

Get ready for AI Infrastructure Certification

Thank you for joining us today!

The webinar will begin shortly.

To improve your viewing experience and audio quality, we recommend ensuring your internet connection is stable and that you close other open applications on your device.

Do you hear music?


If you cannot hear the music, please refresh your browser and ensure your computer speakers (or headset) are turned on and the volume is up.

Still having issues?

Please contact us in the Q&A window.

Before we begin...


- **Questions?** Please submit your questions in the **Q&A window** at any time
- A copy of this presentation is available in the **Resource list**
- **Having trouble with seeing/hearing the presentation?**
 - To maximize the quality of this audio stream, please close any open applications aside from your browser window.
 - If your audio sputters or the slides lag, refresh your browser or try a different browser
 - If you encounter any other technical issues today, let us know in the Q&A window and we will help you troubleshoot.



Agenda

- Introductions
- AI Infrastructure skillset
- Cracking the exams: prep tactics
- Moving toward hands-on based testing
- Next steps
- Live Q&A

Today's Speakers



Marco Alves

Certification Program Lead
NVIDIA



Brad Christian

Chief Enterprise Architect
Computacenter




Gary Fluitt

Technical Certification Manager
NVIDIA

Program Overview

Why Get NVIDIA Certified?



NVIDIA Certification Framework




Based on NVIDIA HW & SW Solutions,
**Designed for Infrastructure and Operations Roles at
Customers, Partners and General Public**

Based on Application Requirements for Accelerated Computing,
AI, and Related Workloads
**Designed for Developer Roles at Customers, Partners
and General Public**


AI Infrastructure Skillset

Gary Fluitt, NVIDIA

NVIDIA AI Infrastructure: AI Factories




AI Factory – skills needed



Based on NVIDIA HW & SW Solutions,
Designed for Infrastructure and Operations Roles at
Customers, Partners and General Public

NVIDIA Provides a Full Stack for AI Factories

Built on NVIDIA customer-validated data center reference architectures



Select Audience and Pick Certification

[Introduction](#) [Certifications](#) [Frequently Asked Questions](#) [Contact Us](#) [Stay Informed](#)

Certifications

[Certifications for Infrastructure Professionals](#) [Certifications for Developers](#)

NVIDIA-Certified Associate: AI Infrastructure and Operations (NCA-AIIO)

An associate-level assessment for IT professionals and others looking to validate fundamental skills in AI infrastructure and operations.

[Learn About This Certification >](#)

NVIDIA-Certified Professional: AI Infrastructure (NCP-AII)

A professional-level assessment to validate a candidate's ability to deploy, manage, and maintain AI infrastructure by NVIDIA.

[Learn About This Certification >](#)

NVIDIA-Certified Professional: AI Operations (NCP-AIO)

A professional-level assessment to validate a candidate's ability to monitor, troubleshoot, and optimize AI infrastructure by NVIDIA.

[Learn About This Certification >](#)

NVIDIA-Certified Professional: AI Networking (NCP-AIN)

A professional-level assessment to validate a candidate's ability to deploy and configure environments that leverage NVIDIA's advanced networking technologies.

[Learn About This Certification >](#)

NVIDIA-Certified Associate

AI Infrastructure and Operations

About This Certification:


An entry-level credential that validates the foundational concepts of AI computing in infrastructure and operations.

Audience: Data center technicians, Developer operations (DevOps) engineers, IT managers, Networking engineers, Systems administration staff, Solution architects

Prerequisites: A basic understanding of data center infrastructure

Recommended Training (Optional)

- [AI Infrastructure & Operations Fundamentals Self-Paced Course](#)
- [Free Study Guide](#)



Topics Covered

NVIDIA-Certified Professional AI Networking

About This Certification


The NCP-AI Networking certification is an intermediate-level credential that validates a candidate's ability to deploy and configure environments that leverage NVIDIA's advanced networking technologies.

Audience: Data center administrators, Infrastructure administrators, Network administrators, Network engineers, Storage administrators, System administrators, Solution architects

Prerequisites: Two to three years of operational experience working in a data center with NVIDIA hardware solutions. The candidate should be able to deploy and manage NVIDIA AI networking infrastructure in support of AI workloads.

Recommended Training (Optional)

- [Data Center Management Made Easy with NVIDIA UFM](#)
 - [BlueField DPU Administration](#)
 - [Fundamentals of RDMA Programming](#)
 - [Cumulus Linux Bootcamp](#)
 - [InfiniBand Network Administration](#)
 - [Spectrum-X Networking Platform Administration](#)
-
- [Free Study Guide](#)



Topics Covered

NVIDIA-Certified Professional AI Infrastructure

About This Certification


The intermediate-level credential that validates a candidate's ability to deploy, manage, and maintain AI infrastructure by NVIDIA.

