

Intel[®] Developer Cloud Overview

Vladimir Kilyazov

AI Software Solutions Engineer



Notices and Disclaimers

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.

Intel contributes to the development of benchmarks by participating in, sponsoring, and/or contributing technical support to various benchmarking groups, including the BenchmarkXPRT Development Community administered by Principled Technologies.

Your costs and results may vary.

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

Some results may have been estimated or simulated.

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

All product plans and roadmaps are subject to change without notice.

Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. For more information on the factors that could cause actual results to differ materially, see our most recent earnings release and SEC filings at www.intc.com.

© 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.

Intel® Developer Cloud

How it works - access

Register your account at <http://cloud.intel.com>

The screenshot shows the Intel Developer Cloud website homepage. At the top, there is a navigation bar with the Intel logo, links for PRODUCTS, SUPPORT, SOLUTIONS, and MORE +, a user profile icon, a language selector set to ENGLISH, and a search bar labeled "Search Intel.com". Below the navigation bar, there is a breadcrumb trail: Developers > Tools > Intel® Developer Cloud > Overview. The main content area features a large hero image of a server room with the text "Intel® Developer Cloud" and "Accelerate AI development using Intel®-optimized software on the latest Intel® Xeon® processors and GPU compute." A prominent "Get Started" button is centered, with a link "Already a Member? Sign In →" below it. At the bottom of the hero section, there are five tabs: Software, Platforms, Resources, Customers, and Sign Up. Below the tabs, there are three main sections: "Get Started with Intel" (with a rocket icon), "Early Technology Access" (with a cloud and server icon), and "Deploy AI at Scale" (with a monitor and graph icon). Each section includes a brief description of the service.

Video: [Get Started with Intel® Developer Cloud | Intel Software](#)

Access resources at <http://console.cloud.intel.com>

The screenshot shows the Intel Developer Cloud console home page. The top navigation bar includes the Intel logo, "Developer Cloud", a region selector set to "us-region-1", a "Help" dropdown, a notification bell, and a user profile icon. A left-hand navigation menu is visible, with "Home" selected. The main content area is titled "Console Home" and is divided into several sections. The "COMPUTE" section includes links for Hardware Catalog, Compute Instances, Instance Groups, and Intel K8s Service. The "SOFTWARE" section includes links for Training and Software Catalog. The "Quick Start" section features links for Hardware catalog, Compute instances, Instance groups, Intel k8s service, Keys, Training Catalog, and Redeem coupon. The "Learning and Support" section includes links for Getting started, Tutorials, and What's new?. The "Gen AI Essentials" section features three cards: "Text-to-Image with Stable Diffusion", "Image-to-Image Generation with Stable Diffusion", and "Simple LLM Inference: Playing with Language Models". Each card includes a brief description and a "Launch" button. A "Notifications" section on the right shows "No notifications yet" with a bell icon and a message: "Stay tuned for exciting updates! No new notifications at the moment." At the bottom of the page, there is a footer with links for Company Overview, Contact Intel, Newsroom, Investors, and Careers.

Intel® Developer Cloud

How it works - resources

Select from Intel portfolio: <https://console.cloud.intel.com/hardware>

The screenshot shows the Intel Developer Cloud Hardware Catalog. The left sidebar contains navigation options: Home, COMPUTE (Hardware Catalog, Compute Instances, Instance Groups, Intel K8s Service, Keys), and SOFTWARE (Training, Software Catalog). The main content area is titled "Hardware Catalog" and features a search bar and a filter sidebar. The filter sidebar includes sections for CATEGORY (Released), TYPE (Virtual Machine, Bare Metal), PROCESSOR (CPU, GPU, AI processors), RECOMMENDED USE CASE (Core compute, GPU, HPC, AI), and a "Collapse" button at the bottom. The main area displays "Available platforms" categorized into Core compute, GPU, HPC, and AI. Each category lists specific hardware configurations with "Select" and "Released" buttons. For example, under Core compute, there are options for 4th and 5th Generation Intel® Xeon® Scalable processors. Under GPU, there are options for Intel® Max Series GPU and Intel® Data Center GPU Flex Series. Under HPC, there is an option for 4th Gen Intel® Xeon® processors with high bandwidth memory (HBM). Under AI, there is an option for Gaudi2® Deep Learning Server.

Get access to Service desk: <https://supporttickets.intel.com/>

The screenshot shows the Intel Support Request form. The top navigation bar includes "PRODUCTS", "SUPPORT", "SOLUTIONS", "DEVELOPERS", and "PARTNERS", along with a search bar and language selection (ENGLISH). The main heading is "Support request" with a progress indicator showing three steps: Product Selection, Check for Answers, and Request Creation. Below this, the form is titled "Support Request for Intel® Developer Cloud" and contains a "Describe Your Issue" text area, a "Choose an Issue Severity" dropdown menu, and several action buttons: "Back to Product Selection", "Cancel Request", "Save for Later", "Check for Answers", and "Give Feedback".

