



NVIDIA AI Enterprise Packaging, Pricing, and Licensing Guide

Application Note

Document History

SWE-NVAIE-011-APNT

Version	Date	Authors	Description of Change
01	August 2021	MR, CW	Initial release
02	November 2021	MR, RR	Added support for A2, A100X and A30X
03	March 2022	MR,NS	Added STS and Inception pricing models
04	September 2022	MR, AT	Added support for H100
05	December 2022	MR, AT	Added support for CSP marketplace
06	February 2023	MR, AH	Added support for H800
07	April 2023	MR, AT	Added support for L4, L40. Updated Licensing model.
08	July 2023	MR	Added section on DGX
09	August 2023	DH	Added section on ISV Pricing Program
10	September 2023	MR, PL	Added L40S, H100 NVL support. Added support for Base Command Manager Essentials
11	October 2023	MR, SR	Updated pricing for Perpetual licensing & SUMS

Table of Contents

Chapter 1. Terminology	5
Chapter 2. Overview	6
2.1 General Purchasing Information	10
2.2 NVIDIA AI Enterprise	11
2.3 NVIDIA AI Enterprise Products and Entitlement	11
Chapter 3. NVIDIA AI Enterprise Licensing and Pricing	14
3.1 NVIDIA AI Enterprise Pricing	14
3.2 Subscription License	15
3.3 NVIDIA H100 GPU Includes NVIDIA AI Enterprise Essentials	16
3.4 DGX software bundle & NVIDIA AI Enterprise	17
3.5 Cloud Deployment of NVIDIA AI Enterprise	17
Chapter 4. Discounted Pricing	18
4.1 NVIDIA Education Pricing Program	18
4.2 NVIDIA Inception Pricing Program	18
4.3 NVIDIA Short Term Support Pricing	18
Chapter 5. NVIDIA License System	19

Chapter 1. Terminology

The following table lists the terms that will be used throughout this guide.

Table 1. Terminology

Term	Meaning
Support Services	Includes technical support, upgrade, and maintenance.
Perpetual License	A non-expiring, permanent software license that can be used on a perpetual basis without the need to renew. Support Services are required and are available in three-year or five-year increments. One-year Support Services available only for renewals.
Subscription	A software license that is active for a fixed period as defined by the terms of the subscription. A subscription includes Support Services for the duration of the subscription term.
License Server	An application that manages license allocation, installed on a physical or virtual server.
GPU	Graphics processing unit.
CPU	The computer's central processing unit (CPU) is the portion of a computer that retrieves and executes instructions.
CPU Socket Licensing	How NVIDIA AI Enterprise was licensed until April 2023. (1) For on-premise deployments, the number of physical processors in the computing environment on which NVIDIA AI Enterprise was installed, or (2) in a cloud computing environment, the number of virtual CPUs for the compute instance on which NVIDIA AI Enterprise runs. NVIDIA required one license per CPU socket.
Per GPU licensing	How NVIDIA AI Enterprise is licensed from April 2023 and later.

Chapter 2. Overview

[NVIDIA AI Enterprise](#), the software layer of the NVIDIA AI platform, is an end-to-end, secure, cloud-native software platform for production AI, enabling organizations to solve new challenges while increasing operational efficiency. It accelerates the data science pipeline and streamlines the development and deployment of AI models to automate essential processes and gain rapid insights from data. With an extensive library of full-stack software including AI workflows, frameworks, pretrained models and infrastructure optimization, the possibilities are endless. Available in the cloud, the data center and at the edge, NVIDIA AI Enterprise enables organizations to develop once and run anywhere. Global enterprise support, API stability and regular security reviews ensure business continuity and AI projects stay on track.

This guide covers the entitlement, packaging, and licensing of NVIDIA AI Enterprise. It is intended to be a quick reference to understand the product at a high level, with the corresponding SKU information. Additional product information can be obtained from the NVIDIA AI Enterprise website at:

<https://www.nvidia.com/en-us/data-center/products/ai-enterprise-suite/>

This guide is not intended to replace or contradict the End User License Agreement (EULA). Refer to the [EULA](#) for more detailed information.