Audience: Data center administrators, Infrastructure administrators, MLOps Engineers, Network Administrators, Network Engineers, Storage Administrators, System Administrators, Solution Architects

Prerequisites: Two to three years of operational experience working in a data center with NVIDIA hardware solutions. The candidate should be able to deploy all the parts of a data center infrastructure in support of AI workloads.

Recommended Training (Optional)

- [AI Infrastructure & Operations Fundamentals Self-Paced Course](#)
- [AI Infrastructure Professional Instructor-led Workshop](#)
- [Free Study Guide](#)



Topics Covered

NVIDIA-Certified Professional AI Operations

About This Certification


Intermediate-level credential that validates a candidate's ability to monitor, troubleshoot, and optimize AI infrastructure by NVIDIA.

Audience: MLOps engineers, DevOps engineers, Solution architects, System architects, AI Infrastructure engineers

Prerequisites: Two to three years of operational experience working in a data center with NVIDIA hardware solutions. The candidate should be able to monitor and manage all the parts of a data center infrastructure in support of AI workloads.

Recommended Training (Optional)

- [AI Infrastructure & Operations Fundamentals Self-Paced Course](#)
- [AI Operations Professional Instructor-led Workshop](#)
- [Free Study Guide](#)



Topics Covered

NVIDIA-Certified Professional: AI Operations Exam Study Guide

Study guides

Administration: Exam Weight 23%

This section assesses an individual's proficiency in administering various AI workload environments. It covers tasks like managing Slurm and Kubernetes clusters, which are essential for scheduling and orchestrating tasks. Additionally, it tests knowledge of administering Run:ai, an AI workload management platform, and configuring MIG (Multi-Instance GPU). Finally, it evaluates the ability to describe data center architecture specifically for AI workloads.

- 1.1 Administer Slurm cluster.
- 1.2 Describe data center architecture for AI workloads
- 1.3 Administer Run:ai
- 1.4 Administer Kubernetes
- 1.5 Configure MIG for AI and HPC.

Sub-topics and topic weighting

Recommended Training (Optional)

Course reference: [AI Operations](#)

- > Module 6, AI Data Center Management, Unit 2: Data Center Architecture for AI Workloads
- > Module 6, AI Data Center Management, Unit 3: Scheduling AI Workloads With Slurm
- > Module 6, AI Data Center Management, Unit 5: Managing AI at the Edge With Fleet Command
- > Module 7, Virtualizing GPU Resources

Recommended training (optional)

Suggested Readings

- > [Advanced Cluster Configuration](#)
- > [NVIDIA DGX Cloud Create Documentation](#)
- > [NVIDIA Multi-Instance GPU User Guide](#)
- > [Monitoring GPUs in Kubernetes with DCGM | NVIDIA Technical Blog](#)
- > [NVIDIA Spectrum-X™ Ethernet networking platform](#)
- > [NVIDIA Cloud Native Technologies](#)
- > [Kubernetes Documentation: Administer a Cluster](#)
- > [Debugging and Troubleshooting — NVIDIA DCGM Documentation latest documentation](#)
- > [DCGM Diagnostics](#)
- > [NVIDIA System Management Interface](#)


Suggested readings (free)

Cracking the Exams: Prep Tactics

Brad Christian, Computacenter

Prep Tactics

- Commit publicly to big goals.
- Find fellow students and pool resources.
- Cut the marketing noise - the exam blueprint is key.
- Lean on the ecosystem.
- Plain text is king & GitHub.
- Don't overuse AI, thoughtful note taking has its place.
- Fail fast! A failed attempt is a valuable scouting mission.
- Finally, AI Infrastructure is a lot of Linux and a lot of networking, be prepared to leave your comfort zone.



Moving Toward Hands-On Based Testing

Gary Fluitt, NVIDIA

AI Operations Exam Lab

head-node-ex1worker-node001-ex1head-node-ex2worker-node001-ex2head-node-ex3 (bcm11)worker-node001-ex3

Ubuntu 22.04.4 LTS bcm tty1
bcm login: [OK] Started BCM daemon.
[OK] Started Snap Daemon.
Starting Time & Date Service...
[OK] Started Time & Date Service.
[OK] Finished Wait until snapd is fully seeded.
[OK] Finished Service for snap application lxd.activate.
[OK] Reached target Multi-User System.
[OK] Reached target Graphical Interface.
Starting Record Runlevel Change in UTM...
[OK] Finished Record Runlevel Change in UTM.

