



DCAI - IMPLEMENTING CISCO DATA CENTER AI INFRASTRUCTURE V1.0

Course Code: 860064

The Implementing Cisco Data Center AI Infrastructure (DCAI) course is designed to equip professionals with the skills to support, secure, and optimize AI workloads within modern data center environments.

This comprehensive program delves into the unique characteristics of AI/ML applications, their influence on infrastructure design, and best practices for automated provisioning. Participants will gain in-depth knowledge of security considerations for AI deployments and master day-2 operations, including monitoring and advanced troubleshooting techniques such as log correlation and telemetry analysis. Through hands-on experience, including practical application with tools like Splunk, learners will be prepared to efficiently monitor, diagnose, and resolve issues in AI/ML-enabled data centers, ensuring optimal uptime and performance for critical organizational workloads.

This training prepares you for the 300-640 DCAI v1.0 exam. If passed, you earn the Cisco Certified Specialist - Data Center AI Infrastructure certification and satisfy the concentration exam requirement for the Cisco Certified Network Professional (CCNP) Data Center certification.

This training also earns you 38 Continuing Education (CE) credits toward recertification.

This training combines content from Operate and Troubleshoot AI Solutions on Cisco Infrastructure (DCAIAOT) and AI Solutions on Cisco Infrastructure Essentials (DCAIE) training.

What You'll Learn

This training will help you:

- Acquire comprehensive skills to support, secure, and optimize AI workloads within modern data center environments
- Understand the design, implementation, and advanced troubleshooting of AI infrastructure, including network challenges and specialized hardware
- Gain in-depth knowledge of AI/ML concepts, generative AI, and their practical application in network management and automation
- Apply hands-on techniques for monitoring, diagnosing, and resolving issues, leveraging tools like Splunk and utilizing AI for enhanced productivity in

- network operations
- Prepare for the 300-640 DCAI v1.0 exam
- Earn 38 CE credits toward recertification

Who Needs to Attend

- Network Designers
- Network Administrators
- Storage Administrators
- Network Engineers
- Systems Engineers
- Data Center Engineers
- Consulting Systems Engineers
- Technical Solutions Architects
- Cisco Integrators/Partners
- Field Engineers
- Server Administrators
- Network Managers
- Program Managers
- Project Managers

Prerequisites

There are no prerequisites for this training. However, the knowledge and skills you are recommended to have before attending this training are:

- Cisco UCS compute architecture and operations
- Cisco Nexus switch portfolio and features
- Data Center core technologies

These skills can be found in the following Cisco Learning Offerings:



DCAI - IMPLEMENTING CISCO DATA CENTER AI INFRASTRUCTURE V1.0

Course Code: 860064

CLASSROOM LIVE

\$5,795 CAD

5 Day

Classroom Live Outline

Section 1: Fundamentals of AI

- Introduction to Artificial Intelligence
- Traditional AI
- Traditional AI Process Flow
- Traditional AI Challenges
- Modern Applications of Traditional AI
- Machine Learning vs. Deep Learning
- ML vs. DL Techniques and Methodologies
- ML vs. DL Applications and Use Cases

Section 2: Generative AI

- Generative AI
- Generative Adversarial Frameworks
- GenAI Use Cases
- Generative AI Inference Challenges
- GenAI Challenges and Limitations
- GenAI Bias and Fairness
- GenAI Resource Optimization
- Generative AI vs. Traditional AI
- Future Trends in AI
- AI Language Models
- LLMs vs. SLMs

Section 3: AI Use Cases

- Analytics
- Network Optimization
- Network Automation and Self-Healing Networks

- Capacity Planning and Forecasting
- Cybersecurity
- Predictive Risk Management
- Threat Detection
- Incident Response
- Collaboration and Communication
- Internet of Things (IoT)

Section 4: AI-ML Clusters and Models

- AI-ML Compute Clusters
- AI-ML Cluster Use Cases
- Custom AI Models-Process
- Custom AI Models-Tools
- Prebuilt AI Model Optimization
- Pre-Trained AI Models
- AI Model Parameters
- Service Placements - On-Premises vs. Cloud vs. Distributed

Section 5: AI Toolset-Jupyter Notebook

- AI Toolset-Jupyter Notebook

Section 6: AI Infrastructure

- Traditional AI Infrastructure
- Modern AI Infrastructure

Section 7: AI Workloads Placement and Interoperability

- Workload Mobility
- Multi-Cloud Implementation
- Vendor Lock-In Risks
- Vendor Lock-In Mitigation

