



ThinkSystem NVIDIA HGX B200 180GB 1000W GPU Product Guide

The NVIDIA Blackwell HGX B200 platform propels the data center into a new era of accelerating computing and generative AI, integrating NVIDIA Blackwell GPUs with high-speed interconnect to accelerate AI performance at scale. As a premier accelerated scale-up x86 platform with up to 15X faster real-time inference performance compared to NVIDIA H100, with 12X lower cost and 12X less energy, the B200 platform is designed for the most demanding AI, data analytics, and high-performance computing (HPC) workloads.

The NVIDIA B200 180GB 1000W GPU is offered in the SXM6 form factor, where GPU modules are implemented with a fully-connected NVLink topology in supported ThinkSystem servers. In the ThinkSystem SR780a V3, the HGX B200 GPUs are water-cooled using Lenovo Neptune water cooling.



Figure 1. ThinkSystem NVIDIA HGX B200 1000W 180GB 8-GPU Liquid-Cooled Board

Did you know?

HGX B200 supports NVIDIA advanced networking options to deliver the highest AI performance. In the Lenovo ThinkSystem SR780a V3, the B200 GPUs can be paired with 8x NVIDIA BlueField-3 B3140H VPI 400GbE adapters or 8x NVIDIA ConnectX-7 NDR OSFP400 InfiniBand adapters.

Part number information

The following table shows the ordering information.

Part number	Feature code	Description	Includes NVAIE	Controlled GPU status
HGX form	factor GPUs	3		
CTO only	C696	ThinkSystem NVIDIA HGX B200 1000W 180GB 8-GPU Liquid-Cooled Board	No	Controlled

Table 1. Ordering information

The NVIDIA B200 GPU is Controlled which means the GPU is not offered in certain markets, as determined by the US Government.

Features

Real-Time Large Language Model Inference

As part of HGX, the Blackwell GPU's cutting-edge capabilities with the second generation Transformer Engine mark a massive leap forward in accelerating inference workloads, making real-time performance a possibility for resource-intensive multi-trillion-parameter language models. HGX B200 achieves 15X acceleration, 12X lower cost, and 12X less energy than the H100 HGX GPU for massive models such as the GPT-MoE-1.8T. This advancement is made possible with a new generation of Tensor Cores, which introduce new precisions including FP4.

Supercharged AI Training

Blackwell GPUs include a faster second generation Transformer Engine featuring FP8 precision, enabling a remarkable 3X faster training for large language models like GPT-MoE-1.8T compared to the H100 HGX GPU. This breakthrough is complemented by the fifth-generation NVLink, which provides 1.8TB/s of GPU-to-GPU interconnect, InfiniBand networking, and Magnum IO software. Together, these ensure efficient scalability for enterprises and extensive GPU computing clusters.

Advancing Data Analytics

Databases and Apache Spark play critical roles in handling, processing, and analyzing large volumes of data for data analytics. Blackwell's new dedicated Decompression Engine accelerates the full pipeline of database queries for the highest performance in data analytics and data science. With support for the latest compression formats such as LZ4, Snappy, and Deflate, HGX B200 GPU performs 6X faster than CPUs and 2X faster than H100 for query benchmarks.

Sustainable Computing

By adopting sustainable computing practices, data centers can lower their carbon footprint and energy consumption while improving their bottom line. The goal of sustainable computing can be realized with efficiency gains using accelerated computing with HGX. For LLM inference performance, HGX B200 improves energy efficiency by 12X and lowers costs by 12X compared to H100.

Technical specifications

The following table lists the NVIDIA B200 GPU specifications.

Table 2. GF	PU specifications
-------------	-------------------

Specification	NVIDIA B200
Form Factor	SXM6
FP64	40 teraFLOPS
FP64 Tensor Core	40 teraFLOPS
FP32	80 teraFLOPS
TF32 Tensor Core	1.1 / 2.2 petaFLOPS*
BFLOAT16 Tensor	2.25 / 4.5 petaFLOPS*
FP16 Tensor Core	2.25 / 4.5 petaFLOPS*
FP8 Tensor Core	4.5 / 9 petaFLOPS*
INT8 Tensor Core	4.5 / 9 petaOPS*
GPU Memory	180 GB HBM3e
GPU Memory Bandwidth	8 TB/s
Total Graphics Power (TGP)	1000W
Multi-Instance GPUs	Up to 7 MIGs @ 23 GB
Interconnect	NVLink: 900 GB/s PCle Gen5: 128 GB/s

* Without / with structural sparsity enabled

Server support

The following tables list the ThinkSystem servers that are compatible.

