Catalog

NVIDIA Training Course Catalog

August 2023

Introduction

NVIDIA offers training for diverse needs, giving individuals and teams across organizations what they need to advance their knowledge in AI, accelerated computing, data science, data center administration, graphics and simulation, networking, and more.

With access to high-performance computing, you'll learn how to train, optimize, and deploy neural networks using the latest deep learning tools, frameworks, and SDKs. You'll also learn how to assess, parallelize, optimize, and deploy GPU-accelerated computing applications.

Our training program offers both self-paced online courses and instructor-led, prescheduled workshops. The self-paced courses range from 10 minutes to 8 hours and guide you through applying a specific technology, setting up a project, or administering solutions in a data center. Instructor-led workshops and boot camps go deeper into topic areas, teaching you how to implement a project or solution from end to end. Both types of courses give you valuable hands-on experience using the latest technologies.

Why Choose NVIDIA for Training?

- Learn how to build deep learning and accelerated computing applications for industries such as healthcare, robotics, autonomous driving, manufacturing, and more.
- Gain hands-on experience with the most widely used, industry-standard platforms including software, hardware, tools, and frameworks. Each student will have access to a fully configured, GPU-accelerated server in the cloud or access to NVIDIA solutions in our training lab.
- Become proficient in administering NVIDIA hardware and software solutions such as DGX™, InfiniBand, Cumulus, NVIDIA AI Enterprise, and more.
- Access instructor-led workshops and online courses from anywhere using just a laptop and internet connection.
- Acquire real-world expertise through content designed in collaboration with industry leaders such as Children's Hospital of Los Angeles, Mayo Clinic, and PwC.
- Earn NVIDIA certifications and course completion certificates to indicate subject matter competency and support your career growth.

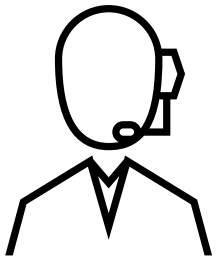

For team training, contact an **NVIDIA training advisor**, who will work with you to create a customized plan that addresses your team's specific training needs and is aligned to your business objectives and priorities.

Table of Contents

Instructor-Led Workshops for Developers

Online, Self-Paced Courses for Developers

Data Science

Deep Learning

Graphics and Simulation

Infrastructure

Instructor-Led Workshops for Administrators

Online, Self-Paced Courses for Administrators

Management

Network

RDMA

Certifications

Instructor-Led Workshops for Developers

Accelerated Computing

Workshop Name	Description	Prerequisites	Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
Accelerating CUDA® C++ Applications With Multiple GPUs	Discover how to write CUDA C++ applications that efficiently and correctly use all available GPUs in a single node, dramatically improving the performance of applications and making the most cost-effective use of systems with multiple GPUs. > Learn More	Professional experience programming CUDA C/C++ applications, including the use of the NVIDIA CUDA Compiler (NVCC), kernel launches, grid-stride loops, host-to-device and device-to-host memory transfers, and CUDA error handling. Familiarity with the Linux command line and experience using makefiles to compile C/C++ code.	CUDA C++, NVCC, Nsight Systems	English, Simplified Chinese	8 hours	$500 (excludes tax, if applicable)	Yes
Fundamentals of Accelerated Computing With CUDA C/C++	Learn how to accelerate and optimize existing C/C++ CPU-only applications to apply the power of GPUs using the most essential CUDA techniques and the NVIDIA Nsight Systems profiler. > Learn More	Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of CUDA programming is assumed.	NVIDIA Nsight Systems, nsys	English, Korean, Japanese, Simplified Chinese, Traditional Chinese	8 hours	$500 (excludes tax, if applicable)	Yes
Fundamentals of Accelerated Computing With CUDA Python	Explore how to use Numba—the just-in-time, type-specializing Python function compiler—to create and launch CUDA kernels to accelerate Python programs on massively parallel NVIDIA GPUs. > Learn More	Basic Python competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. Also, must have NumPy competency, including the use of ndarrays and ufuncs.	CUDA, Python, Numba, NumPy	English, Simplified Chinese, Traditional Chinese	8 hours	$500 (excludes tax, if applicable)	Yes
Fundamentals of Accelerated Computing With OpenACC®	Find out how to write and configure code parallelization with OpenACC, optimize memory movements between the CPU and GPU accelerator, and apply the techniques to accelerate a CPU-only Laplace heat equation to achieve performance gains. > Learn More	Basic C/C++ or Fortran competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of GPU programming is assumed.	NVIDIA Nsight, OpenACC	English	8 hours	$500 (excludes tax, if applicable)	Yes

Back

Workshop Name	Description	Prerequisites			
Scaling CUDA C++ Applications to Multiple Nodes	Learn the tools and techniques needed to write CUDA C++ applications that can scale efficiently to clusters of NVIDIA GPUs. > Learn More	Intermediate experience writing CUDA C/C++ applications.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	C++, CUDA, MPI, NVSHMEM	English, Simplified Chinese	8 hours	$500 (excludes tax, if applicable)	Yes

Data Science

Accelerating Data Engineering Pipelines	Explore how to employ advanced data engineering tools and techniques with GPUs to significantly improve data engineering pipelines. > Learn More	Intermediate knowledge of **Python** (list comprehension, objects). Familiarity with **pandas** and **introductory statistics** (mean, median, mode) a plus.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	pandas, cuDF, Dask, NVTabular, Plotly	English	8 hours	$500 (excludes tax, if applicable)	Yes
Fundamentals of Accelerated Data Science	Learn how to perform multiple analysis tasks on large datasets using NVIDIA RAPIDS™, a collection of data science libraries that allows end-to-end GPU acceleration for data science workflows. > Learn More	Professional data science experience with Python, including proficiency in pandas and NumPy. Also, must have familiarity with common machine learning algorithms, including XGBoost, linear regression, DBSCAN, K-Means, and SSSP.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	RAPIDS, cuDF, XGBoost, cuML, cuGraph, Dask, cuPy, pandas, NumPy, Bokeh	English, Traditional Chinese, Japanese	8 hours	$500 (excludes tax, if applicable)	Yes