Section 8: AI Policies

- Data Sovereignty
- Compliance, Governance, and Regulations

Section 9: AI Sustainability

- Green AI vs. Red AI
- Cost Optimization
- AI Accelerators
- Power and Cooling

Section 10: AI Infrastructure Design

- Project Description
- Your Role

Section 11: Key Network Challenges and Requirements for AI Workloads

- Bandwidth and Latency Considerations
- Scalability Considerations
- Redundancy and Resiliency Considerations

- Visibility
- Nonblocking Lossless Fabric
- Congestion Management Considerations

Section 12: AI Transport

- Optical and Copper Cabling
- Organizing Data Center Cabling
- Ethernet Cables
- InfiniBand Cables
- Ethernet Connectivity
- InfiniBand Connectivity
- Hybrid Connectivity

Section 13: Connectivity Models

- Network Types: Isolated vs. Purpose-Built Network
- Network Architectures: Two-Tier vs. Three-Tier Hierarchical Model
- Networking Considerations: Single-Site vs. Multi-Site Network Architecture

Section 14: AI Network

- Layer 2 Protocols
- Layer 3 Protocols
- Scalability Considerations for Deploying AI Workloads
- Fog Computing for AI Distributed Processing

Section 15: Architecture Migration to AI/ML Network

- Project Description
- Your Role

Section 16: Application-Level Protocols

- RDMA Fundamentals
- RDMA Architecture
- RDMA Operations
- RDMA over Converged Ethernet > NEW title RoCE/RoCEv2

Objectives:

- Understand the RDMA operations over Ethernet

Section 17: High-Throughput Converged Fabrics

- InfiniBand-to-Ethernet Transition
- Cisco Nexus 9000 Series Switches Portfolio

Section 18: Building Lossless Fabrics

- Traditional QoS Toolset
- Enhanced Transmission Selection
- Intelligent Buffer Management on Cisco Nexus 9000 Series Switches
- AFD with ETRAP
- Dynamic Packet Prioritization
- Data Center Bridging Exchange
- Lossless Ethernet Fabric Using RoCEv2
- Advanced Congestion Management with AFD

Section 19: Congestion Visibility

- Explicit Congestion Notification
- Priority Flow Control
- Congestion Visibility in AI/ML Cluster Networks Using Cisco Nexus Dashboard Insights
- Pipeline Considerations

Section 20: Data Preparation for AI

- Data Processing Workflow Overview
- Data Processing Workflow Phases

Section 21: AI/ML Workload Data Performance

- Use Cisco Nexus Dashboard Insights for monitoring AI/ML traffic flows

Section 22: AI-Enabling Hardware

- CPUs, GPUs, and DPUs
- GPU Overview
- NVIDIA GPUs for AI/ML
- Intel GPUs for AI/ML
- DPU Overview
- SmartNIC Overview
- Cisco Nexus SmartNIC Family
- NVIDIA BlueField SuperNIC

Section 23: Compute Resources

- Compute Hardware Overview
- Intel Xeon Scalable Processor Family Overview
- Cisco UCS C-Series Rack Servers
- Cisco UCS X-Series Modular System
- Mapping AI/ML Workloads to Cisco UCS Servers
- GPU Sharing
- Compute Resources Sharing
- Total Cost of Ownership
- AI/ML Clustering

Section 24: Compute Resource Solutions

- Cisco Hyperconverged Infrastructure Solutions Overview
- Cisco Hyperconverged Solution Components
- FlashStack Data Center
- Nutanix GPT-in-a-Box
- Run:ai on Cisco UCS

Section 25: Virtual Resources

- Virtual Infrastructure
- Device Virtualization
- Server Virtualization Defined
- Virtual Machine
- Hypervisor

- Container Engine
- Storage Virtualization
- Virtual Networks
- Virtual Infrastructure Deployment Options
- Hyperconverged Infrastructure
- HCI and Virtual Infrastructure Deployment

Section 26: Storage Resources

- Data Storage Strategy
- Fibre Channel and FCoE
- NVMe and NVMe over Fabrics
- Software-Defined Storage

Section 27: Setting Up AI Cluster

- Use NDFC to configure a fabric optimized for AI/ML workloads.

Section 28: Deploy and Use Open Source GPT Models for RAG

- Use locally-hosted GPT models with RAG for network engineering tasks.