			٩MI	D V3 2S Intel V3/V4							4 In	S 8 tel '	S V3	N N V	/luli lod '3/V	ti e '4	1S V3				
Part Number	Description	SR635 V3 (7D9H / 7D9G)	SR655 V3 (7D9F / 7D9E)	SR645 V3 (7D9D / 7D9C)	SR665 V3 (7D9B / 7D9A)	ST650 V3 (7D7B / 7D7A)	SR630 V3 (7D72 / 7D73)	SR650 V3 (7D75 / 7D76)	SR630 V4 (7DG8 / 7DG9)	SR650 V4 (7DGC / 7DGD)	SR650a V4 (7DGC / 7DGD)	SR850 V3 (7D97 / 7D96)	SR860 V3 (7D94 / 7D93)	SR950 V3 (7DC5 / 7DC4)	SD535 V3 (7DD8 / 7DD1)	SD530 V3 (7DDA / 7DD3)	SD550 V3 (7DD9 / 7DD2)	ST45 V3 (7DH4 / 7DH5)	ST50 V3 (7DF4 / 7DF3)	ST250 V3 (7DCF / 7DCE)	SR250 V3 (7DCM / 7DCL)
C696	ThinkSystem NVIDIA HGX B200 1000W 180GB 8-GPU Liquid-Cooled Board	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	Ν	N	N

Table 4. Server support (Part 2 of 4)

			GPU Rich			Edge					Super Computing								
Part Number	Description	SR670 V2 (7Z22 / 7Z23)	SR675 V3 (7D9Q / 7D9R)	SR680a V3 (7DHE)	SR685a V3 (7DHC)	SR780a V3 (7DJ5)	SE100 (7DGR)	SE350 (7Z46 / 7D1X)	SE350 V2 (7DA9)	SE360 V2 (7DAM)	SE450 (7D8T)	SE455 V3 (7DBY)	SC750 V4 (7DDJ)	SC777 V4 (7DKA)	SD665 V3 (7D9P)	SD665-N V3 (7DAZ)	SD650 V3 (7D7M)	SD650-I V3 (7D7L)	SD650-N V3 (7D7N)
C696	ThinkSystem NVIDIA HGX B200 1000W 180GB 8-GPU Liquid-Cooled Board	Ν	Ν	Ν	Ν	1	Ν	N	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν

Table 5. Server support (Part 3 of 4)

		15	S In V2	tel	2S Inte V2				A	MD	V1		Dense V2				4 V	8S	
Part Number	Description	ST50 V2 (7D8K / 7D8J)	ST250 V2 (7D8G / 7D8F)	SR250 V2 (7D7R / 7D7Q)	ST650 V2 (7Z75 / 7Z74)	SR630 V2 (7Z70 / 7Z71)	SR650 V2 (7Z72 / 7Z73)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR655 Client OS	SR645 (7D2Y / 7D2X)	SR665 (7D2W / 7D2V)	SD630 V2 (7D1K)	SD650 V2 (7D1M)	SD650-N V2 (7D1N)	SN550 V2 (7Z69)	SR850 V2 (7D31 / 7D32)	SR860 V2 (7259 / 7260)	SR950 (7X11 / 7X12)
C696	ThinkSystem NVIDIA HGX B200 1000W 180GB 8-GPU Liquid-Cooled Board	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν

Table 6. Server support (Part 4 of 4)

		4S V1			1S Intel V1				2S Intel V1								D	/1		
Part Number	Description	SR850 (7X18 / 7X19)	SR850P (7D2F / 2D2G)	SR860 (7X69 / 7X70)	ST50 (7Y48 / 7Y50)	ST250 (7Y45 / 7Y46)	SR150 (7Y54)	SR250 (7Y52 / 7Y51)	ST550 (7X09 / 7X10)	SR530 (7X07 / 7X08)	SR550 (7X03 / 7X04)	SR570 (7Y02 / 7Y03)	SR590 (7X98 / 7X99)	SR630 (7X01 / 7X02)	SR650 (7X05 / 7X06)	SR670 (7Y36 / 7Y37)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
C696	ThinkSystem NVIDIA HGX B200 1000W 180GB 8-GPU Liquid-Cooled Board	N	N	N	Ν	Ν	Ν	N	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	N	Ν

Operating system support

For SXM GPUs, operating system support is based on that of the supported servers.