Deep Learning

Applications of AI for Anomaly Detection	Learn to detect anomalies in large datasets to identify network intrusions using supervised and unsupervised machine learning techniques, such as accelerated XGBoost, autoencoders, and generative adversarial networks (GANs). > Learn More	Experience with convolutional neural networks (CNNs) and Python.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	NVIDIA RAPIDS, XGBoost, TensorFlow, Keras, pandas, autoencoders, GANs	English	8 hours	$500 (excludes tax, if applicable)	Yes
Applications of AI for Predictive Maintenance	Discover how to identify anomalies and failures in time-series data, estimate the remaining useful life of the corresponding parts, and use this information to map anomalies to failure conditions. > Learn More	Experience with Python and deep networks.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Python, TensorFlow, Keras, XGBoost, RAPIDS, cuDF, long short-term memory (LSTM), autoencoders	English, Simplified Chinese	8 hours	$500 (excludes tax, if applicable)	Yes

Back

Workshop Name	Description	Prerequisites
Building AI-Based Cybersecurity Pipelines	Traditional cybersecurity methods include creating barriers around your infrastructure to protect it from intruders. However, as enterprises continue to digitally transform, they're faced with a proliferation of devices, more sophisticated cybersecurity attacks, and an incredibly vast network of data to protect—which means new cybersecurity methodologies must be explored. An alternative approach is to address cybersecurity as a data science problem: Better understand all the users and activities across your network so that you can identify which transactions are typical and which are potentially nefarious. The **NVIDIA Morpheus AI framework** lets cybersecurity developers and practitioners harness the power of GPU computing to implement cybersecurity solutions that perform on a scale never before possible. With Morpheus, cybersecurity developers can create optimized AI pipelines for filtering, processing, and classifying large volumes of real-time data. Bringing a new level of information security to data centers, Morpheus enables dynamic protection, real-time telemetry, and adaptive defenses for detecting and remediating cybersecurity threats. **> Learn More**	> Familiarity with defensive cybersecurity themes. > Professional data science and/or data analysis experience. > Competency with the Python programming language. > Competency with the Linux command line.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA Morpheus, NVIDIA Triton™ Inference Server, NVIDIA RAPIDS, CLX, Helm, Kubernetes	English	8 hours	$500 (excludes tax, if applicable)	Yes

Workshop Name	Description	Prerequisites
Building Conversational AI Applications V2.0	Discover how to quickly build and deploy production-quality speech AI applications with real-time transcription and natural language processing capabilities. **> Learn More**	Experience with Python coding and use of library functions and parameters. Also, a fundamental understanding of a deep learning framework, such as TensorFlow, PyTorch, or Keras, and a basic understanding of neural networks.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA Riva, NVIDIA TAO Toolkit, Kubernetes	English	8 hours	$500 (excludes tax, if applicable)	Yes

Workshop Name	Description	Prerequisites
Building Intelligent Recommender Systems	Explore the fundamental tools and techniques for building highly effective recommender systems, as well as how to deploy GPU-accelerated solutions for real-time recommendations. **> Learn More**	Intermediate knowledge of Python, including an understanding of list comprehension. Data science experience using Python and familiarity with NumPy and matrix mathematics.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
CuDF, CuPy, TensorFlow 2, NVIDIA Merlin™, NVTabular, and NVIDIA Triton Inference Server	English	8 hours	$500 (excludes tax, if applicable)	Yes

Back

Workshop Name	Description	Prerequisites
Building Transformer-Based Natural Language Processing	In this workshop, you'll learn how Transformers are used as the building blocks of modern large language models (LLMs). You'll then use these models for various NLP tasks, including text classification, named-entity recognition (NER), author attribution, and question answering. You'll also learn how to analyze various model features, constraints, and characteristics to determine which model is best suited for a particular use case based on metrics, domain specificity, and available resources. > **Learn More**	Experience with Python coding and use of library functions and parameters. Fundamental understanding of a deep learning framework, such as TensorFlow, PyTorch, or Keras. And basic understanding of neural networks.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
PyTorch, pandas, NVIDIA NeMo, NVIDIA Triton Inference Server	English, Simplified Chinese	8 hours	$500 (excludes tax, if applicable)	Yes

Workshop Name	Description	Prerequisites
Computer Vision for Industrial Inspection	In this workshop, you'll learn how to quickly develop and deploy a machine learning model that uses deep learning for computer vision to perform defect classification and other visual recognition tasks. Using NVIDIA's own real production dataset as an example, this workshop illustrates how the solution can be easily applied to a variety of manufacturing and industrial inspection use cases. > **Learn More**	> Experience with Python; basic understanding of data processing and deep learning > To gain experience with Python, we suggest this **Python tutorial** > For a basic understanding of data processing and deep learning, we suggest **Fundamentals of Deep Learning.**

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
Python, pandas, DALI, NVIDIA TAO Toolkit, NVIDIA TensorRT™, and NVIDIA Triton Inference Server	English, Simplified Chinese	8 hours	$500 (excludes tax, if applicable)	Yes

Workshop Name	Description	Prerequisites
Data Parallelism: How to Train Deep Learning Models on Multiple GPUs	This workshop teaches you techniques for data-parallel deep learning training on multiple GPUs to shorten the training time required for data-intensive applications. Working with deep learning tools, frameworks, and workflows to perform neural network training, you'll learn how to decrease model training time by distributing data to multiple GPUs, while retaining the accuracy of training on a single GPU. > **Learn More**	Experience with deep learning training using Python. See the Fundamentals of Deep Learning self-paced course **here**.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
PyTorch, PyTorch Distributed Data Parallel, NCCL	English, Simplified Chinese	8 hours	$500 (excludes tax, if applicable)	Yes

Back

Workshop Name	Description	Prerequisites
Fundamentals of Deep Learning	Learn how deep learning (DL) works through hands-on exercises in computer vision and natural language processing (NLP). You'll train deep learning models from scratch and pick up tricks and tools for achieving highly accurate results along the way. You'll also learn to leverage freely available, state-of-the-art pretrained models to save time and get your deep learning application up and running quickly. > **Learn More**	An understanding of fundamental programming concepts in **Python 3**, such as functions, loops, dictionaries, and arrays. Also, familiarity with **pandas data structures** and an understanding of how to compute a **regression line**. > Suggested materials to satisfy prerequisites: **Python Beginner's Guide**