Intel® Developer Cloud

Free tier access

Select Training tab: <https://console.cloud.intel.com/training>

The screenshot displays the Intel Developer Cloud console interface. On the left is a dark sidebar with the Intel logo and 'Developer Cloud' text. Below this, there are two main sections: 'COMPUTE' and 'SOFTWARE'. The 'COMPUTE' section includes links for 'Hardware Catalog', 'Compute Instances', 'Instance Groups', 'Intel K8s Service', and 'Keys'. The 'SOFTWARE' section includes a link for 'Training', which is currently selected and highlighted in blue. The main content area has a light gray header with 'us-region-1', 'Help', a notification bell, and a user profile icon. Below the header is a blue horizontal bar with the text 'Training and Workshops' and a 'Launch JupyterLab' button. The main content area is divided into two sections: 'AI' and 'C++ SYCL'. The 'AI' section contains three cards, each with a title, a brief description, and a 'Launch' button with an external link icon. The first card is 'AI Kit XGBoost Predictive Modeling', the second is 'Heterogeneous Programming Using Data Parallel Extension for Numba® for AI and HPC', and the third is 'Machine Learning Using oneAPI'. The 'C++ SYCL' section is partially visible at the bottom.

intel. Developer Cloud

us-region-1 Help

Training and Workshops

Launch JupyterLab

AI

AI Kit XGBoost Predictive Modeling
Learn predictive modeling with decision trees using Intel® AI Analytics Toolkit

Launch

Heterogeneous Programming Using Data Parallel Extension for Numba® for AI and HPC
Data Parallel Extension for Numba accelerates Python® code on Intel® XPU's

Launch

Machine Learning Using oneAPI
Intel® AI Analytics Toolkit accelerates data science and analytics with Python®

Launch

C++ SYCL

Intel® Developer Cloud

Free tier access

Select Training tab: <https://console.cloud.intel.com/training>

■ / Training / AI /

Name

- GenAI
- Introduction_to_Machine_Learning
- Numba_dpex_Essentials_training
- Predictive_Modeling_Training

intel. Developer Cloud

COMPUTE

- Hardware Catalog
- Compute Instances
- Instance Groups
- Intel K8s Service
- Keys

SOFTWARE

- Training

us-region-1 ▾ ? Help ▾  


Training and Workshops

 Launch JupyterLab

AI


AI Kit XGBoost Predictive Modeling

Learn predictive modeling with decision trees using Intel® AI Analytics Toolkit

 Launch 


Heterogeneous Programming Using Data Parallel Extension for Numba® for AI and HPC

Data Parallel Extension for Numba accelerates Python® code on Intel® XPU

 Launch 

Machine Learning Using oneAPI

Intel® AI Analytics Toolkit accelerates data science and analytics with Python®

 Launch 

C++ SYCL

Intel® Developer Cloud

Deploying a Dedicated System

Hardware Catalog

Hardware Catalog

Available platforms

Search

Filter by:

CATEGORY

Released

TYPE

Virtual Machine

Bare Metal

PROCESSOR

CPU

GPU

AI processors

RECOMMENDED USE CASE

Core compute

GPU

HPC

AI

Core compute

4th Generation Intel® Xeon® Scalable processors
Virtual machines and Bare Metal servers available

Select Released

5th Generation Intel® Xeon® Scalable processors
2 sockets, 128 cores, 1 TB memory, 1 TB disk

Select Released

GPU

Intel® Max Series GPU
4th Gen CPU, 2 sockets, 256 GB memory, 2 TB disk

Select Released

Intel® Data Center GPU Flex Series with 4th Gen Intel® Xeon® processors
GPUs with 4th Gen CPU, 2 sockets

Select Released

HPC

4th Gen Intel® Xeon® processors with high bandwidth memory (HBM)
2 sockets, 2 TB disk, SNC4 enabled

Select Released

AI

Gaudi2® Deep Learning Server Cluster
Cluster of 8 Gaudi2® HL-225H mezzanine cards with 3rd Gen Xeon® processors, 1 TB RAM, 30 TB disk

Select Released

Gaudi2® Deep Learning Server
8 Gaudi2® HL-225H mezzanine cards with 3rd Gen Xeon® processors, 1 TB RAM, 30 TB disk

Select Released

[← Back to Instance types](#)

Launch a compute instance

Specify the needed Information

Instance configuration

Instance family: *

Gaudi2® Deep Learning Server

More information:

The Habana® Gaudi2® processor is designed to maximize training throughput and efficiency, while providing developers with optimized software and tools that scale to many workloads and systems. For more information see the [Gaudi2 product page](#). You can find more information on Gaudi and SynapseAI® on [Habana's Developer Site](#).