See the SR780a V3 server product guide for details: https://lenovopress.lenovo.com/lp1980-thinksystemsr780a-v3-server

NVIDIA GPU software

This section lists the NVIDIA software that is available from Lenovo.

- NVIDIA AI Enterprise Software
- NVIDIA HPC Compiler Software

As listed in Table 1 in the Part number information section, the PCIe adapter version of the H200 includes a 5year subscription to NVIDIA AI Enterprise Software (NVAIE).

NVIDIA AI Enterprise Software

Lenovo offers the NVIDIA AI Enterprise (NVAIE) cloud-native enterprise software. NVIDIA AI Enterprise is an end-to-end, cloud-native suite of AI and data analytics software, optimized, certified, and supported by NVIDIA to run on VMware vSphere and bare-metal with NVIDIA-Certified Systems[™]. It includes key enabling technologies from NVIDIA for rapid deployment, management, and scaling of AI workloads in the modern hybrid cloud.

NVIDIA AI Enterprise is licensed on a per-GPU basis. NVIDIA AI Enterprise products can be purchased as either a perpetual license with support services, or as an annual or multi-year subscription.

- The perpetual license provides the right to use the NVIDIA AI Enterprise software indefinitely, with no expiration. NVIDIA AI Enterprise with perpetual licenses must be purchased in conjunction with one-year, three-year, or five-year support services. A one-year support service is also available for renewals.
- The subscription offerings are an affordable option to allow IT departments to better manage the flexibility of license volumes. NVIDIA AI Enterprise software products with subscription includes support services for the duration of the software's subscription license

The features of NVIDIA AI Enterprise Software are listed in the following table.

Features	Supported in NVIDIA AI Enterprise
Per GPU Licensing	Yes
Compute Virtualization	Supported
Windows Guest OS Support	No support
Linux Guest OS Support	Supported
Maximum Displays	1
Maximum Resolution	4096 x 2160 (4K)
OpenGL and Vulkan	In-situ Graphics only
CUDA and OpenCL Support	Supported
ECC and Page Retirement	Supported
MIG GPU Support	Supported
Multi-vGPU	Supported
NVIDIA GPUDirect	Supported
Peer-to-Peer over NVLink	Supported
GPU Pass Through Support	Supported
Baremetal Support	Supported
Al and Data Science applications and Frameworks	Supported
Cloud Native ready	Supported

Table 7. Features of NVIDIA AI Enterprise Software (NVAIE)

Note: Maximum 10 concurrent VMs per product license

The following table lists the ordering part numbers and feature codes.

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
AI Enterprise Pe	rpetual License		
7S02001BWW	S6YY	731-AI7004+P3CMI60	NVIDIA AI Enterprise Perpetual License and Support per GPU Socket, 5 Years
7S02001EWW	S6Z1	731-AI7004+P3EDI60	NVIDIA AI Enterprise Perpetual License and Support per GPU Socket, EDU, 5 Years
AI Enterprise Su	bscription License	•	
7S02001FWW	S6Z2	731-AI7003+P3CMI12	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 1 Year
7S02001GWW	S6Z3	731-AI7003+P3CMI36	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 3 Years
7S02001HWW	S6Z4	731-AI7003+P3CMI60	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 5 Years
7S02001JWW	S6Z5	731-AI7003+P3EDI12	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 1 Year
7S02001KWW	S6Z6	731-AI7003+P3EDI36	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 3 Years
7S02001LWW	S6Z7	731-AI7003+P3EDI60	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 5 Years

Table 8. NVIDIA AI Enterprise Software (N	VAIE)
---	-------

Find more information in the NVIDIA AI Enterprise Sizing Guide.