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
Tensorflow, Keras, Pandas, NumPy	English, Simplified Chinese, Japanese	8 hours	$500 (excludes tax, if applicable)	Yes

Workshop Name	Description	Prerequisites
Generative AI with Diffusion Models	Get started with gen AI application development with this hands-on course where students will learn how to build a text-to-image generative AI application using the latest techniques. Generate images with diffusion models and refine the output with various optimizations. Build a denoising diffusion model from the U-Net architecture to context embeddings for greater user control. > **Learn More**	> Good understanding of **PyTorch** > Good understanding of **deep learning**

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
PyTorch, CLIP	English	8 hours	$500 (excludes tax, if applicable)	Yes

Workshop Name	Description	Prerequisites
Model Parallelism: Building and Deploying Large Neural Networks	In this workshop, you'll learn how to scale training and deployment of LLMs and neural networks across multiple nodes, use various forms of model parallelism to overcome the challenges associated with large-model memory footprint, capture and understand training performance characteristics to optimize model architecture and deploy very large multi-GPU, multi-node models to production using NVIDIA Triton™ Inference Server. > **Learn More**	> Good understanding of **PyTorch**, **deep learning**, and **data parallel** training concepts > Practice with **multi-GPU training** and **natural language processing** is useful, but optional.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
PyTorch, Megatron-LM, DeepSpeed, Slurm, NVIDIA Triton Inference Server, NVIDIA Nsight	English, Korean, Simplified Chinese	8 hours	$500 (excludes tax, if applicable)	Yes

Back

Graphics and Simulation

Bootstrapping Computer Vision Models with Synthetic Data	Learn how to use NVIDIA Omniverse Replicator, a core Omniverse extension, to accelerate the development of computer vision models. Generate accurate, photorealistic, physics-conforming synthetic data to ease the expensive, time-consuming task of labeling real-world data. Omniverse Replicator accelerates AI development at scale and reduces time to production. > Learn More	> Intermediate understanding of Python (including classes, objects, and decorators). > Basic understanding of Machine Learning and Deep Learning concepts and pipelines.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Omniverse Replicator, Omniverse Defect Extension	English	8 hours	$500 (excludes tax, if applicable)	Yes

Online, Self-Paced Courses for Developers

Course Name	Description	Prerequisites			
Accelerated Computing Fundamentals					
Accelerating CUDA C++ Applications With Concurrent Streams	Discover how to improve performance for your CUDA C/C++ applications by overlapping memory transfers to and from the GPU with computations on the GPU. > Learn More	Professional experience programming CUDA C/C++ applications, including the use of the nvcc compiler, kernel launches, grid-stride loops, host-to-device and device-to-host memory transfers, and CUDA error handling; Experience using Makefiles to compile C/C++ code.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
		English	4 hours	$30 (excludes tax, if applicable)	Yes
An Even Easier Introduction to CUDA	Learn the basics of writing parallel CUDA kernels to run on NVIDIA GPUs. > Learn More	Competency writing applications in CUDA C/C++.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	C/C++	English	1 hour	Free	N/A
Fundamentals of Accelerated Computing With CUDA C/C++	Discover how to accelerate and optimize existing C/C++ CPU-only applications to leverage the power of GPUs using the most essential CUDA techniques and the Nsight Systems profiler. > Learn More	Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of CUDA programming is assumed.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	C/C++, CUDA	English, Japanese, Korean, Simplified Chinese, Traditional Chinese	8 hours	$90 (excludes tax, if applicable)	Yes

Back

Course Name	Description	Prerequisites			
Fundamentals of Accelerated Computing With CUDA Python	Explore how to use Numba—the just-in-time, type-specializing Python function compiler—to create and launch CUDA kernels to accelerate Python programs on massively parallel NVIDIA GPUs. **> Learn More**	Basic Python competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. Also, must have NumPy competency, including the use of ndarrays and ufuncs.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	CUDA, Python, Numba, NumPy	English, Simplified Chinese, Traditional Chinese	8 hours	$90 (excludes tax, if applicable)	Yes
Fundamentals of Accelerated Computing With OpenACC	Find out how to build and optimize accelerated heterogeneous applications on multiple GPU clusters using a combination of OpenACC, CUDA-aware MPI, and NVIDIA profiling tools. **> Learn More**	Basic experience with C/C++			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	OpenACC, C/C++	English	8 hours	$90 (excludes tax, if applicable)	N/A
GPU Acceleration With the C++ Standard Library	Learn to write simple, portable, parallel-first applications using only standard C++ language features that can be compiled without modification to take advantage of NVIDIA GPU-accelerated environments. **> Learn More**	Beginner-level experience with **C++11**. Comfort working with **C++ lambdas and standard library algorithms**.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	C++, **NVIDIA HPC SDK**	English	2 hours	$30 (excludes tax, if applicable)	N/A
High-Performance Computing With Containers	Learn how to reduce complexity and improve portability and efficiency of your code by using a containerized environment for HPC application development. **> Learn More**	Proficiency in programming in C/C++ and professional experience working on HPC applications.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Docker, Singularity, HPC Container Maker (HPCCM), C/C++	English	2 hours	$30 (excludes tax, if applicable)	N/A

Back

Course Name	Description	Prerequisites
Optimizing CUDA Machine Learning Codes With NVIDIA Nsight™ Profiling Tools	NVIDIA Developer Tools are a collection of applications, spanning desktop and mobile targets, that enable developers to build, debug, profile, and develop class-leading and cutting-edge software using the latest visual computing hardware from NVIDIA. In this course, you'll learn the effective use of two powerful NVIDIA developer tools: **Nsight Systems** and **Nsight Compute**. Nsight Systems provide developers with a system-wide visualization of an application's performance. Developers can optimize bottlenecks to scale efficiently across any number or size of CPU and GPU—from large servers to the smallest systems on chip. Nsight Compute is an interactive kernel profiler for CUDA applications. It provides detailed performance metrics and API debugging via a user interface and command-line tool. By the time you complete this course, you'll be able to use Nsight Systems and Nsight Compute to analyze and optimize CUDA applications. Following best practices, you'll begin by using Nsight Systems to analyze overall application structure and explore parallelization opportunities before turning to Nsight Compute to analyze and optimize individual CUDA kernels. **> Learn More**	Familiarity with machine learning applications using CUDA. We suggest **Fundamentals of Accelerated Computing with CUDA C/C++.**