Instance type: *

Gaudi2® Deep Learning Server featuring eight Gaudi2® HL-225H mezzanine cards with 3rd Gen Xeon® Platinum 8380 Processors \$39 / hour
8 Gaudi2® HL-225H mezzanine cards with 3rd Gen Xeon® Platinum 8380 Processors, 1 TB RAM, 30 TB disk

[Compare instance types](#)

Machine image: *

ubuntu-22.04-server-cloudimg-amd64-latest

Version: 22.04

Architecture: X86_64 (Baremetal only)

More information:

Ubuntu Server is a version of the Ubuntu operating system designed and engineered as a backbone for the internet.

Instance name: *

Instance name

Name must be 63 characters or less, and can include letters, numbers, and '-' only. It should start and end with an alphanumeric character.

Public Keys

Select keys *

No keys found. Please create a key to continue.

[+ Create Key](#)

[↻ Refresh Keys](#)

[Launch](#)

[Cancel](#)

Intel® Developer Cloud Online Documentation

- Popular topics:
 - User account types
 - How to use SSH keys
 - Tutorials

[Guides — Developer Cloud Docs documentation \(intel.com\)](https://www.intel.com/developer-cloud/docs/guides)

Tutorials

Ready to train and evaluate models? Dive into machine learning, or learn SYCL C++ essentials for cross-platform development. Explore our tutorials.

Visual Studio Code Dev Set up Visual Studio Code* app on an instance	Gaudi2 Deep Learning Server Train with Gaudi2 Deep Learning Server	Jupyter Notebook training Run popular use cases in Jupyter* Notebooks
XPU Verify Tool Run a suite of tests for discrete GPUs on Linux* operating systems.	Processor Finder Find the right processor by AI/ML use case.	

The screenshot displays the Intel Developer Cloud Docs website. The navigation bar includes the Intel logo, 'Developer Cloud Docs', a home icon, and links for 'Guides', 'Tutorials', 'Reference', 'Support', and 'About'. A 'Sign in to the console' button is located in the top right. A search bar is positioned below the navigation. The left sidebar lists navigation options: Home, Guides (expanded), Get Started, SSH Keys, Manage Instance, User Accounts, Billing and usage, Support Matrix, Intel Kubernetes Service, Tutorials, Reference, Support, and About. The main content area is titled 'Guides' and features a sub-header: 'Learn the benefits of using Intel processors integrated with AI/ML software stacks.' Below this, there are six guide cards: 'Get started' (with sub-points: Sign up, Set up SSH Keys, Manage access), 'SSH keys' (Learn how to add and use SSH keys to connect to instances.), 'Manage instance' (Manage a compute instance (bare metal or virtual machine) from creation to deletion.), 'User accounts' (Learn about our service tiers: standard, premium, and enterprise.), 'Intel Kubernetes Service' (with sub-points: Launch a Kubernetes cluster, Add nodes to cluster, Connect and develop), and 'Billing usage' (Track your usage and choose how to pay your invoices.). A 'Support Matrix' card is also visible at the bottom left of the main content area.

Intel® Developer Cloud

Additional Resources

How-to Videos

- [Get Started with Intel® Developer Cloud](#)
- [Gaudi2 Running Hugging-Face on Intel® Developer Cloud](#)
- [Gaudi2 on Intel® Developer Cloud](#)

Article & video: [Touring the Intel AI Playground – Inside the Intel Developer Cloud](#)

- *STH*, by Patrick Kennedy, Nov. 2023

[IDC Infograph](#)

Intel blogs

- [Democratizing AI Access with Intel® Developer Cloud](#) - Nov. 2023
- GenAI Essentials
 - Part 1: [LLM Chatbots with Camel 5B & Open Llama 3B v2 on Latest Intel GPU](#)
 - Part 2: [Text-to-Image Stable Diffusion with Stability AI & CompVis models on Latest Intel GPU](#)
 - Part 3: [Image-To-Image Stable Diffusion with Runway ML & Stability AI on Latest Intel GPU](#)
- [Text-to-SQL Generation Using LLMs Fine-Tuned with QLoRA on Intel GPUs](#)

The Intel logo is centered on a solid blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®).

intel®

Legal Notices and Disclaimers

For notices, disclaimers, and details about performance claims, visit www.intel.com/PerformanceIndex or scan the QR code:



© 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.