

Case Study

Intel® Xeon® Processors
Habana® Gaudi®2 HPUs



Speed AI Development and Deployment Using 4th Gen Intel® Xeon® Processors, Habana® Gaudi®2 HPUs, and Hugging Face Open Source Libraries

AWS instances featuring Intel® AI acceleration technologies, with Optimum Intel and Optimum Habana libraries, give companies powerful tools for generative AI implementation.

Solution Summary

- Intel® Xeon® Processors
- Intel® Advanced Matrix Extensions (Intel® AMX)
- Habana® Gaudi®2 HPUs
- Intel® Extension for PyTorch
- OpenVINO™ toolkit
- Amazon EC2 M7i, M7i-flex, and C7i instances



GAUDI®2




Executive Summary

Hugging Face is the leading open platform for AI builders. Its mission is to democratize good machine learning through open science and open source, including the [Optimum Intel library](#), an extension of the Hugging Face Transformers library. From a developer perspective, Hugging Face is the go-to place for free resources, including over a million open source models, data sets, and applications that make generative AI and large language model development easier. The combination of Hugging Face's tools and AI-acceleration features built into 4th Gen Intel® Xeon® processors underlying Amazon EC2 instances prove ideal for companies seeking a more turnkey approach for developing performant and scalable AI solutions. Compared with 3rd Gen Intel Xeon processors, the latest CPUs featuring Intel® Advanced Matrix Extensions (Intel® AMX) can deliver 3x to 10x higher inference and training performance.¹ Hugging Face also performed benchmark tests on the Habana® Gaudi®2 Habana Processing Units (HPUs), finding them roughly twice the speed of Nvidia A100 80GB processors for training and inference.²

Challenge

While Hugging Face's offerings simplify and streamline the process of developing AI applications, companies often face another obstacle to deployment—their infrastructure. Running AI models in-house requires compute architectures capable of accommodating heavy workloads. Organizations needed an easier way to access the latest Intel hardware and software in the cloud to optimize their generative AI and LLM implementations.



Hugging Face's offerings simplify and streamline the process of developing AI applications that companies often face.

Solution

In partnership with Intel, Hugging Face created the Optimum Intel library. The library makes the latest Intel Xeon processor hardware and software available to any Hugging Face user. When working with Amazon EC2 M7i, M7i-flex, and C7i instances with 4th Gen Intel Xeon processors, users can benefit from software tools like the Intel® Neural Compressor and built-in accelerators like Intel® Advanced Matrix Extensions (Intel® AMX). The OpenVINO™ toolkit also eases generative AI deployments with high performance inference optimization choices for PyTorch users.

Users planning to implement large language models (LLMs) with a billion or more parameters will appreciate technologies produced by Habana, an Intel company. The custom-designed Habana Gaudi2 deep learning accelerator is available through the Intel® Developer Cloud. Hugging Face and Intel also offer the [Optimum Habana library](#).

“Amazon instances featuring 4th Gen Intel Xeon processors or Habana Gaudi HPU accelerators—combined with the Hugging Face Optimum Intel and Optimum Habana open source libraries—provide incredibly efficient solutions for companies deploying AI models of all sizes.”

— Jeff Boudier, product director, Hugging Face

Results

Users adopting Amazon instances featuring 4th Gen Intel Xeon processors with the Optimum library can gain a significant performance boost. Testing found that CPUs with Intel AMX can deliver 3x to 10x higher inference and training performance than the previous generation of Intel Xeon processors.¹ Performance advantages like these help users deploy AI solutions faster and more cost-effectively. Benchmark tests also found Habana Gaudi2 processors about twice as fast as Nvidia A100 80GB GPUs for both training and inference.²



Key Takeaways

- Don't try to create an AI model from scratch. Simplify the process using an applicable, pre-trained model from Hugging Face.
- Experiment with models and determine their effectiveness using the free Hugging Face API.
- Take advantage of advanced technologies like Habana Gaudi2 in the Intel Developer Cloud.
- Hugging Face's Parameter Efficient Fine-Tuning (PEFT) library can save users significant time when fine-tuning language models.
- Deploy models with APIs to build upon, like the Hugging Face Inference Endpoints service.

Where to Get More Information

[Explore Intel Xeon processors.](#)

[Read more about Habana Gaudi2 HPUs.](#)

[Learn about Intel AMX.](#)

[Explore Hugging Face resources.](#)

[Read about best practices for Amazon EC2 instances.](#)

¹ <https://www.intel.com/content/www/us/en/products/docs/accelerator-engines/advanced-matrix-extensions/ai-solution-brief.html>

² <https://huggingface.co/blog/habana-gaudi-2-benchmark>

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

For workloads and configurations visit www.intel.com/PerformanceIndex. Results may vary.

Intel does not control or audit thirdparty data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.