 NVIDIA

Exit Lab

InstructionsResources

NVIDIA-Certified Professional: AI Operations (NCP-AIO)

Introduction

Welcome to the NVIDIA-Certified Professional: AI Operations assessment. In this lab you'll demonstrate your knowledge of Base Command Manager through various administrative tasks.

Objectives

- Troubleshoot and manage a BCM-managed Slurm cluster by restoring a non-operational compute node, verifying job execution, and controlling node states.
- Configure and verify a BCM-managed Kubernetes cluster by setting up the control plane, ensuring worker node visibility, and deploying a GPU-enabled workload.
- Provision a new compute node in a BCM-managed cluster by creating and assigning a software image and node category, adding the node to the cluster, and verifying its operational status.

← PreviousNext →

Troubleshoot a cluster in Slurm

Use Kubernetes to run a workload

Adding nodes to clusters using BCM

AI Operations Exam Lab

head-node-ex1worker-node001-ex1head-node-ex2worker-node001-ex2head-node-ex3

Welcome to
Base Command Manager trunk

Use Alt+F2 for console access.

Node Info #####

Head Node: bcm
Hostname: slurm-worker
Kernel: Linux 5.15.0-113-generic x86_64
Memory: 23 GiB

CPU:
16 logical processors - 1 physical CPU - 16 cores per CPU - 1 thread per CPU
2095 MHz - GenuineIntel - Intel(R) Xeon(R) Gold 6230 CPU @ 2.10GHz

Storage:
vendor model rev size
1 x - - 4K
1 x VMware Virtual disk 2.0 250G

 NVIDIA

Exit Lab

InstructionsResources

NVIDIA-Certified Professional: AI Operations (NCP-AIO)

Exercise 1: Slurm administration using Base Command Manager

You're a cluster administrator and one of your users submitted a

Environment

- For this exercise, you'll use the following systems:

System	Username	Password
head-node-ex1	root	bcm123
worker-node001-ex1	root	bcm123

Task list

☐ 1. Diagnose the issue on the worker node and return it to the idle state

Previous

Next

Exercise Instructions

Task List

AI Operations Exam Lab

head-node-ex1worker-node001-ex1head-node-ex2worker-node001-ex2head-node-ex3 (bcm11)worker-node

Welcome to
Base Command Manager trunk

Use Alt+F2 for console access.

Node Info #####

Head Node: bcm
Hostname: slurm-worker
Kernel: Linux 5.15.0-113-generic x86_64
Memory: 23 GiB

CPU:
16 logical processors - 1 physical CPU - 16 cores per CPU - 1 thread per CPU
2095 MHz - GenuineIntel - Intel(R) Xeon(R) Gold 6230 CPU @ 2.10GHz

Storage:

	vendor	model	rev	size
1	x -	-	-	4K
1	x VMware	Virtual disk	2.0	250G

 NVIDIA

InstructionsResources

?

⚙

head-node-ex1

worker-node001-ex1

Username

Password

DVD Drive

Network Adapter

Ctrl+Alt+Delete

Open in New Window

head-node-ex2

worker-node001-ex2

head-node-ex3 (bcm11)

Screen Unavailableworker-node001-ex3

Exercise Resources

21
 NVIDIA

AI Operations Exam Lab

Score Report

Show As End Screen: On

Email Score Report Link: On

Congratulaitons John Doe you've passed the NCP - AI Operations exam!

Section scores are meant only as a suggestion for further study and do not reflect absolute scores in each section.

Section scores are an average of your overall exam score and your section score. Please use these only as study guidance.

Topic Name	Score
Administration	58 %
Installation and Deployment	71 %
Troubleshooting and Optimization	66 %
Workload Management	64 %
Slurm Troubleshooting	82%
Kubernetes Verification	72%
Provisioning using BCM	69%

Your badge will be issued on Credly (look for an email from Credly to accept your digital badge, it may be in your junk mail). Congratulations on achieving this NVIDIA credential!

NVIDIA credentials are valid for 2 years. Before your credential expires, you will receive a notice from Credly. If there is a renewal path you will receive instructions at that time.

Once again, congratulations on your achievement!

Score Report

Hands-on Lab Exam FAQ

- How much time for the lab portion of the exam?
- Can I fail the lab and still pass the exam?
- Can I take the lab exam at home?
- Can I take the lab exam at a different time from the multiple-choice?
- What happens if I lose connectivity to the exam / what about bandwidth issues?
- Do I need a GPU-enabled computer to take the lab exam?
- If I fail the exam, can I retake only the part that I failed (lab or multiple-choice)?
- How can I practice the lab before the exam?