NVIDIA HPC Compiler Software

Part number	Feature code 7S09CTO6WW	Description	
HPC Compiler Support Services			
7S090014WW	S924	NVIDIA HPC Compiler Support Services, 1 Year	
7S090015WW	S925	NVIDIA HPC Compiler Support Services, 3 Years	
7S09002GWW	S9UQ	NVIDIA HPC Compiler Support Services, 5 Years	
7S090016WW	S926	NVIDIA HPC Compiler Support Services, EDU, 1 Year	
7S090017WW	S927	NVIDIA HPC Compiler Support Services, EDU, 3 Years	
7S09002HWW	S9UR	NVIDIA HPC Compiler Support Services, EDU, 5 Years	
7S090018WW	S928	NVIDIA HPC Compiler Support Services - Additional Contact, 1 Year	
7S09002JWW	S9US	NVIDIA HPC Compiler Support Services - Additional Contact, 3 Years	
7S09002KWW	S9UT	NVIDIA HPC Compiler Support Services - Additional Contact, 5 Years	
7S090019WW	S929	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 1 Year	
7S09002LWW	S9UU	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 3 Years	
7S09002MWW	S9UV	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 5 Years	
HPC Compiler Premier Support Services			
7S09001AWW	S92A	NVIDIA HPC Compiler Premier Support Services, 1 Year	

Part number	Feature code 7S09CTO6WW	Description
7S09002NWW	S9UW	NVIDIA HPC Compiler Premier Support Services, 3 Years
7S09002PWW	S9UX	NVIDIA HPC Compiler Premier Support Services, 5 Years
7S09001BWW	S92B	NVIDIA HPC Compiler Premier Support Services, EDU, 1 Year
7S09002QWW	S9UY	NVIDIA HPC Compiler Premier Support Services, EDU, 3 Years
7S09002RWW	S9UZ	NVIDIA HPC Compiler Premier Support Services, EDU, 5 Years
7S09001CWW	S92C	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 1 Year
7S09002SWW	S9V0	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 3 Years
7S09002TWW	S9V1	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 5 Years
7S09001DWW	S92D	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 1 Year
7S09002UWW	S9V2	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 3 Years
7S09002VWW	S9V3	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 5 Years

Regulatory approvals

The NVIDIA B200 GPU has the following regulatory approvals:

- RCM
- BSMI
- CE
- FCC
- ICES
- KCC
- cUL, UL
- VCCI

Warranty

The NVIDIA B200 GPU assumes the server's base warranty and any warranty upgrades.

Seller training courses

The following sales training courses are offered for employees and partners (login required). Courses are listed in date order.

1. VTT AI: Introducing the Lenovo Hybrid AI 285 Platform April 2025

2025-04-30 | 60 minutes | Employees Only

The Lenovo Hybrid AI 285 Platform enables enterprises of all sizes to quickly deploy AI infrastructures supporting use cases as either new greenfield environments or as an extension to current infrastructures. The 285 Platform enables the use of the NVIDIA AI Enterprise software stack. The AI Hybrid 285 platform is the perfect foundation supporting Lenovo Validated Designs.

- Technical overview of the Hybrid AI 285 platform
- AI Hybrid platforms as infrastructure frameworks for LVDs addressing data center-based AI solutions.
- Accelerate AI adoption and reduce deployment risks

Tags: Artificial Intelligence (AI), Nvidia, Technical Sales, Lenovo Hybrid AI 285

Published: 2025-04-30 Length: 60 minutes

Start the training: Employee link: Grow@Lenovo

Course code: DVAI215

2. Lenovo Cloud Architecture VTT: Supercharge Your Enterprise AI with NVIDIA AI Enterprise on Lenovo Hybrid AI Platform

2025-04-17 | 75 minutes | Employees Only

Join us for an in-depth webinar with Justin King, Principal Product Marketing Manager for Enterprise AI exploring the power of NVIDIA AI Enterprise, delivering Generative and Agentic AI outcomes deployed with Lenovo Hybrid AI platform environments.

In today's data-driven landscape, AI is evolving at high speed, with new techniques delivering more accurate responses. Enterprises are seeking not just an understanding but also how they can achieve AI-driven business outcomes.

With this, the demand for secure, scalable, and high-performing AI operations-and the skills to deliver them-is top of mind for many. Learn how NVIDIA AI Enterprise, a comprehensive software suite optimized for NVIDIA GPUs, provides the tools and frameworks, including NVIDIA NIM, NeMo, and Blueprints, to accelerate AI development and deployment while reducing risk-all within the control and security of your Lenovo customer's hybrid AI environment.