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA Nsight Systems, **NVIDIA Nsight Compute**	English	2 hours	$30 (excludes tax, if applicable)	N/A

Course Name	Description	Prerequisites
Scaling GPU-Accelerated Applications With the C++ Standard Library	In this interactive, hands-on workshop, which is the followup to GPU Acceleration With the C++ Standard Library, you'll learn how to write scalable, GPU-accelerated, hybrid applications using C++ standard language features alongside MPI. **> Learn More**	Beginner-level experience with C++11; comfort working with C++ lambdas and standard library algorithms; experience developing C++/MPI hybrid applications that require inter-rank communication; comfort working with C++ concurrency primitives such as std::thread, std::barrier, and andstd::thread.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
C++, NVIDIA HPC SDK, MPI	English	2 hours	$30 (excludes tax, if applicable)	N/A

Course Name	Description	Prerequisites
Scaling Workloads Across Multiple GPUs With CUDA C++	Learn how to build robust and efficient CUDA C++ applications that can take advantage of all available GPUs on a single node. **> Learn More**	Competency writing applications in CUDA C/C++.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
C/C++, accelerated computing, CUDA	English	4 hours	$30 (excludes tax, if applicable)	Yes

Back

Course Name	Description	Prerequisites			
Data Science					
Accelerating End-to-End Data Science Workflows	Explore how to perform multiple analysis tasks on large datasets using RAPIDS, a collection of data science libraries that allows end-to-end GPU acceleration for data science workflows. > Learn More	Experience with Python, ideally including pandas and NumPy.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	RAPIDS, cuDF, cuML, cuGraph, Apache Arrow	English, Simplified Chinese	6 hours	$90 (excludes tax, if applicable)	Yes
Deep Learning					
Building a Brain in 10 Minutes	This one-click notebook explores the biological and psychological inspirations for the world's first neural networks. > Learn More	An understanding of fundamental programming concepts in Python 3 such as functions, loops, dictionaries, and arrays.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	N/A	English	10 minutes	Free	N/A
Building Real-Time Video AI Applications	Gain the knowledge and skills needed to enable the real-time transformation of raw video data from widely deployed camera sensors into deep learning-based insights. > Learn More	Competency in the Python 3, programming language, some experience manipulating data using pandas DataFrames, and familiarity with deep networks (specifically variations of CNNs).			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	NVIDIA DeepStream, NVIDIA TAO Toolkit, and NVIDIA TensorRT	English, Simplified Chinese	8 hours	$90.00 (excludes tax, if applicable)	N/A
Building Video AI Applications at the Edge on NVIDIA® Jetson Nano™	Use JupyterLab notebooks and Python application samples on your Jetson Nano to build new projects that extract meaningful insights from video streams through deep learning video analytics. > Learn More	Basic familiarity with the Linux command line and an understanding of fundamental programming concepts in Python 3, such as functions, loops, dictionaries, and arrays.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	DeepStream, TensorRT, Jetson Nano, Python	English, Simplified Chinese	8 hours	Free (hardware required)	N/A
Deploying a Model for Inference at Production Scale	Learn how to deploy your own machine learning models on a GPU server. > Learn More	Familiarity with at least one machine learning framework, such as PyTorch, TensorFlow, ONNX, or TensorRT.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	NVIDIA Triton	English	4 hours	$30 (excludes tax, if applicable)	N/A

Back

Course Name	Description	Prerequisites			
Digital Fingerprinting With Morpheus	In this course, you'll get hands-on experience developing and deploying the NVIDIA digital fingerprinting AI workflow that enables 100% data visibility and drastically reduces the time to detect threats. You'll also hear from cybersecurity experts from a variety of institutions about how to use NVIDIA AI frameworks and tools to architect cybersecurity solutions. > **Learn More**	This tutorial doesn't have any prerequisites, but familiarity with defensive cybersecurity themes and the Linux command line are a plus.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	NVIDIA Morpheus AI framework, **NVIDIA Triton Inference Server**	English	1 hour	Free	N/A
Disaster Risk Monitoring Using Satellite Imagery	Learn how to build and deploy a deep learning model to automate the detection of flood events using satellite imagery. This workflow can be applied to lower the cost, improve the efficiency, and significantly enhance the effectiveness of various natural disaster management use cases. > **Learn More**	> Competency in the **Python 3** programming language. > Basic understanding of machine learning and deep learning concepts, specifically variations of convolutional neural networks (**CNNs**), and pipelines. > Interest in understanding how to manipulate satellite imagery using modern methods.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	NVIDIA DALI®, the NVIDIA TAO Toolkit, NVIDIA TensorRT, NVIDIA Triton Inference Server	English, Simplified Chinese	10 hours	Free	Yes
Generative AI Explained	Generative AI describes technologies that are used to generate new content based on a variety of inputs. In this course, you will learn Generative AI concepts, applications, as well as the challenges and opportunities in this exciting field. > **Learn More**	Basic understanding of Machine Learning and Deep Learning concepts			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	N/A	English	2 hours	Free	N/A
Get Started With Highly Accurate Custom ASR for Speech AI	Learn to build, train, fine-tune, and deploy a GPU-accelerated automatic speech recognition service with NVIDIA Riva that includes customized features. > **Learn More**	Basic understanding of machine learning and deep learning concepts and pipelines. In addition, this lab requires that the user have an NVIDIA NGC account and API key. To fulfill this requirement, please **register and activate a free NGC account.** > **Generate your NGC API key and save it in a safe location In**			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Riva, TAO Toolkit, Kubernetes	English	2 hours	$30 (excludes tax, if applicable)	N/A