Table 2. NVIDIA AI Enterprise Licensed Products

Product	Description
NVIDIA AI Enterprise - Essentials	The NVIDIA AI Enterprise Essentials edition includes AI frameworks and pretrained models for the development and deployment of production AI. In addition to providing security, reliability and support, NVIDIA AI Enterprise also includes features only available with subscription. This includes - AI workflows which are reference applications for common AI use cases, unencrypted pretrained models and NVIDIA vGPU software.

NVIDIA AI Enterprise - Riva	NVIDIA Riva is a GPU-accelerated speech and translation AI SDK for building fully customizable, real-time, multilingual conversational AI pipelines. Riva is a premium edition of the NVIDIA AI Enterprise platform, and includes NVIDIA AI Enterprise Essentials.
-----------------------------	--

NVIDIA AI Enterprise software editions are supported on prior and current generations of NVIDIA GPUs for deployment flexibility. Customers also have the option to deploy on bare metal systems as well as virtualized environments. Refer to the below Tables 3 and 4 for the NVIDIA GPUs supported with NVIDIA AI Enterprise.

Find NVIDIA Certified Systems compatible with NVIDIA AI Enterprise software at:

<https://docs.nvidia.com/ngc/ngc-deploy-on-premises/nvidia-certified-systems/index.html>

Table 3. Supported NVIDIA GPUs Optimized for Compute Workloads

	NVIDIA H100NVL	NVIDIA H100	NVIDIA HGX A100	NVIDIA A100	NVIDIA A30	NVIDIA A100X	NVIDIA A30X	NVIDIA V100S/V100
Recommended Use case	Compute Optimized	Compute Optimized	Compute Optimized	Compute Optimized	Compute Optimized	Data intensive AI and compute	Data intensive AI and compute	Compute Optimized
Number of GPUs	1 NVIDIA H100	1 NVIDIA H100	4 NVIDIA A100	1 NVIDIA A100	1 NVIDIA A30	1 NVIDIA A100	1 NVIDIA A30	1 NVIDIA V100
FP32 Cores / GPU	16896	14592	6912	6912	3584	6912	3584	5120
Tensor Cores / GPU	528	456	432	432	224	432	224	640
RT Cores			-	-	-			-
Total Memory Size / GPU	94 GB HBM3	80 GB HBM2e	40 GB HBM2/80GB HBM2	40 GB HBM2/80 GB HBM2e	24 GB HBM2	80 GB HBM2e	24 GB HBM2e	32 GB HBM2
MIG Instances/GPU	7	7	7	7	4	7	4	N/A
Max GPU Power / GPU	400W	310 W / 350W	400 W	250W/300 W	165W	300W	230W	250 W / 300 W
Form Factor	PCIe 5.0 2-Slot FHFL	PCIe 5.0 2-Slot FHFL	4x SXM4 GPUs	PCIe 4.0 2-Slot FHFL	PCIe 4.0 2-Slot FHFL	PCIe 4.0 2-Slot FHFL	PCIe 4.0 2-Slot FHFL	PCIe, SXM2, Full Height/Length
Board Dimensions	10.5" x 4.4"	10.5" x 4.4"		10.5" x 4.4"	10.5" x 4.4"	10.5" x 4.4"	10.5" x 4.4"	10.5" x 4.4" 5.5" x 3.1" x .5"
Cooling Solution	Passive	Passive	Passive	Passive / Liquid Cool	Passive	Passive	Passive	Passive



Note: NVIDIA A800, AX800, A30LC, H800 and H800 NVL GPUs are also supported with NVIDIA AI Enterprise. For information on H800 and A800 Tensor Core GPUs that meet regulations for China region, please contact your local NVIDIA partner.