Tags: Artificial Intelligence (AI), Cloud, Data Management, Nvidia, Technical Sales

Published: 2025-04-17 Length: 75 minutes

Start the training: Employee link: Grow@Lenovo

Course code: DVCLD221

3. AI VTT: GTC Update and The Lenovo LLM Sizing Guide

2025-03-12 | 86 minutes | Employees Only

Please view this session that is two parts. Part one is Robert Daigle, Director, Global AI Solutions and Hande Sahin-Bahceci, AI Solutions Marketing Leader explaining the upcoming announcements for NVIDIA GTC. Part Two is Sachin Wani, AI Data Scientist explaining the Lenovo LLM Sizing Guide with these topics:

- · Minimum GPU requirements for fine-tuning/training and inference
- · Gathering requirements for the customer's use case
- · LLMs from a technical perspective

Tags: Artificial Intelligence (AI), Technical Sales

Published: 2025-03-12 Length: 86 minutes

Start the training: Employee link: Grow@Lenovo

Course code: DVAI214

4. Partner Technical Webinar - NVIDIA Portfolio 2024-11-06 | 60 minutes | Employees and Partners

In this 60-minute replay, Jason Knudsen of NVIDIA presented the NVIDIA Computing Platform. Jason talked about the full portfolio from GPUs to Networking to AI Enterprise and NIMs.

Tags: Artificial Intelligence (AI), Nvidia

Published: 2024-11-06 Length: 60 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: 110124

5. NVIDIA Data Center GPU Portfolio

2024-09-26 | 11 minutes | Employees and Partners

This course equips Lenovo and partner technical sellers with the knowledge to effectively communicate the positioning of NVIDIA's data center GPU portfolio, enhancing your ability to showcase its key advantages to clients.

Upon completion of this training, you will be familiar with the following:

- Data Center GPUs for AI and HPC
- •Data Center GPUs for Graphics
- •GPU comparisons

Tags: Artificial Intelligence (AI), High-Performance Computing (HPC), Nvidia

Published: 2024-09-26 Length: 11 minutes

Start the training:

Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DAINVD201

6. Q2 Solutions Launch TruScale GPU Next Generation Management in the AI Era Quick Hit 2024-09-10 | 6 minutes | Employees and Partners

This Quick Hit focuses on Lenovo announcing additional ways to help you build, scale, and evolve your customer's private AI faster for improved ROI with TruScale GPU as a Service, AI-driven systems management, and infrastructure transformation services.

Tags: Artificial Intelligence (AI), Services, TruScale

Published: 2024-09-10 Length: 6 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: SXXW2543a

7. VTT AI: The NetApp AIPod with Lenovo for NVIDIA OVX

2024-08-13 | 38 minutes | Employees and Partners

AI, for some organizations, is out of reach, due to cost, integration complexity, and time to deployment. Previously, organizations relied on frequently retraining their LLMs with the latest data, a costly and time-consuming process. The NetApp AIPod with Lenovo for NVIDIA OVX combines NVIDIA-Certified OVX Lenovo ThinkSystem SR675 V3 servers with validated NetApp storage to create a converged infrastructure specifically designed for AI workloads. Using this solution, customers will be able to conduct AI RAG and inferencing operations for use cases like chatbots, knowledge management, and object recognition.

Topics covered in this VTT session include:

- •Where Lenovo fits in the solution
- •NetApp AIPod with Lenovo for NVIDIA OVX Solution Overview
- •Challenges/pain points that this solution solves for enterprises deploying AI
- •Solution value/benefits of the combined NetApp, Lenovo, and NVIDIA OVX-Certified Solution

Tags: Artificial Intelligence (AI), Nvidia, Sales, Technical Sales, ThinkSystem

Published: 2024-08-13 Length: 38 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DVAI206

8. Retrieval Augmented Generation

2024-08-02 | 15 minutes | Employees and Partners

IMPORTANT: If you receive the following error message:

"There is an issue with this slide content. Please contact your administrator", please change your VPN location setting and try again. We are actively working on fixing this issue. Thank you for your understanding!

In this NVIDIA course, Dave Barry, Senior Solutions Architect, talks about a technique known as Retrieval Augmented Generation (RAG). It is a powerful tool for enhancing the accuracy and reliability of Generative AI models with facts fetched from external sources.

This course requires prior knowledge of Generative AI concepts, such as the difference between model training and inference. Please refer to relevant courses within this curriculum.