Back

Course Name	Description	Prerequisites			
Getting Started With AI on Jetson Nano	Discover how to build a deep learning classification project with computer vision models using the NVIDIA Jetson Nano Developer Kit. > **Learn More**	Basic familiarity with Python (helpful, not required).			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	PyTorch, Jetson Nano	English, Simplified Chinese, Japanese, Korean	8 hours	Free (hardware required)	Yes
Getting Started With Deep Learning	Explore the fundamentals of deep learning by training neural networks and using results to improve performance and capabilities. > **Learn More**	> An understanding of fundamental programming concepts in **Python 3**, such as functions, loops, dictionaries, and arrays. > Familiarity with **pandas data structures** and an understanding of how to compute a **regression line** > Suggested materials to satisfy prerequisites: **Python Beginner's Guide**			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	TensorFlow 2 with Keras, pandas	English, Simplified Chinese	8 hours	$90 (excludes tax, if applicable)	Yes
Getting Started With Image Segmentation	Learn how to categorize segments of an image. > **Learn More**	Basic experience training neural networks.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	TensorFlow 2 with Keras	English	2 hours	$30 (excludes tax, if applicable)	N/A
Integrating Sensors With NVIDIA DRIVE	Find out how to integrate automotive sensors into your applications using NVIDIA DRIVE. > **Learn More**	Basic experience in C++ and Linux terminal commands.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	C++, NVIDIA DriveWorks	English	2 hours	$30 (excludes tax, if applicable)	N/A
Introduction to Graph Neural Networks	Learn the basic concepts, models, and applications of graph neural networks. > **Learn More**	Competency in the **Python 3** programming language. Experience with deep neural networks (specifically variations of **CNNs**).			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Deep Graph Library, PyTorch	English	2 hours	$30 (excludes tax, if applicable)	N/A

Back

Course Name	Description	Prerequisites
Introduction to Physics-Informed Machine Learning With NVIDIA Modulus	High-fidelity simulations in science and engineering are computationally expensive and time-prohibitive for quick iterative use cases, from design analysis to optimization. NVIDIA Modulus, the physics machine learning platform, turbocharges such use cases by building physics-based deep learning models that are 100,000X faster than traditional methods and offer high-fidelity simulation results. Upon completion, you'll understand the various building blocks of Modulus and the basics of physics-informed deep learning. You'll also understand how the Modulus framework integrates with the overall Omniverse platform. > **Learn More**	> Familiarity with the Python programming language > An understanding of partial differential equations and their use in physics. > Familiarity with machine learning concepts like training and inference.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA Modulus	English	4 hours	$30 (excludes tax, if applicable)	N/A

Course Name	Description	Prerequisites
Modeling Time-Series Data With Recurrent Neural Networks in Keras	Explore how to classify and forecast time-series data using recurrent neural networks (RNNs), such as modeling a patient's health over time. > **Learn More**	Basic experience with deep learning.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
Keras	English	2 hours	$30 (excludes tax, if applicable)	N/A

Course Name	Description	Prerequisites
Optimized Vehicle Routing	NVIDIA cuOpt™ is a GPU-accelerated logistics solver that uses heuristics and optimizations to calculate complex vehicle-routing problems with a wide range of constraints. In this self-paced course, you'll work through a demonstration of a common vehicle-routing optimization problem. Upon completion, participants will be able to preprocess input data for use by NVIDIA cuOpt and compose variants of the problem that reflect real-world business constraints. > **Learn More**	Anyone can run the code to see how it works, but to get the most out of this content, we recommend: > An understanding of fundamental programming concepts in Python 3, such as functions, loops, dictionaries, and arrays. > A familiarity of matrix-based Python libraries, such as NumPy and **pandas.** > A familiarity with **NVIDIA RAPIDS**, in particular **cuDF**, is nice to have but not required.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA cuOpt, cuDF, SciPy, NumPy, pandas, GeoPandas, VeRoViz	English	1 hour	Free	N/A

Graphics and Simulation

Course Name	Description	Prerequisites
Assemble a Simple Robot in NVIDIA Isaac Sim™	In this course, you'll step through the "Assemble a Simple Robot" tutorial to rig a two-wheel mobile robot in a live NVIDIA Isaac Sim GPU environment. > **Learn More**	A Windows or Linux computer with the ability to install Omniverse Launcher and Omniverse applications; internet bandwidth sufficient to support the Isaac Sim client/server stream (performance will vary).

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA Isaac Sim	English	30 minutes	Free	N/A

Back

Course Name	Description	Prerequisites	Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
Build Beautiful, Custom UI for 3D Tools on NVIDIA Omniverse™	Experience the NVIDIA Omniverse development platform for builders and creators of virtual worlds. Become a master in UI with a deep dive into NVIDIA Omniverse Kit's powerful omni.ui suite of tools and frameworks. In this self-paced course, you'll build your own custom UI for workflows in Omniverse with hands-on exercises. > **Learn More**	Basic familiarity with Python (helpful, not required). Suggested materials to satisfy prerequisites: The Python Tutorial.	Omniverse Code, Visual Studio Code, Python, and the Python Extension	English, Simplified Chinese	90 minutes	Free	N/A
Develop, Customize, and Publish in NVIDIA Omniverse With Extensions	Want to change the functionality and user interface (UI) of NVIDIA Omniverse? Learn how to customize the Omniverse experience with extensions using Python code. > **Learn More**	A basic understanding of Python. A basic understanding of computer graphics is useful but not required.	Omniverse Code, Visual Studio Code, Python, and the Python Extension	English	8 hours	Free	Yes
Developing Omniverse Kit Applications	NVIDIA has built a number of large reference applications like Create, Drive Sim, and Isaac Sim to show some of the capabilities of Omniverse Kit. Everything that you see in the reference apps is possible for your custom app and you can leverage many of the existing NVIDIA extensions to kickstart your own apps. Apps are made up of many extensions working together to address specific workflows. While end-users and content creators leverage the Omniverse platform to connect and accelerate their 3D workflows, developers can plug into the platform layer of the Omniverse stack to easily build extensions, apps and microservices on Omniverse Kit. > **Learn More**	A basic understanding of Python. A basic understanding of computer graphics is useful but not required. Familiarity with: > Creating an extension for Omniverse. > Using Github. > How to use terminal commands.	Visual Studio Code and Python	English	90 minutes	Free	N/A
Easily Develop Advanced 3D Layout Tools on NVIDIA Omniverse	Get hands-on experience with NVIDIA Omniverse—the platform for connecting and creating physically accurate, 3D virtual worlds. See how easy it is to create your own custom scene layout tools in Omniverse Code with a few lines of Python script. In this self-paced course, you'll build your own custom scene layout in Omniverse with hands-on exercises in Omniverse Code and Python. > **Learn More**	A basic understanding of computer graphics concepts–such as vertices, meshes, and RGB values–and an understanding of fundamental programming concepts in Python like functions, loops, dictionaries, and arrays.	**Universal Scene Description**	English, Simplified Chinese	2 hours	Free	N/A