Table 4. Supported NVIDIA GPUs Optimized for Mixed Graphics and Compute Workloads

	NVIDIA T4	NVIDIA A40	NVIDIA A10	NVIDIA A16	NVIDIA A2	NVIDIA L4	NVIDIA L40	NVIDIA L40S
Recommended Use case	Mixed Professional Graphics & Compute (entry)	Mixed Professional Graphics and Compute (high-end)	Mixed Professional Graphics and Compute (midrange)	VDI (Density Optimized)	Mixed Professional Graphics & Compute (entry)	Mixed Professional Graphics (entry-mid) vPC(density Optimized) Compute Optimized	NVIDIA vWS -Performance Optimized (High end). Compute optimized	NVIDIA vWS -Performance Optimized (High end). Compute optimized
Number of GPUs	1 NVIDIA T4	1 NVIDIA A40	1 NVIDIA A10	4 NVIDIA A16	1 NVIDIA A2	1 NVIDIA L4	1 NVIDIA L40	1 NVIDIA L40S
FP32 Cores/GPU	2560	10572	9216	4x 1280	1280	7,424	18,176	18,176
Tensor Cores/GPU	320	336	288	4x 40	40	232	568	568
RT Cores	40	84	72	4x 10	10	58	142	142
Total Memory Size	16 GB GDDR6	48 GB GDDR6	24 GB GDDR6	64GB (4x 16GB) GDDR6	16 GB GDDR6	24 GB GDDR6	48 GB GDDR6	48 GB GDDR6
Max GPU Power	70 W	300 W	150 W	250W	60W	72 W	300 W	350 W
Form Factor	PCIe 3.0 1-Slot HHHL	PCIe 4.0 2-Slot FHFL	PCIe 4.0 1-Slot FHFL	PCIe 4.0 1-Slot FHFL	PCIe 4.0 1-Slot HHHL	PCIe 4.0 Single-Slot LP	PCIe 4.0 Dual-Slot FHFL	PCIe 4.0 Dual-Slot FHFL
Board Dimensions	2.7" x 6.6"	10.5" x 4.4"	10.5" x 4.4"	10.5" x 4.4"	2.7" x 6.6"	6.6" x 2.7"	10.5" x 4.4"	10.5" x 4.4"

Cooling Solution	Passive	Passive	Passive	Passive	Passive	Passive	Passive	Passive
------------------	---------	---------	---------	---------	---------	---------	---------	---------



Note: NVIDIA RTX 6000 Ada, RTX 5000 Ada, A6000, RTX A5000, RTX 6000 and RTX 5000 are also supported with NVIDIA AI Enterprise.

2.1 General Purchasing Information

NVIDIA AI Enterprise editions are licensed on a per-GPU basis and can be purchased through NVIDIA Preferred Partners, select server OEMs and via cloud marketplaces. A list of these Preferred Partners and OEMs is available at:

<https://www.nvidia.com/en-us/about-nvidia/partners/>

All NVIDIA AI Enterprise subscriptions include NVIDIA Business Standard Support and can be purchased as either a perpetual license as an annual or multi-year subscription and on an hourly consumption basis via cloud marketplaces. NVIDIA AI Enterprise with perpetual licenses must be purchased in conjunction with five-year support services. A one-year support service is also available for renewals.

The subscription and consumption offerings are an affordable option to allow IT departments to better manage the flexibility of license volumes. NVIDIA AI Enterprise software purchased with subscription includes support services for the duration of the software's subscription license.

Table 5. General Purchasing Information

Entitlement	NVIDIA AI Enterprise - Business Standard Support
Technical Support	NVIDIA Business Standard support including: <ul style="list-style-type: none"> • 24x7 availability for case filing • 8 am-5 pm support coverage, local business hours, excluding observed holidays • 4-hour initial response
Maintenance	Access to all maintenance releases, defect resolutions, and security patches for flexibility in upgrading as per the NVIDIA AI Enterprise Software Lifecycle Policy
Upgrades	Access to all new major version releases including feature enhancements and new hardware support
Long-term branch maintenance	Available for up to 3 years from general availability as per the NVIDIA AI Enterprise Software Lifecycle Policy
Direct support	Direct access to NVIDIA support engineering for timely resolution of customer-specific issues
Knowledgebase access	✓

Web support	✓
E-mail support	✓
Phone support	✓

2.2 NVIDIA AI Enterprise

The software layer of the NVIDIA AI platform, NVIDIA AI Enterprise accelerates the data science pipeline and streamlines development and deployment of production AI including generative AI, computer vision, speech AI and more. With over 100 frameworks, pretrained models and development tools, NVIDIA AI Enterprise is designed to accelerate enterprises to the leading edge of AI, while also simplifying AI to make it accessible to every enterprise.

