

Accelerate your Enterprise AI journey with OpenShift AI & Intel technology



Sonja Pierer
Managing Director
Intel Germany



Stefanie Chiras
Senior Vice President
Partner Ecosystem Success
Red Hat

Innovation sponsor



The Beating Heart of Enterprise Open Source


Together, Red Hat and Intel provide agile, cloud-ready network architectures based on high-performance, industry-standard platforms and open, software-defined infrastructure.

[Intel and Red Hat: Leaders in Open Technology](#)



25 years in
partnership

Top contributor to
Linux kernel and
open source



Solutions
optimized
through joint
roadmaps

The AI pipeline runs on Intel

Database, data warehouse, data lake, streaming data, feature store, and model registry

Data

For data collection, ingestion, analysis, cleanup, labelling, validation, preparation, and transformation.



The majority of data prep already happens on Intel Xeon processors.

Train

For feature creation, model selection, training, validation, and evaluation.



Intel Xeon's large memory and fast cores excel running **machine learning** workloads. Intel Gaudi accelerator are well-suited for **deep learning** and Intel Xeon's built-in AI capabilities deliver compelling TCO.

Deploy


For model optimization, packaging, endpoint config, performance monitoring, and model serving.



$\frac{2}{3}$ of data center inferring is run on Intel Xeon Processors¹.

Intel® Gaudi® 3 AI Accelerator

Giant Leap in Performance
and Productivity for AI Training & Inference



DELL Technologies


Hewlett Packard
Enterprise



ASUS

GIGABYTE
TECHNOLOGY

ingrasys[®]

Inventec

QCT
Quanta CLOUD TECHNOLOGY

wlstron


wiwynn[®]

 **IBM Cloud**

intel[®]

Intel® Gaudi® 3 AI Accelerator

Giant Leap in Performance
and Productivity for AI Training & Inference



DELL Technologies


Hewlett Packard
Enterprise



ASUS

GIGABYTE
TECHNOLOGY

ingrasys[®]

Inventec

QCT
Quanta CLOUD TECHNOLOGY

wlstron

wiwynn[®]

 **IBM Cloud**

intel[®]

Open Ecosystem to accelerate Enterprise GenAI Deployments

RETRIEVAL AUGMENTED GENERATION (RAG)

Enterprise
Data




Mature, More Secure and Confidential

x86 CPU Based



Embedding
And
Vectorization

AI
Models



Rapidly Evolving and Disruptive

Accelerator Based

Open Ecosystem to accelerate Enterprise GenAI Deployments

RETRIEVAL AUGMENTED GENERATION (RAG)

Enterprise
Data




Mature, More Secure and Confidential

x86 CPU Based



Embedding
And
Vectorization

AI
Models



Rapidly Evolving and Disruptive

Accelerator Based

Come try: Preview access to Intel® Gaudi®
Accelerators on Intel® Tiber™ AI Cloud

Come visit: OPEA-based Retrieval Augmented Generation
(RAG) on Intel® Gaudi with OpenShift AI tech session

intel®



Red Hat



Open Platform
for Enterprise AI

intel®

Red Hat
Summit

Connect

Zurich | January 15, 2025

Bring your
curiosity