Course Objectives:

By the end of this training, you should be able to:

- 1. Explain the limitations of large language models to customers
- 2. Articulate the value of RAG to enterprises
- 3. Demo an NVIDIA RAG workflow with a video
- 4. Drive TCO conversations using an authentic use case

Tags: Nvidia

Published: 2024-08-02 Length: 15 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DAINVD107

9. Al Industry Use Cases & Solutions

2024-08-02 | 25 minutes | Employees and Partners

IMPORTANT: If you receive the following error message:

"There is an issue with this slide content. Please contact your administrator", please change your VPN location setting and try again. We are actively working on fixing this issue. Thank you for your understanding!

This NVIDIA course aims to answer the question:

· How does NVIDIA bring AI solutions to market with and through the partner ecosystem?

Course Objectives:

By the end of this training, you should be able to:

- 1. Think of solutions in terms of an industry and use case approach
- 2. Develop solutions that address the industry-specific challenges (with FSI as the illustrative model)
- 3. Engage customers with their conversations and advance deals with stakeholder's concerns in mind
- 4. Replicate NVIDIA's best practices and ecosystem engagement strategies appropriately

Tags: Nvidia

Published: 2024-08-02 Length: 25 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DAINVD105r2

10. GPU Fundamentals

2024-08-02 | 10 minutes | Employees and Partners

IMPORTANT: If you receive the following error message: "There is an issue with this slide content. Please contact your administrator", please change your VPN location setting and try again. We are actively working on fixing this issue. Thank you for your understanding.

This NVIDIA course introduces you to two devices that a computer typically uses to process information – the CPU and the GPU. We'll discuss their differences and look at how the GPU overcomes the limitations of the CPU. We will also talk about the value GPUs bring to modern-day enterprise computing.

Course Objectives:

By the end of this training, you should be able to:

- 1. Distinguish between serial and parallel processing
- 2.Explain what a GPU is and what it does at a high level
- 3.Articulate the value of GPU computing for enterprises
- 4.List three typical GPU-accelerated workloads and a few uses cases
- 5. Recommend the appropriate NVIDIA GPU for its corresponding enterprise computing workloads

Tags: Nvidia

Published: 2024-08-02 Length: 10 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DAINVD103r2

11. Generative Al Overview

2024-08-02 | 17 minutes | Employees and Partners

IMPORTANT: If you receive the following error message:

"There is an issue with this slide content. Please contact your administrator", please change your VPN location setting and try again. We are actively working on fixing this issue. Thank you for your understanding!

Since ChatGPTs debut in November of 2022, it has become clear that Generative AI has the potential to revolutionize many aspects of our personal and professional lives. This NVIDIA course aims to answer questions such as:

• What are the Generative AI market trends?

• What is generative AI and how does it work?

Course Objectives:

By the end of this training, you should be able to:

Discuss the Generative AI market trends and the challenges in this space with your customers.
Explain what Generative AI is and how the technology works to help enterprises to unlock new opportunities for the business.

3. Present a high-level overview of the steps involved in building a Generative AI application.

Tags: Nvidia

Published: 2024-08-02 Length: 17 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DAINVD106r2

12. Key NVIDIA Use Cases for Industry Verticals

2024-08-02 | 32 minutes | Employees and Partners

IMPORTANT: If you receive the following error message: "There is an issue with this slide content. Please contact your administrator", please change your VPN location setting and try again. We are actively working on fixing this issue. Thank you for your understanding.

In this NVIDIA course, you will learn about key AI use cases driving innovation and change across Automotive, Financial Services, Energy, Healthcare, Higher Education, Manufacturing, Retail and Telco industries.

Course Objectives:

By the end of this training, you should be able to:

1. Discuss common AI use cases across a broad range of industry verticals

2. Explain how NVIDIA's AI software stack speeds up time to production for AI projects in multiple industry verticals

Tags: Nvidia

Published: 2024-08-02 Length: 32 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DAINVD108

13. Introduction to Artificial Intelligence

2024-08-02 | 11 minutes | Employees and Partners

IMPORTANT: If you receive the following error message:

"There is an issue with this slide content. Please contact your administrator", please change your VPN location setting and try again. We are actively working on fixing this issue. Thank you for your understanding!

This NVIDIA course aims to answer questions such as:

- What is AI?
- · Why are enterprises so interested in it?
- How does AI happen?
- Why are GPUs so important for it?
- What does a good AI solution look like?

