

Accelerate Your Enterprise AI Journey with OpenShift AI & Intel Technology



Piotr Grabuszyński

Cloud Systems Architect,
Intel Corporation



Agnes Fabok

General Manager, Central and
Eastern Europe (CEE)
Red Hat

Innovation sponsor




Bringing AI EVERYWHERE



AI PC Node
AI Developer Productivity &
Light Inference


AI PC
Broadest AI SW Ecosystem



Node
Fine-tuning,
Inference

Cluster
Light Training, Tuning,
Peak Inference

ENTERPRISE AI & EDGE AI
Open Standard, "Ready to Use"



Super Cluster
Training, Tuning,
Peak Inference

Mega Cluster
Large Scale Training
& Inference

DATA CENTER AI
AI Open, Scalable Systems & Reference Arch



Intel Enterprise AI for Red Hat® OpenShift® AI and Red Hat® Enterprise Linux

Customer and Ecosystem Managed Software

Intel® Tiber™
Edge Platform

 PyTorch

 OpenVINO™

Intel® Tiber™
AI Studio

 **Red Hat**
OpenShift AI

 **Red Hat**
Ansible Automation
Platform


 **Red Hat**
OpenShift
Container Platform

 **Red Hat**
Enterprise Linux



Intel® Gaudi® 3 AI Accelerator

Giant Leap in Performance
and Productivity for AI Training &
Inference



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

CONFIDENTIAL Computing & AI

Now enabled with Red Hat*



intel.
XEON®

Intel® Trust Domain Extensions (Intel® TDX)

Extended to AI workloads and soon Gaudi3

Intel® Tiber™ Trust Services

intel.

*Guest support available from select CSPs. Host support coming soon

Come try: Preview access to Intel® Gaudi® 2 & 3*
Accelerators on Intel® Tiber™ Developer 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

Come by: Our booth for more information