Section 29: AI Infrastructure Operations and Monitoring

- The Need for AI Infrastructure Monitoring
- Monitoring Compute
- Monitoring Storage
- Monitoring the Runtime Layer
- Monitoring AI Fabrics
- The Need for AI Infrastructure Lifecycle Management
- Compute Lifecycle Upgrades
- Fabric Lifecycle Upgrades

Section 30: Troubleshooting AI Infrastructure

- Log Correlation for AI Applications
- Telemetry Analysis for AI Workloads
- Hands-On Telemetry for AI Workloads
- Timing Protocols

Section 31: Troubleshoot Common Issues in AI/ML Fabric

- Overview of Splunk Enterprise and Splunk Cloud
- Data Ingestion Methods
- Splunk Applications
- Basics of Splunk SPL

Classroom Live Labs

- Discovery 1: AI Toolset-Jupyter Notebook
- Design Practice Activity 1: AI Infrastructure Design
- Design Practice Activity 2: Architecture Migration to AI/ML Network
- Discovery 2: AI/ML Workload Data Performance
- Discovery 3: Setting Up AI Cluster

- Discovery 4: Deploy and Use Open Source GPT Models for RAG
- Discovery 5: Troubleshoot Common Issues in AI/ML Fabric



DCAI - IMPLEMENTING CISCO DATA CENTER AI INFRASTRUCTURE V1.0

Course Code: 860064

VIRTUAL CLASSROOM LIVE

\$5,795 CAD

5 Day

Virtual Classroom Live Outline

Section 1: Fundamentals of AI

- Introduction to Artificial Intelligence
- Traditional AI
- Traditional AI Process Flow
- Traditional AI Challenges
- Modern Applications of Traditional AI
- Machine Learning vs. Deep Learning
- ML vs. DL Techniques and Methodologies
- ML vs. DL Applications and Use Cases

Section 2: Generative AI

- Generative AI
- Generative Adversarial Frameworks
- GenAI Use Cases
- Generative AI Inference Challenges
- GenAI Challenges and Limitations
- GenAI Bias and Fairness
- GenAI Resource Optimization
- Generative AI vs. Traditional AI
- Future Trends in AI
- AI Language Models
- LLMs vs. SLMs

Section 3: AI Use Cases

- Analytics
- Network Optimization
- Network Automation and Self-Healing Networks

- Capacity Planning and Forecasting
- Cybersecurity
- Predictive Risk Management
- Threat Detection
- Incident Response
- Collaboration and Communication
- Internet of Things (IoT)

Section 4: AI-ML Clusters and Models

- AI-ML Compute Clusters
- AI-ML Cluster Use Cases
- Custom AI Models-Process
- Custom AI Models-Tools
- Prebuilt AI Model Optimization
- Pre-Trained AI Models
- AI Model Parameters
- Service Placements - On-Premises vs. Cloud vs. Distributed

Section 5: AI Toolset-Jupyter Notebook

- AI Toolset-Jupyter Notebook

Section 6: AI Infrastructure

- Traditional AI Infrastructure
- Modern AI Infrastructure

Section 7: AI Workloads Placement and Interoperability

- Workload Mobility
- Multi-Cloud Implementation
- Vendor Lock-In Risks
- Vendor Lock-In Mitigation

Section 8: AI Policies

- Data Sovereignty
- Compliance, Governance, and Regulations

Section 9: AI Sustainability

- Green AI vs. Red AI
- Cost Optimization
- AI Accelerators
- Power and Cooling

Section 10: AI Infrastructure Design

- Project Description
- Your Role

Section 11: Key Network Challenges and Requirements for AI Workloads

- Bandwidth and Latency Considerations
- Scalability Considerations
- Redundancy and Resiliency Considerations

- Visibility
- Nonblocking Lossless Fabric
- Congestion Management Considerations

Section 12: AI Transport

- Optical and Copper Cabling
- Organizing Data Center Cabling
- Ethernet Cables
- InfiniBand Cables
- Ethernet Connectivity
- InfiniBand Connectivity
- Hybrid Connectivity

Section 13: Connectivity Models

- Network Types: Isolated vs. Purpose-Built Network
- Network Architectures: Two-Tier vs. Three-Tier Hierarchical Model
- Networking Considerations: Single-Site vs. Multi-Site Network Architecture

Section 14: AI Network

- Layer 2 Protocols
- Layer 3 Protocols
- Scalability Considerations for Deploying AI Workloads
- Fog Computing for AI Distributed Processing

Section 15: Architecture Migration to AI/ML Network

- Project Description
- Your Role

Section 16: Application-Level Protocols

- RDMA Fundamentals
- RDMA Architecture
- RDMA Operations
- RDMA over Converged Ethernet > NEW title RoCE/RoCEv2

Objectives:

- Understand the RDMA operations over Ethernet

Section 17: High-Throughput Converged Fabrics

- InfiniBand-to-Ethernet Transition
- Cisco Nexus 9000 Series Switches Portfolio

Section 18: Building Lossless Fabrics

- Traditional QoS Toolset
- Enhanced Transmission Selection
- Intelligent Buffer Management on Cisco Nexus 9000 Series Switches
- AFD with ETRAP
- Dynamic Packet Prioritization
- Data Center Bridging Exchange
- Lossless Ethernet Fabric Using RoCEv2
- Advanced Congestion Management with AFD

Section 19: Congestion Visibility

- Explicit Congestion Notification
- Priority Flow Control
- Congestion Visibility in AI/ML Cluster Networks Using Cisco Nexus Dashboard Insights
- Pipeline Considerations

Section 20: Data Preparation for AI

- Data Processing Workflow Overview
- Data Processing Workflow Phases

Section 21: AI/ML Workload Data Performance

- Use Cisco Nexus Dashboard Insights for monitoring AI/ML traffic flows

Section 22: AI-Enabling Hardware

- CPUs, GPUs, and DPUs
- GPU Overview
- NVIDIA GPUs for AI/ML
- Intel GPUs for AI/ML
- DPU Overview
- SmartNIC Overview
- Cisco Nexus SmartNIC Family
- NVIDIA BlueField SuperNIC

Section 23: Compute Resources

- Compute Hardware Overview
- Intel Xeon Scalable Processor Family Overview
- Cisco UCS C-Series Rack Servers
- Cisco UCS X-Series Modular System
- Mapping AI/ML Workloads to Cisco UCS Servers
- GPU Sharing
- Compute Resources Sharing
- Total Cost of Ownership
- AI/ML Clustering

Section 24: Compute Resource Solutions

- Cisco Hyperconverged Infrastructure Solutions Overview
- Cisco Hyperconverged Solution Components
- FlashStack Data Center
- Nutanix GPT-in-a-Box
- Run:ai on Cisco UCS

Section 25: Virtual Resources

- Virtual Infrastructure
- Device Virtualization
- Server Virtualization Defined
- Virtual Machine
- Hypervisor

- Container Engine
- Storage Virtualization
- Virtual Networks
- Virtual Infrastructure Deployment Options
- Hyperconverged Infrastructure
- HCI and Virtual Infrastructure Deployment

Section 26: Storage Resources

- Data Storage Strategy
- Fibre Channel and FCoE
- NVMe and NVMe over Fabrics
- Software-Defined Storage

Section 27: Setting Up AI Cluster

- Use NDFC to configure a fabric optimized for AI/ML workloads.

Section 28: Deploy and Use Open Source GPT Models for RAG

- Use locally-hosted GPT models with RAG for network engineering tasks.

Section 29: AI Infrastructure Operations and Monitoring

- The Need for AI Infrastructure Monitoring
- Monitoring Compute
- Monitoring Storage
- Monitoring the Runtime Layer
- Monitoring AI Fabrics
- The Need for AI Infrastructure Lifecycle Management
- Compute Lifecycle Upgrades
- Fabric Lifecycle Upgrades

Section 30: Troubleshooting AI Infrastructure

- Log Correlation for AI Applications
- Telemetry Analysis for AI Workloads
- Hands-On Telemetry for AI Workloads
- Timing Protocols

Section 31: Troubleshoot Common Issues in AI/ML Fabric

- Overview of Splunk Enterprise and Splunk Cloud
- Data Ingestion Methods
- Splunk Applications
- Basics of Splunk SPL

Virtual Classroom Live Labs

- Discovery 1: AI Toolset-Jupyter Notebook
- Design Practice Activity 1: AI Infrastructure Design
- Design Practice Activity 2: Architecture Migration to AI/ML Network
- Discovery 2: AI/ML Workload Data Performance
- Discovery 3: Setting Up AI Cluster

- Discovery 4: Deploy and Use Open Source GPT Models for RAG
- Discovery 5: Troubleshoot Common Issues in AI/ML Fabric

Mar 23 - 27, 2026 | 9:00 AM - 5:00 PM EDT

May 11 - 15, 2026 | 9:00 AM - 5:00 PM EDT

Sep 28 - Oct 2, 2026 | 9:00 AM - 5:00 PM EDT

Visit us at www.globalknowledge.com or call us at 1-866-716-6688.

Date created: 1/10/2026 4:44:04 PM

Copyright © 2026 Global Knowledge Training LLC. All Rights Reserved.