Next Steps

Marco Alves, NVIDIA

Next Steps

Visit nvidia.com/certification to:

- Learn how to prepare for your NVIDIA exam
- Redeem your special webinar promo code
- Join our Certification Champions program



How to Register

NVIDIA-Certified Associate

AI Infrastructure and Operations

(NCA-AIIO)

[Register for Exam](#)

[About This Certification](#) [Topics Covered](#) [Preparation Materials](#) [Stay Informed](#) [Register for Exam](#)

About This Certification

The NCA-AIIO certification is an entry-level credential that validates the foundational concepts of AI computing related to infrastructure and operations. The exam is online and proctored remotely, includes 50 questions, and has a 60-minute time limit.

Please carefully review [NVIDIA's examination policy](#) before scheduling your exam.

If you have any questions, please contact us [here](#).

Certification Exam Details

Duration: One hour

Price: \$125

Certification level: Associate

Subject: AI Infrastructure and operations

Number of questions: 50

Prerequisites: A basic understanding of data center infrastructure.

Language: English

50% Webinar Promo Code

EXAM

NVIDIA-Certified Associate: Generative AI LLMs (NCA-GENL)

DURATION

60 Minutes

LANGUAGE

English

AMOUNT DUE

\$125.00

RESERVATION DATE AND TIME

Take the test on: Jun 04, 2025 at 03:30 PM

TERMS AND CONDITIONS

Please carefully review NVIDIA's [Terms and Conditions](#) and [FAQs](#) before scheduling your exam.

To ensure a secure and fair testing environment, please follow the **room preparation** and **behavior requirements**:

- Take your exam in a private room where no other person is allowed entry during the exam session.
- Your desk and work area should be free from all materials, such as papers, books, or electronics.
- No phones, smartwatches, or similar devices are permitted in the room.
- Headphones, earbuds, and similar devices are prohibited.
- Only one mouse, keyboard, and monitor/display are allowed during the exam.
- Keep your head and shoulders fully visible throughout the exam to ensure clear proctor visibility.
- Sharing, reproducing, or transmitting exam content is prohibited and may result in disqualification and potential legal action.

Please read the full [Certiverse Testing Rules](#) and [Compliance Requirements](#) before completing your exam scheduling.

PAYMENT

PRICE

\$125.00

PAYMENT METHOD ⓘ

☒ CREDIT CARD

☐ VOUCHER CODE

Credit Card details will be entered in the payment processing system, after confirming the details of your booking.

PROMO CODE

Promo Code

Optional

APPLY


☐ Employer Information for Reimbursement ⓘ

BACK

CONTINUE

Enter Promo Code





Thank You for Attending!

Please take a few minutes to complete the feedback survey.

Sample AI Infrastructure Exam Questions

Sample Exam Question 1

NVIDIA-Certified Associate: AI Infrastructure and Operations

An operations team is building a multi-tenant AI cluster and will use Slurm together with NVIDIA Base Command Manager (BCM).

Which requirement is best addressed by adding BCM instead of relying only on Slurm's basic OS scheduling?

- A. Enforcing GPU- and project-level quotas for fair multi-tenant resource sharing.
- B. Performing rolling Linux kernel upgrades across all compute nodes.
- C. Selecting the lowest-cost cloud region based on power and energy prices.
- D. Managing BIOS configuration settings during initial server provisioning.

Sample Exam Question 2

NVIDIA-Certified Professional: AI Infrastructure

An engineer plans to install more GPUs into an existing AI server.

Which checks are most important to perform on the server before adding the new GPUs?

- A. Network switch speed, internet bandwidth, and backup schedule.
- B. Hypervisor version, container runtime version, and GPU driver version.
- C. PCIe slots, PSU capacity/connectors, and server cooling.
- D. RAID level, storage capacity, and backup retention period.

Sample Exam Question 3

NVIDIA-Certified Professional: AI Networking

An AI team is designing a Spectrum-X based fabric for GPU training and wants to keep job performance stable even if one of the uplinks between leaf and spine switches goes down.

Which Spectrum-X capability primarily helps maintain high effective bandwidth when a link in the fabric fails?

A. RoCE Performance Isolation

B. RoCE Adaptive Routing

C. RoCE Congestion Control

D. NVIDIA NetQ

Sample Exam Question 4

NVIDIA-Certified Professional: AI Operations

A user launches a GPU-enabled NGC container, but inside the container the command `nvidia-smi` reports “no CUDA-capable device is detected.” Which quick check is most likely to identify the problem?

- A. Rebuild the container image so its CUDA toolkit matches the host kernel exactly.
- C. Increase the container memory limit so CUDA libraries can allocate additional RAM.
- C. Run the container in privileged mode so it can modify required kernel modules.
- D. Confirm the host NVIDIA driver, CUDA version, and GPU runtime are correctly configured.**