Back

Course Name	Description	Prerequisites			
Essentials of Developing Omniverse Kit Applications	In this course, participants will learn about kit files and how to create one, how to add extensions to applications, how to define the layout of an application and how to package and distribute an application.' **> Learn More**	> A basic understanding of Python > A basic understanding of computer graphics is useful but not required. > Creating an extension for Omniverse. > Using Github. > How to use terminal commands.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Visual Studio Code and Python	English	90 minutes	Free	N/A
Getting Started With USD for Collaborative 3D Workflows	Learn how to generate a scene using human-readable Universal Scene Description ASCII (.USDA) files. Upon completion, you'll be able to create your own scenes within the USD framework and will have a strong foundation to use it in applications, such as **NVIDIA Omniverse**, Maya, Unity, and Unreal Engine. **> Learn More**	A basic understanding of computer graphics concepts–such as vertices, meshes, and RGB values–and an understanding of fundamental programming concepts in Python like functions, loops, dictionaries, and arrays.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Universal Scene Description	English, Simplified Chinese	2 hours	Free	N/A
How to Build Customer 3D Scene Manipulator Tools on NVIDIA Omniverse	See how you can build advanced tools on the modular, easily extensible Omniverse platform. You'll learn from the Omniverse developer ecosystem team how you can extend and enhance the 3D tools you know and love today. In this self-paced course, you'll build your own custom scene manipulator tools in Omniverse with hands-on exercises writing a few lines of Python code. **> Learn More**	Basic familiarity with Python (helpful, not required). Suggested material to satisfy prerequisites: **The Python Tutorial.**			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certificate**
	Omniverse Code, Visual Studio Code, Python, and the Python Extension	English, Simplified Chinese	90 minutes	Free	N/A

Back

Course Name	Description	Prerequisites
Introduction to Robotic Simulations in NVIDIA Isaac Sim	Robotic automation has enjoyed great success in recent years with increasing hardware capabilities driving innovation in simulation and machine learning. In this course, we introduce you to Isaac Sim, NVIDIA Omniverse's solution for simulation and robotics. You'll learn how to tap into the simulation loop of a 3D engine and initialize experiments with objects, robots, and physics logic. This can be done programmatically using Omniverse Kit and Pixar USD commands, but the course will use Isaac Sim Core to wrap these low-level operations in an object-oriented fashion. By the end of the course, you'll be able to simulate and control NVIDIA JetBot™ and Franka Emika robots and coordinate them together to perform a handoff. The skills covered in this course are direct prerequisites for working with Isaac Gym and create a good starting point for exploring Isaac Sim and other Omniverse applications. The course is great for those interested in 3D scene specification and robotic simulation, but it's also useful for researchers looking to expand their toolkits and seasoned developers interested in exploring design patterns for Omniverse Kit development. **> Learn More**	> Intermediate knowledge and general comfort with Python 3. This includes familiarity with functions, classes, and basic design patterns. > Comfort with NumPy arrays and basic matrix operations. > A Windows or Linux machine with NVIDIA Omniverse and the Omniverse Streaming Client app.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
Isaac Sim, Omniverse Kit, NumPy	English, Simplified Chinese	4 hours	$30 (excludes tax, if applicable)	N/A

Course Name	Description	Prerequisites
Synthetic Data Generation for Training Computer Vision Models	How much data is enough? This is a common question when fine-tuning or training computer vision models. In cases where data collection is a limiting factor, we can use synthetic data! NVIDIA Omniverse Replicator streamlines synthetic data generation (SDG) using 3D assets into a single application, with the ability to modify the appearance and format of the data. This lab highlights one of the ways deep learning tools and Omniverse can be used together to streamline deep learning workloads. **> Learn More**	> Intermediate understanding of Python (including classes, objects, and decorators): learn about this topic from the Python.org tutorials > Basic understanding of Machine Learning and Deep Learning concepts and pipelines:learn about this topic from the "Deep Learning Demystified" video

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA Omniverse Replicator, NVIDIA Triton Inference Server, PyTorch	English	3 hours	$30	N/A

Back

Course Name	Description	Prerequisites

Infrastructure

Introduction to AI in the Data Center

Explore AI, GPU computing, NVIDIA AI software architectures, and how to implement and scale AI workloads in the enterprise data center.

> **Learn More**

Basic knowledge of enterprise networking, storage, and data center operations

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
Artificial intelligence, machine learning, deep learning, GPU hardware and software	English	4 hours	$49 (excludes tax, if applicable)	Available

Introduction to NVIDIA DOCA™ for DPUs

The NVIDIA DOCA Software Framework lets developers rapidly create applications and services on top of NVIDIA BlueField data processing units (DPUs). Together, DOCA and the BlueField DPU deliver breakthrough networking, security, and storage performance with a comprehensive, open development platform.

In this self-paced course, you'll learn the basic concepts of DOCA as a platform for accelerated data center computing on BlueField DPUs. Upon completion, participants will be equipped with introductory knowledge that will enable you to begin using DOCA and DPUs to develop applications that accelerate your data centers services.

> **Learn More**

> Familiarity with software architecture and how it relates to and executes on hardware.

> Suggested materials to satisfy prerequisite:

- **Enterprise Data Center Networking**
- **Data Center: Overview**
- **Data Center: Virtualization**

> Some working knowledge of data center networking.

> Suggested materials to satisfy prerequisite:

- **Introducing How Computers Work**
- **Hardware Acceleration**
- **Software Execution and Computing**

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
NVIDIA DOCA SDK	English, Simplified Chinese	2 hours	Free	N/A

Back

Workshop Name	Description	Prerequisites
Getting Started With DOCA Flow	NVIDIA DOCA is the key to unlocking the potential of the NVIDIA BlueField DPU, enabling you to offload, accelerate, and isolate data center workloads. With DOCA, developers can program the data center infrastructure of tomorrow by creating software-defined, cloud-native, DPU-accelerated services with zero-trust protection to address the increasing performance and security demands of modern data centers. DOCA Flow is the most fundamental API for building generic execution pipes in hardware. The library provides an API for building a set of pipes, where each pipe consists of match criteria, monitoring, and a set of actions. Pipes can be chained so that after a pipe-defined action is executed, the packet may proceed to another pipe. In this course, you'll be introduced to DOCA Flow programming by building an "ARP Storm Control" application, which prevents network failures caused by broadcast storms. It does so through the creation of a DOCA Flow pipeline that can dampen malicious broadcast network activity without impacting well-behaved traffic. **> Learn More**	A working knowledge of networking basics.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certificate
DOCA Flow	English, Simplified Chinese	2 hours	Free	N/A

Back

Instructor-Led Workshops for Administrators

Workshop Name	Description	Prerequisites

AI and Data Science

NVIDIA AI Enterprise Administration: Public Training

This hands-on training course explores architecture, installation, configuration, operation, and management of NVIDIA AI Enterprise.