Course Objectives:

By the end of this training, you should be able to:

- 1. Describe AI on a high level and list a few common enterprise use cases
- 2. List how enterprises benefit from AI
- 3. Distinguish between Training and Inference
- 4. Say how GPUs address known bottlenecks in a typical AI pipeline
- 5. Tell a customer why NVIDIA's AI solutions are well-respected in the market

Tags: Nvidia

Published: 2024-08-02 Length: 11 minutes

Start the training:

Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DAINVD104r2

14. Partner Technical Webinar - NVIDIA Smart Spaces

2024-07-24 | 60 minutes | Employees and Partners

In this 60-minute replay, Alex Pazos, NVIDIA BDM for Smart Spaces, reviewed the NVIDIA AI for Smart Spaces framework and use cases. Alex reviewed the Metropolus Framework and the Smart Spaces ecosystem. Then he reviewed several use cases including sports stadiums, warehouses, airports, and roadways.

Tags: Artificial Intelligence (AI)

Published: 2024-07-24 Length: 60 minutes

Start the training:

Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: 071924

15. Guidance for Selling NVIDIA Products at Lenovo for ISG

2024-07-01 | 25 minutes | Employees and Partners

This course gives key talking points about the Lenovo and NVIDIA partnership in the Data Center. Details are included on where to find the products that are included in the partnership and what to do if NVIDIA products are needed that are not included in the partnership. Contact information is included if help is needed in choosing which product is best for your customer. At the end of this session sellers should be able to explain the Lenovo and NVIDIA partnership, describe the products Lenovo can sell through the partnership with NVIDIA, help a customer purchase other NVIDIA product, and get assistance with choosing NVIDIA products to fit customer needs.

Tags: Artificial Intelligence (AI), Nvidia

Published: 2024-07-01 Length: 25 minutes

Start the training: Employee link: Grow@Lenovo Partner link: Lenovo Partner Learning

Course code: DNVIS102

16. Think Al Weekly: Lenovo Al PCs & Al Workstations

2024-05-23 | 60 minutes | Employees Only

Join Mike Leach, Sr. Manager, Workstations Solutions and Pooja Sathe, Director Commercial AI PCs as they discuss why Lenovo AI Developer Workstations and AI PCs are the most powerful, where they fit into the device to cloud ecosystem, and this week's Microsoft announcement, Copilot+PC

Tags: Artificial Intelligence (AI), ThinkStation

Published: 2024-05-23 Length: 60 minutes

Start the training:

Employee link: Grow@Lenovo

Course code: DTAIW105

17. VTT Cloud Architecture: NVIDIA Using Cloud for GPUs and AI

2024-05-22 | 60 minutes | Employees Only

Join JD Dupont, NVIDIA Head of Americas Sales, Lenovo partnership and Veer Mehta, NVIDIA Solution Architect on an interactive discussion about cloud to edge, designing cloud Solutions with NVIDIA GPUs and minimizing private\hybrid cloud OPEX with GPUs. Discover how you can use what is done at big public cloud providers for your customers. We will also walk through use cases and see a demo you can use to help your customers.

Tags: Artificial Intelligence (AI), Cloud, Nvidia, Software Defined Infrastruture (SDI), Technical Sales

Published: 2024-05-22 Length: 60 minutes

Start the training: Employee link: Grow@Lenovo

Course code: DVCLD212

Related links

For more information, refer to these documents:

- ThinkSystem and ThinkAgile GPU Summary: https://lenovopress.lenovo.com/lp0768-thinksystem-thinkagile-gpu-summary
- ServerProven compatibility: https://serverproven.lenovo.com/
- NVIDIA HGX Platform product page: https://www.nvidia.com/en-us/data-center/hgx//
- ThinkSystem SR780a V3 product guide https://lenovopress.lenovo.com/lp1980-thinksystem-sr780a-v3-server

Related product families

Product families related to this document are the following:

• GPU adapters

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc. 8001 Development Drive Morrisville, NC 27560 U.S.A. Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2025. All rights reserved.

This document, LP2226, was created or updated on May 28, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at: https://lenovopress.lenovo.com/LP2226
- Send your comments in an e-mail to: comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP2226.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both: Lenovo® Neptune® ServerProven® ThinkAgile® ThinkSystem®

The following terms are trademarks of other companies:

AMD is a trademark of Advanced Micro Devices, Inc.

Intel® is a trademark of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Windows® is a trademark of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.