

OUR LEARNING EXCLUSIVE

- Custom exam prep software and materials
- Course-specific thinQtank[®] Learning publications to make the learning fun and engaging
- Hands-on labs administering GPU clusters, workload managers and Kubernetes
- **We do not teach the test.** We always deliver valuable hands-on experience
- All reading material and study guides sent to you when you register
- Every course taught by AI engineers experienced with NVIDIA platforms
- Six months mentoring access to our engineers after the course

TRAINING CAMP FORMAT

thinQtank[®] Learning delivers this course as a **four-day instructor-led training camp**, Monday to Thursday. About 50% lecture and 50% hands-on labs. Contact us for the schedule covering your dates.

PREREQUISITES

- No certification is required
- Recommended: Two to three years of data center operations with NVIDIA hardware
- NCA-AIO or equivalent data center experience is a strong foundation

TARGET AUDIENCE

- MLOps Engineer
- DevOps Engineer
- AI Infrastructure Engineer
- Solution / System Architect

CERTIFICATION

This course prepares you for the NVIDIA-Certified Professional: AI Operations (NCP-AIO) exam.

Exam	NCP-AIO, 30 multiple choice questions plus 3 hands-on lab exercises in 120 minutes
Delivery	Online, remote proctored (Certiverse)
Exam fee	\$500, set by NVIDIA
Renewal	Valid two years; retake to renew

COURSE OVERVIEW

thinQtank[®] Learning is offering a unique four-day training camp comprised of four days of instructor-led learning that gives students the knowledge and skills needed to install and deploy the NVIDIA AI software stack, administer GPU clusters, manage AI workloads with schedulers and Kubernetes, and troubleshoot and optimize AI operations.

The camp follows the official NVIDIA NCP-AIO exam topics and study guide, with hands-on labs administering GPU clusters, workload managers and Kubernetes throughout.

This course prepares you for **the NVIDIA-Certified Professional: AI Operations (NCP-AIO) exam**.

WHAT YOU WILL BE ABLE TO DO

- Use Mission Control and BCM Base View to monitor cluster performance, utilization and node health
- Install Kubernetes, Slurm and Run:ai and deploy DOCA services on DPUs
- Administer Slurm, Kubernetes and Run:ai clusters, users, roles and node categories
- Configure MIG and apply patches, firmware and software images across nodes
- Deploy training and inference workloads and allocate GPU resources between teams
- Deploy containers from NGC
- Troubleshoot Docker, Fabric Manager (NVLink/NVSwitch), BCM, Magnum IO and storage performance

EXCLUSIVE LEARNING PACKAGE

- NCP-AIO reading plan and review guide
- NCP-AIO practice exams
- Lab notebooks you keep after the course
- Course video recordings for review

DELIVERY METHODS

- Instructor-led training
- Immersive live-online training
- On-site and custom delivery

NVIDIA CERTIFICATION PATHS

AI Infrastructure

NCA-AIO · NCP-AII · NCP-AIO ·
NCP-AIN · NCP-ARI

Data center, operations and
networking

Data Science

NCA-ADS · NCP-ADS

GPU-accelerated analytics and ML

Generative AI

NCA-GENL · NCA-GENM · NCP-
GENL · NCP-AAI

LLMs, multimodal and agents

3D and Simulation

NCP-OUSD

OpenUSD development

COURSE MODULES

01 AI Data Center Architecture for Operations

03 Installing Slurm, Kubernetes and Run:ai

05 Patching, Firmware and Software Image Management

07 Slurm Administration

09 MIG Configuration

11 NGC Containers and Docker

13 Hands-On Lab Practice

02 NVIDIA Mission Control and BCM Base View Monitoring

04 Deploying DOCA Services on BlueField DPUs

06 Users, Roles, Node Categories and Reporting

08 Kubernetes and Run:ai Administration

10 Training and Inference Workloads and Resource Allocation

12 Troubleshooting: Fabric Manager, Magnum IO and Storage

14 Exam Review

EXAM TOPIC AREAS

01 Installation and Deployment

03 Workload Management

02 Administration

04 Troubleshooting and Optimization



TERMS

© 2026 thinqtank® Global, Inc. All rights reserved. The product or information materials are protected by U.S. and intellectual property laws. thinqtank Global, thinqtank Learning and the thinqtank logo are registered trademarks of thinqtank Global, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. NVIDIA, CUDA, RAPIDS, InfiniBand, Spectrum, NeMo, NIM and Omniverse are trademarks of NVIDIA Corporation. OpenUSD is maintained by the Alliance for OpenUSD. thinqtank Global, Inc. warrants that it will perform these training services in a reasonable manner using generally accepted industry standards and practices. THE EXPRESS WARRANTY SET FORTH IS IN LIEU OF ALL OTHER WARRANTIES, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A

PARTICULAR PURPOSE WITH RESPECT TO THE SERVICES AND DELIVERABLES PROVIDED BY THINQTANK GLOBAL, INC., OR AS TO THE RESULTS WHICH MAY BE OBTAINED THEREFROM. THINQTANK GLOBAL, INC. WILL NOT BE LIABLE FOR ANY THIRD-PARTY SERVICES OR PRODUCTS IDENTIFIED OR REFERRED TO CUSTOMER. All materials provided in this training are copyrighted by thinqtank Global, Inc. ("Learning Materials"). thinqtank Global, Inc. grants the customer of this learning a license to use Learning Materials strictly for the purpose of facilitating such company's internal understanding, utilization and operation of the technology covered herein. Except as set forth expressly in the sentence above, there is no transfer of any intellectual property rights or any other license granted under the terms of this training.