



Anders Huge

Nordic Public Sector & Innovation center lead
Intel



Jonas Widell

Senior Manager, Partner Ecosystem Sales
Red Hat





Stockholm Red Hat Summit Connect: 2025

Intel[®] AI for Enterprise Inference

Open-source software:

Intel is committed

Intel® has a long history with Linux®, actively participating in open-source development and collaboration with the Linux community, to ensure hardware is well-supported and delivers optimal performance on Linux-based systems.

Intel contributes to more than 100 different open-source projects, from the Linux kernel to cloud orchestration and plugins for Kubernetes.



Intel is one of the key contributors to Linux
Linux Kernel Corporate Contributor
since 2007¹

Intel in the news

SoftBank will invest \$2 billion into Intel

UPDATED AUG 19, 2025

By Ramishah Maruf



US to take 10% equity stake in Intel

By Aditya Soni, David Shepardson, Andrea Shalal and Max A. Cherney

August 23, 2025 2:19 AM GMT+2 · Updated August 23, 2025




REUTERS

Nvidia investing \$5 billion in Intel, weeks after U.S. government takes stake

The AI leader's infusion comes after the Trump administration engineered a deal to assume a 10 percent stake in the struggling chipmaker.

Updated September 18, 2025


6 min Summary 44



The Washington Post

Exponential growth in enterprise inference

AI inference market growth



Source: marketsandmarkets.com/Market-Reports/ai-inference-market-189921964.html

Through 2028, at least 50% of GenAI projects will overrun their budgeted costs due to poor architectural choices and lack of operational know-how.

Gartner, "10 Best Practices for Optimizing Generative AI Costs," 2024

Intel Xeon 6 the World's Best CPU for AI

Up to 128 cores per CPU

Increased Memory BW

Up to 2.3x higher memory bandwidth
w/MRDIMM memory


vs. 5th Gen Intel® Xeon® processors¹

Increased LLC

L3 cache as large as 504 MB and with
exceptionally low latency at large L3
access sizes

Intel® AI Software

Variety of tools available for AI
development across GenAI, Edge
deployment and classical Machine
Learning



Intel® Advanced Matrix Extensions
(Intel® AMX)

FP16-based models to enhance AI
performance

Advanced Vector Extensions 512
(AVX-512)

Unique instructions and two 512-bit fused-
multiply add (FMA) units per core

Advanced Vector Extensions 2
(AVX2)

New VNNI instructions and fast up/down
convert for BF16 and FP16

¹See IFA 2024 section of [intel.com/performanceindex](https://www.intel.com/performanceindex) for workloads and configurations. Your results may vary.

Apply for a **free** Gaudi 3 Proof of Concept in **30 seconds**

Choose your GenAI or Virtualization PoC:

- ▶ Building Inference, RAG, AgenticAI, Model-as-a-Service, and other AI Use Cases with Intel Gaudi and Xeon
- ▶ Optimize finetuning with intel Gaudi

Why work with Intel + Red Hat?:

- ▶ Benefit from access to free highly qualified experts from Red Hat and Intel and free access to the latest hardware to build your AI use case / application.



If selected, a Intel / Red Hat representative will contact you via email.

The Intel logo is centered on a solid dark blue background. It features the word "intel" in a white, lowercase, sans-serif typeface. A small, bright blue square is positioned directly above the first vertical stroke of the letter "i".

intel