Exclusive software and features only available with the purchase of NVIDIA AI Enterprise such as:

- NVIDIA AI Workflows which are reference applications for common AI use cases
- Foundation models for industry use cases such as a multimodal foundational model backbone for various vision AI use cases, defect detection in PCBs, and retail loss prevention
- End-to-end management software including Base Command Manager Essentials for infrastructure management across cloud and data center environments and Triton Management Service (TMS) for automated model deployment
- NVIDIA vGPU software that enables the NVIDIA GPU to be virtualized and is supported to run on leading hypervisor platforms.



For Graphics products, refer to the [vGPU Software packaging, pricing, and licensing guide](#).

2.3 NVIDIA AI Enterprise Products and Entitlement

Table 6. Feature Entitlements

Feature	Essentials and Riva Editions
---------	------------------------------

Per GPU Licensing	✓ ¹
Compute Virtualization	✓
Windows Guest OS Support	
Linux Guest OS Support	✓
Maximum Displays	1
Maximum Resolution	4096 × 2160 (4K)
OpenGL and Vulkan	In-situ Graphics only
CUDA and OpenCL Support	✓
ECC and Page Retirement	✓
MIG GPU Support	✓
Multi-vGPU	✓
NVIDIA GPUDirect	✓
Peer-to-Peer over NVLink	✓
GPU Pass Through Support	✓
Bare Metal Support	✓
AI and Data Science applications and Frameworks	✓
Cloud Native ready	✓
Base Command Manager Essentials	✓
4 GB	✓
5 GB	✓
6 GB	✓
8 GB	✓
10 GB	✓
12 GB	✓
16 GB	✓
20 GB	✓
24 GB	✓
32 GB	✓
40 GB	✓
48 GB	✓
80 GB	✓

Notes:

¹Starting April 2023 maximum 20 concurrent VMs per product license

Table 7. AI and Data Science applications and frameworks supported with NVIDIA AI Enterprise product offerings:

	Essentials	Riva
AI Workflows/Reference Applications		
Intelligent Virtual Assistant		✓
Audio Transcription		✓
Digital Fingerprinting Threat Detection	✓	✓
Next Item Prediction	✓	✓
Generative AI Knowledge Base Chatbot*	✓	✓
Spam Phishing Detection*	✓	✓
Route Optimization	Contact NVIDIA Sales	
Specialized/Use Case		
Clara	✓	✓
Parabricks	✓	✓
Clara MONAI	✓	✓
cuOpt*	Contact NVIDIA Sales	
DeepStream	✓	✓
Maxine	✓	✓
Merlin	✓	✓
Morpheus	✓	✓
NeMo framework	✓	✓
Modulus	✓	✓
Deep Graph Library (DGL)	✓	✓
Riva		✓
Data Science		
RAPIDS Accelerator for Apache Spark	✓	✓
NVIDIA RAPIDS	✓	✓
Model Training		
Apache Next	✓	✓
PaddlePaddle	✓	✓
PyTorch	✓	✓
TAO	✓	✓
TensorFlow	✓	✓

Optimized Inference/Simulation		
TensorRT	✓	✓
Deploy at Scale		
Triton Inference Server	✓	✓
Triton Management Service*	✓	✓
Container Orchestration		
GPU Operator	✓	✓
Network Operator	✓	✓
Cluster Management		
Base Command Manager Essentials	✓	✓
Infrastructure Optimization		
Magnum IO	✓	✓
vGPU	✓	✓
CUDA	✓	✓

