using RAG (Retrieval Augmented Generation) in combination with AWS Bedrock running on Red Hat.

Workshops and demos are available on TD Synnex owned infrastructures or on cloud, where many Red Hat products are already installed and configured

---

## Some Red Hat use cases / references

[Banco Galicia onboards new corporate customers in minutes with intelligent NLP platform](#)

[Boston University sees success with new educational tools built with Red Hat OpenShift AI](#)

[The Basque Government through EJIIE promotes linguistic diversity with high-speed translation tool](#)

# OpenShift AI at Banco Galicia



Authorized Distributor

Banco Galicia onboards new corporate customers in minutes with intelligent NLP platform

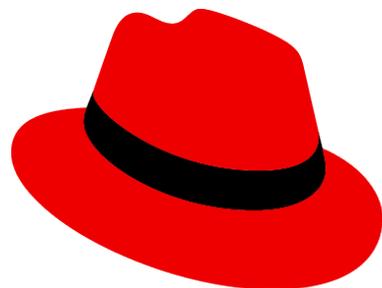
## Building an open source intelligent NLP platform

Banco Galicia worked closely with Red Hat Consulting to build an AI-based intelligent NLP platform. Red Hat OpenShift provides the on-premise platform with a reliable and highly available foundation, with the AI workloads deployed in containers to optimize inference processing.

To develop and train the AI models, Banco Galicia's data scientists use a comprehensive set of open source AI tools from the Open Data Hub (ODH) open source project. This includes Jupyter Notebooks (for development) and TensorFlow and scikit-learn (for training), which all run on Red Hat OpenShift.

Continuous re-training the API-based models allows the bank to improve processes, and then make the models available for other projects and areas.

Banco Galicia also takes advantage of Red Hat Integration and Red Hat Single Sign-on (SSO) to further streamline the platform. Red Hat Integration technologies expose models to any internal workloads not running on the cloud, while Red Hat SSO simplifies and secures access across the Red Hat technologies.



**Red Hat**

Thank you!



**TD SYNEX**