

OUR LEARNING EXCLUSIVE

- Custom exam prep software and materials
- Course-specific thinQtank[®] Learning publications to make the learning fun and engaging
- Hands-on labs on InfiniBand and Spectrum Ethernet fabrics
- **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 and networking
- NCA-AIIO or equivalent data center experience is a strong foundation

TARGET AUDIENCE

- Network Engineer
- Network Administrator
- Data Center Administrator
- Solution Architect

CERTIFICATION

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

Exam	NCP-AIN, 70–75 questions in 120 minutes
Delivery	Online, remote proctored (Certiverse)
Exam fee	\$400, 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 design and optimize AI data center networks, deploy and manage NVIDIA Spectrum Ethernet and InfiniBand fabrics, integrate with Kubernetes, troubleshoot with NVIDIA tools and automate network configuration.

The camp follows the official NVIDIA NCP-AIN exam topics and study guide, with hands-on labs on InfiniBand and Spectrum Ethernet fabrics throughout.

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

WHAT YOU WILL BE ABLE TO DO

- Describe AI factory networking, rail-optimized topologies and GPU-to-GPU communication
- Configure Spectrum-X switches for RoCE with QoS, ECN, PFC, adaptive routing and telemetry
- Configure BGP-EVPN multi-tenancy and simulate networks with NVIDIA Air
- Provision InfiniBand fabrics with HA, PKeys, QoS and adaptive routing, monitored with UFM
- Deploy and verify the NVIDIA Network Operator for RDMA/InfiniBand in Kubernetes
- Troubleshoot with NetQ, WJH, UFM and ib* diagnostic tools
- Automate Spectrum-X configuration with NVUE templates and Ansible

EXCLUSIVE LEARNING PACKAGE

- NCP-AIN reading plan and review guide
- NCP-AIN 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 Networking Architecture	02 Rail-Optimized Topologies and GPU-to-GPU Communication
03 Cumulus Linux and NVUE Fundamentals	04 Spectrum-X and RoCE: QoS, ECN, PFC and Adaptive Routing
05 BGP-EVPN Multi-Tenancy and NVIDIA Air	06 SuperNICs and DOCA
07 Monitoring with NetQ, Telemetry and WJH	08 InfiniBand Fundamentals and Initial Configuration
09 InfiniBand PKeys, QoS, Adaptive Routing and UFM	10 Kubernetes Network Operator
11 InfiniBand and RDMA Troubleshooting Tools	12 Automation with NVUE Templates and Ansible
13 Exam Review	

EXAM DOMAINS		
AI Data Center Design and Optimization		5%
NVIDIA Spectrum Networking		30%
NVIDIA InfiniBand Networking		30%
Kubernetes Integration		5%
Troubleshooting Tools		20%
Automation and Configuration		10%



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.