*Available in late 2023

Chapter 3. NVIDIA AI Enterprise Licensing and Pricing

3.1 NVIDIA AI Enterprise Pricing

NVIDIA AI Enterprise software is licensed on a per-GPU basis. A software license is required for every GPU installed on the server that will host NVIDIA AI Enterprise. NVIDIA AI Enterprise software can be purchased by enterprises as a subscription, on a consumption basis via cloud marketplaces and as a perpetual license with required 5-year support services. For subscription and cloud consumption licenses, **Business Standard** support is included with the annual license cost. Enterprises can upgrade to **Business Critical** support for an additional cost.

Pricing is suggested pricing only. Contact your authorized NVIDIA partner for final pricing.

3.2 Subscription License

An annual enterprise subscription is active for a fixed period as defined by the terms of the subscription license. To be kept active, the license will need to be renewed at the end of the subscription period. The subscription license includes the software license and production level support services for the duration of the license subscription period.

An annual enterprise subscription is active for a fixed period as defined by the terms of the subscription license. To be kept active, the license will need to be renewed at the end of the subscription period. The subscription license includes the software license and production-level support services for the duration of the license subscription period.

A perpetual enterprise license allows for use of the licensed software indefinitely and includes five year Business Standard Support services. The support services can be renewed on a yearly basis after the expiration of the initial subscription.

Table 7. Software License Model for NVIDIA AI Enterprise Essentials

NVIDIA AI ENTERPRISE LICENSE TYPE	SUGGESTED PRICE	Support Services
Subscription, 1-year term	\$4,500 / GPU	Business Standard Included +\$1,100 for Business Critical
Subscription, 3-year term	\$13,500 / GPU	Business Standard Included +\$3,000 for Business Critical
Subscription, 5-year term	\$18,000 / GPU	Business Standard Included +\$5,000 for Business Critical
Consumption via Cloud Marketplaces	\$2 / GPU / Hour	Business Standard Included, 3 support calls
Perpetual License + 5-year support services	\$22,500 / GPU	Business Standard Included +\$5,000 for Business Critical

Suggested pricing shown. Contact your authorized NVIDIA partner for final pricing.

For compute environments or instances without NVIDIA GPUs, NVIDIA requires one NVIDIA AI Enterprise subscription or license per server or instance independent of the number of CPUs.

Contact NVIDIA sales for Riva pricing.

Table 8. NVIDIA Enterprise Support for Essentials and Riva Editions

SUPPORT OPTIONS	SERVICES	SUPPORTED CHANNELS	SERVICE LEVEL AGREEMENT
Business Standard (Required for Perpetual)	Issue resolution, bug fixes, software updates, maintenance	Phone, portal, email	9 x 5 Business Days (4-hour initial response)
Business Critical (Optional Upgrade/Add-on)	Same as Business Standard Support	Same as Business Standard Support	24 x 7 (1-hour initial response for S1 cases)

Contact NVIDIA Enterprise Support at <https://www.nvidia.com/en-us/support/enterprise/>.

3.3 NVIDIA H100 GPU Includes NVIDIA AI Enterprise Essentials

With each H100 PCIe card that is purchased, a 5-year subscription of NVIDIA AI Enterprise Essentials is included, which can only be used with that specific card. This means the NVIDIA AI Enterprise Essential software can be installed on the [NVIDIA Certified Systems](#) in which the H100 PCIe card is installed but not on a different server or instance.

The 5-year NVIDIA AI Enterprise Essential subscription can be redeemed with the serial number of the H100 PCIe card. Information on how to redeem the 5-year subscriptions of NVIDIA AI Enterprise Essentials is provided [here](#).

The start date for each NVIDIA AI Enterprise Essential subscription included with H100 GPUs is based on the ship date of the specific H100 PCIe board to the OEM partner plus 90 days to account for integration and final delivery to the customer site. This start date of the NVIDIA AI Enterprise subscription can not be modified as it is tied to the specific card.