> Learn More

None.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
N/A	English	12 hours	$1,500	N/A

Ethernet Cumulus

Cumulus® Linux: Public Bootcamp

Learn how to install, deploy, configure, and troubleshoot Cumulus-based networks. This course offers a perfect blend of hands-on training and theoretical education.

> Learn More

None.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
Cumulus Linux switches	English	12 hours	$1,500	Available

Cumulus Linux: Private Workshop

In this hands-on private training, you'll learn about NVIDIA Cumulus OS architecture, installation, configuration, operation, and management of Cumulus Linux running on NVIDIA switches.

> Learn More

None.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
Cumulus Linux switches	English	20 hours	Contact us	Available

InfiniBand

InfiniBand Customized Course

In this course, you'll learn about InfiniBand architecture and how to manage, monitor, and troubleshoot your InfiniBand network.

> Learn More

Network administrators and IT professionals that need to install, configure, manage, monitor, and troubleshoot the configuration and performance of InfiniBand networks.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
InfiniBand networks	English	16 hours	Contact us	Available

Back

NVIDIA DGX

Workshop Name	Description	Prerequisites	Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
NVIDIA DGX H100/A100 Administration: Private Workshop	This course provides an overview of the NVIDIA DGX A100 system and NVIDIA DGX Station™ A100, tools for in-band and out-of-band management, NGC, the basics of running workloads, and specific management tools and command-line interface (CLI) commands. In addition, this course includes content on Multi-Instance GPU (MIG), managing storage, performance validation, and other system management tools and concepts. > Learn More	System and network administrators and IT professionals that need to configure and verify the configuration and performance of DGX A100 systems and DGX Station A100.	DGX A100 system and DGX Station A100	English	16 hours	Contact us	N/A
NVIDIA DGX H100/A100 Administration: Public Workshop	This course provides an overview of the DGX A100 system and DGX Station A100's tools for in-band and out-of-band management, the basics of running workloads, specific management tools, and CLI commands. > Learn More	System and network administrators and IT professionals that need to configure and verify the configuration and performance of DGX A100 systems and DGX Station A100.	DGX A100 system and DGX Station A100	English	16 hours	$1,500	N/A
NVIDIA DGX BasePOD Administration: Private Workshop	This course provides an overview of DGX POD components and related processes, including the NVIDIA DGX A100 system, InfiniBand and ethernet networks, tools for in-band and out-of-band management, NGC, the basics of running workloads, and specific management tools and CLI commands. It includes instructions for managing vendor-specific storage per the architecture of your specific POD solution. > Learn More	System and network administrators and IT professionals that need to configure and verify the configuration and performance of DGX A100 POD clusters.	DGX POD cluster	English	16 hours	Contact us	N/A
NVIDIA DGX SuperPOD™ Administration: Private Workshop	This course is designed to help IT professionals successfully administer all aspects of a DGX SuperPOD cluster, including compute, storage, and networking. > Learn More	System and network administrators and IT professionals that need to configure and verify the configuration and performance of DGX SuperPOD clusters.	DGX SuperPOD cluster	English	16 hours	Contact us	N/A

Back

Workshop Name	Description	Prerequisites
Virtualization		
NVIDIA AI Enterprise Administration: Public Bootcamp	This course covers the platform and solution overview, hardware and software architecture, deployment options, licensing, temporal and spatial GPU partitioning, scaling, comprehensive validation, management, maintenance, monitoring, and troubleshooting. **> Learn More**	System administrators and IT professionals that need to install, configure, manage, monitor, and troubleshoot the configuration and performance of their NVIDIA AI Enterprise solution.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
NVIDIA AI Enterprise	English	12 hours	$1,500	N/A

Back

Online, Self-Paced Courses for Administrators

Course Name	Description	Prerequisites			
AI and Data Science					
Introduction to AI in the Data Center	Explore an introduction to AI, GPU computing, NVIDIA AI software architecture, and how to implement and scale AI workloads in the data center. > Learn More	None			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	N/A	English	4 hours	$49	Available
NVIDIA AI Enterprise Administration	This course covers the platform and solution overview, hardware and software architecture, deployment options, licensing, temporal and spatial GPU partitioning, scaling, comprehensive validation, management, maintenance, monitoring, and troubleshooting. > Learn More	To gain the most value from this course, the target audience should have a working knowledge in the following domains: > Data Center Infrastructure: Servers, Storage, Networking, GPUs, Operating Systems. > Virtualization: VMware vSphere. > Containerization: Docker.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	NVIDIA AI Enterprise	English	8 hours	$99	N/A
Cluster Administration					
Bright Cluster Manager Administration	This course is based on NVIDIA Bright Cluster Manager and gives an overview of the cluster management tools, Bright View and cluster management shell (CMSH). > Learn More	None.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	NVIDIA Bright Cluster Manager	English	5 hours	Free	N/A
Bright Cluster Manager Autoscaling Hybrid Cloud	This course is based on NVIDIA Bright Cluster Manager and gives an overview of extending the cluster to the cloud with Cluster as a service and cluster extension (i.e., hybrid cloud). > Learn More	None			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	NVIDIA Bright Cluster Manager	English	3 hours	Free	N/A
Introduction to Bright Cluster Manager	This course is based on NVIDIA Bright Cluster Manager and gives an overview of the usage and components of the software. > Learn More	None			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	NVIDIA Bright Cluster Manager	English	3 hours	Free	N/A