NVIDIA H800 PCIe GPUs also include support for NVIDIA AI Enterprise Essential with the same 5-year subscription redeemable with the serial number per card. For information on H800 Tensor Core GPUs that meet regulations for China region, please contact your local NVIDIA partner

3.4 DGX software bundle & NVIDIA AI Enterprise

The DGX software bundle includes NVIDIA AI Enterprise for streamlined AI development and deployment and enterprise support.

3.5 Cloud Deployment of NVIDIA AI Enterprise

Cloud deployment of NVIDIA AI Enterprise is available on an on-demand basis in the [Google Cloud Platform \(GCP\) marketplace](#) and will be made available in other marketplaces soon. NVIDIA AI Enterprise Essentials in cloud marketplaces is priced at \$2/ GPU/ hour on an on-demand/pay-as-you-go basis. Custom pricing can be provided for a committed term and usage via a private offer (for further details, please visit [the marketplace pages](#)).

The NVIDIA AI Enterprise license, purchased via an [NVIDIA NPN partner](#), also enables organizations to deploy and run NVIDIA AI Enterprise software as a Bring Your Own License (BYOL) on certified cloud providers, such as Microsoft Azure, AWS, Google Cloud and Oracle Cloud Infrastructure. When NVIDIA AI Enterprise is deployed in a cloud instance as BYOL, one NVIDIA AI Enterprise subscription license is required for each GPU on which NVIDIA AI Enterprise runs. If the instance does not have a GPU, one NVIDIA AI Enterprise subscription license is required for each instance.

Chapter 4. Discounted Pricing

4.1 NVIDIA Education Pricing Program

The NVIDIA Education Pricing Program supports the use of AI computing in teaching and research institutions. The program makes it easy to procure and administer NVIDIA solutions,

software licensing, and services for qualified educational institutions and helps reduce their total cost. For more information on eligibility, review the [NVIDIA Education Pricing Program](#) documentation.

4.2 NVIDIA Inception Pricing Program

The NVIDIA Inception Program supports cutting-edge startups in AI, data science and HPC startups with go-to-market support, expertise and technology. The program makes it easy to procure and administer NVIDIA solutions, software licensing, and services for qualified startups and helps reduce their total cost. NVIDIA Inception Members can purchase up to 16 1-year subscriptions of NVIDIA AI Enterprise at a reduced price within a 12-month period. For more information on NVIDIA Inception Program eligibility, review the [NVIDIA Inception for Startup](#) site.

4.3 NVIDIA Short Term Support Pricing

The NVIDIA Short Term Support (STS) offering of NVIDIA AI Enterprise provides top of tree AI innovations for teams building services and products leveraging NVIDIA AI Enterprise. Service providers, ISVs and other organizations can leverage this offering to access NVIDIA AI Enterprise subscriptions with technical support provided limited to the current and preceding (last two releases only) at a reduced subscription price.

Chapter 5. NVIDIA License System

The NVIDIA License System (NLS) provides monitoring and reporting on license usage for capacity planning. NLS supports the two types of service instances:

1. **Cloud License Service (CLS)** instance: a CLS is hosted on the NVIDIA Licensing Portal.
2. **Delegated License Service (DLS)** instance: a DLS instance is hosted on-premises at a location that is accessible from your private network, such as inside your data center.

DLS is available for download with the NVIDIA AI Enterprise packages from the NVIDIA Licensing Portal.

To ensure that the end user experience is not impaired by license overages or connection issues, NVIDIA AI Enterprise software runs with or without a valid license server connection. The license server allows the IT administrator to track license usage and determine correct sizing for their environments. In compliance with the EULA, IT administrators may also use any other method that reliably tracks the software usage to ensure they have enough licenses for their per-CPU usage.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA, the NVIDIA logo, CUDA, NVIDIA AI Enterprise, NVIDIA HGX, NVIDIA Turing, NVIDIA RTX, NVLink, and TensorRT are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2023 NVIDIA Corporation & affiliates. All rights reserved.