Back

Course Name	Description	Prerequisites
DGX		
NVIDIA DGX Cloud	This course is based on NVIDIA DGX Cloud using NVIDIA Base Command Platform. You'll learn to manage users and teams, run single and multi-node jobs, and manage data. > Learn More	None

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
DGX Base Command Manage	English	1 hour	Free	N/A

Course Name	Description	Prerequisites
Ethernet		
Linux Networking Fundamentals	Learn the fundamental concepts and commands behind Linux-based open networking. > Learn More	None

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
Linux networking concepts	English	6 hours	$99	N/A

Course Name	Description	Prerequisites
Network Administration With the NVIDIA Onyx™ Switch System	This course provides the required set of skills to configure and manage NVIDIA Ethernet switch systems. You'll learn in depth layer 2 configurations such as VLAN, STP, LAG, and MLAG, as well as how to configure layer 3 features such as BGP. > Learn More	> Basic understanding of Ethernet network principles. > Basic understanding of switching and routing concepts.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
NVIDIA Onyx	English	3 hours	$99	N/A

Course Name	Description	Prerequisites
RDMA Over Converged Ethernet (RoCE) From A to Z	In this course, you'll learn what RoCE is, how it works, the different network types RoCE can run over, and how to configure RoCE for each network type. > Learn More	Basic understanding of networking concepts and the Open Systems Interconnection (OSI) model.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
RoCE	English	2 hours	Free	N/A

Back

Course Name	Description	Prerequisites			
Graphics and Simulation					
NVIDIA Omniverse Enterprise Administration	The course covers the solution overview, hardware and software architecture, deployment options, installation, configuration, licensing, scaling, comprehensive validation, security, management, maintenance, monitoring, and troubleshooting. The instruction and guidance are based on NVIDIA's best practices and cover the critical knowledge and skills for deploying, administering, and managing your Omniverse solution. > Learn More	None			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	Omniverse	English	6.5 hours	$99	N/A
InfiniBand					
InfiniBand Essentials	This self-paced course covers the fundamental first steps into the world of InfiniBand. If you're looking to become more familiar with InfiniBand's benefits, uses, architecture layers, and management concepts, this is the best place to start. > Learn More	General understanding of networking concepts and principles.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	InfiniBand	English	1.5 hours	Free	N/A
InfiniBand Professional	This course covers the fundamentals of the InfiniBand technology from a usability point of view and builds on the details of the InfiniBand architecture specification. You'll learn how to install, configure, manage, troubleshoot, and monitor your InfiniBand network. > Learn More	General understanding of networking concepts and principles.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	InfiniBand	English	6 hours	$250	Available
Management					
Data Center Management Made Easy With NVIDIA UFM	Learn about NVIDIA Unified Fabric Manager (UFM) and its capabilities, advantages, and components through a set of interactive learning units, videos, and simulators. > Learn More	Understanding of InfiniBand fabrics and management concepts			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	N/A	English	3 hours	$49	N/A

Back

Workshop Name	Description	Prerequisites			
NVIDIA License System	NVIDIA License System (NLS) is a new licensing solution to support the continued expansion of the NVIDIA enterprise software portfolio. This course will help you to learn about NLS and how you can move from your existing licensing solution to NLS. **> Learn More**	> Basic understanding of virtual appliances installation and setup. > Familiarity with web/cloud-based applications. > Familiarity with NVIDIA products like virtual GPU (vGPU) and NVIDIA AI Enterprise.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	Cloud License Service (CLS) and Delegated License Service (DLS)	English	2 hours	Free	N/A

Network

Ansible Essentials for Network Engineers	In this course, you'll explore a variety of Ansible modules and write playbooks specifically adapted to modern data centers. This course includes an exclusive hands-on lab environment and exercises to practice real-world scenarios in real cloud environments. **> Learn More**	> Basic Linux administration. > General understanding of networking concepts and principles.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	Ansible	English	3 hours	$49	N/A
Introduction to Networking	In this course, we'll cover the basics of Ethernet technology and understand how data is forwarded in an Ethernet network. **> Learn More**	None			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	N/A	English	1 hour	Free	N/A
MLXlink and MLXcables Debug Tools	In this course, you'll learn about the MLXlink and MLXcables debug tools. These debug tools are used for both basic link troubleshooting and for analyzing the more complex link characteristics. **> Learn More**	Good technical background and understanding of networking hardware.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	MLXLink and MLXcables	English	2 hours	Free	N/A
NVIDIA BlueField DPU Administration	Learn the basic concepts of BlueField DPUs as a platform for accelerated data center computing. **> Learn More**	> Basic knowledge and experience in networking concepts and principle. > Basic knowledge and experience in Linux administration.			
	Tools, Libraries, Frameworks	**Languages**	**Duration**	**Price**	**Certification Exam**
	N/A	English	3 hours	$49	N/A

Back

Workshop Name	Description	Prerequisites
RDMA		
The Fundamentals of RDMA Programming	This course allows C programmers to dive into the RDMA programming world without requiring previous experience in networking or RDMA programming. We've also added tips and tricks, as well as do's and don'ts, so the skills you acquire will truly serve you when you need them. > **Learn More**	Understanding of C/C++ programming.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
RDMA, C/C++	English	4 hours	$49	N/A

Back

Certifications

Certification Name	Description	Prerequisites
NVIDIA Certified Associate: AI in the Data Center	This is an entry-level certification that validates foundational concepts of adopting artificial intelligence computing by NVIDIA in a data center environment. The exam is online and remote proctored with 50 questions and a time limit of 60 minutes for completion. **> Learn More**	A basic understanding of data center infrastructure.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
N/A	English	1 hour	$135	Available

Certification Name	Description	Prerequisites
NVIDIA Certified Professional: InfiniBand	This is an intermediate level certification that validates core concepts for designing, deploying, and managing NVIDIA InfiniBand fabrics. The exam is online and remote proctored with 40 questions and a time limit of 90 minutes for completion. **> Learn More**	A thorough understanding of data center infrastructure and networking.

Tools, Libraries, Frameworks	Languages	Duration	Price	Certification Exam
NVIDIA InfiniBand fabrics	English	1.5 hours	$220	Available

Ready to Get Started?

To get started with hands-on training, visit **www.nvidia.com/en-us/learn/enterprise**

For questions, contact us at **nvdli@nvidia.com**

NVIDIA