CISCO 300-640

Cisco CCNP Data Center Certification Questions & Answers

Exam Summary – Syllabus – Questions

300-640

Cisco Certified Specialist - Data Center AI Infrastructure

55-65 Questions Exam – Variable (750-850 / 1000 Approx.)% Cut Score – Duration of 90 minutes

Table of Contents:

Know Your 300-640 Certification Well:

The 300-640 is best suitable for candidates who want to gain knowledge in the Cisco Data Center. Before you start your 300-640 preparation you may struggle to get all the crucial CCNP Data Center materials like 300-640 syllabus, sample questions, study guide.

But don't worry the 300-640 PDF is here to help you prepare in a stress free manner.
The PDF is a combination of all your queries like-

- What is in the 300-640 syllabus?
- How many questions are there in the 300-640 exam?
- Which Practice test would help me to pass the 300-640 exam at the first attempt?

Passing the 300-640 exam makes you Cisco Certified Specialist - Data Center AI Infrastructure. Having the CCNP Data Center certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

Cisco 300-640 CCNP Data Center Certification Details:

Exam Name	Implementing Cisco Data Center AI Infrastructure
Exam Number	300-640 DCAI
Exam Price	$300 USD
Duration	90 minutes
Number of Questions	55-65
Passing Score	Variable (750-850 / 1000 Approx.)
Exam Registration	**PEARSON VUE**
Sample Questions	**Cisco 300-640 Sample Questions**
Practice Exam	**Cisco Certified Specialist - Data Center AI Infrastructure Practice Test**

300-640 Syllabus:

Section	Weight	Objectives
AI Fundamentals and Applications	20%	- Describe AI/ML workload types • RAG • Training • Inference • Generative AI - Describe the AI lifecycle - Describe AI use cases - Describe the types of AI infrastructure • Cloud • Hybrid • On-premises • Edge AI - Describe the components used for AI environments • Network • Compute and GPUs deployment (NVLink) • Virtualization and containerization • Orchestration • Monitoring • Storage such as SAN, Fibre Channel, NVMe, Block and File - Describe Cisco AI solutions • AI PODs • AI Canvas • Hyperfabric AI
AI Infrastructure Components and Architecture	30%	- Evaluate network deployment based on AI workload requirements such as bandwidth, latency, redundancy, scalability, and security - Evaluate compute deployment based on AI workload requirements such as CPU resources, GPU resources and connectivity, memory, virtualization support, scalability redundancy, and workload types - Evaluate storage deployment based on AI workload requirements such as capacity,

Section	Weight	Objectives
		performance, redundancy and availability, and scalability - Evaluate power, efficiency, and sustainability based on AI workload requirements such as power and cooling, power usage effectiveness, and renewable energy - Evaluate hybrid AI deployment with cloud integration such as secure connectivity, data synchronization, and workload mobility
AI Infrastructure Deployment and Data Management	30%	- Configure high-performance networks to support AI workloads using Cisco Data Center • Congestion control mechanisms (PFC, ECN, ETS) • RDMA over Converged Ethernet (RoCE, RoCEv2) • Quality of service (QoS) • Load distribution - Configure high-performance compute and storage to support AI workloads using Cisco UCS • Domain profiles • Power policy • Storage policies • LAN connectivity and vNIC policies • QoS policies and system classes • NTP policy - Deploy AI-ready fabrics using Cisco orchestration tools • Nexus Dashboard • APIC • Hyperfabric • Intersight
AI Infrastructure Operations and Troubleshooting	20%	- Implement benchmarks to evaluate AI infrastructure performance - Implement monitoring of AI data center infrastructures using Cisco solutions such as Nexus Dashboard and Intersight - Monitor AI infrastructure using system messages

Section	Weight	Objectives
		and management tools to ensure reliability, scalability and performance • Operational telemetry • System health • Alerts • Log correlation - Troubleshoot AI infrastructure using system messages and management tools

Cisco 300-640 Sample Questions:

Question: 1

During which phase of the AI lifecycle would a data engineer perform "Feature Engineering" and "Data Labeling"?

a) Model Monitoring
b) Data Preparation
c) Model Evaluation
d) Model Deployment

Answer: d

Question: 2

Which AI use case is specifically designed to detect anomalies in network traffic patterns to prevent zero-day exploits?

a) Predictive Maintenance
b) Natural Language Processing (NLP)
c) AI-driven Cybersecurity
d) Computer Vision

Answer: c

Question: 3

Why is Redundancy (N+1 or 2N) more complex in AI clusters than in standard web-tier clusters?

a) AI jobs are stateful and long-running; a single node failure can crash a month-long training job.

a) b) AI nodes don't use IP addresses.
b) GPUs don't support power redundancy.
c) AI traffic is only one-way.

Answer: a

Question: 4

In Nexus Dashboard Fabric Controller (NDFC), which deployment "Persona" is specifically optimized for managing the leaf-spine fabrics typically used in AI PODs?

a) LAN Fabric
b) IP Fabric for Media
c) SDN Fabric (ACI)
d) SAN Controller

Answer: a

Question: 5

In a Hybrid AI deployment, which Cisco solution would be used to provide secure, high-speed connectivity between an on-premises Nexus fabric and a VPC in AWS?

a) Cisco SD-WAN or Cloud Router (CSR/Catalyst 8000V)
b) Cisco Umbrella
c) Cisco Duo
d) Cisco AnyConnect (Standard VPN)

Answer: a

Question: 6

What is the role of a Vector Database in a RAG architecture?

a) It stores the weights of the neural network.
b) It stores high-dimensional "embeddings" of data for fast similarity searches.
c) It acts as the primary DHCP server for the GPUs.
d) It manages the power distribution units (PDUs).

Answer: b

Question: 7

When configuring a storage policy for an AI cluster in Cisco Intersight, which setting ensures that the UCS server can utilize NVMe over Fabrics (NVMe-oF) for high-speed data ingestion?

a) Enabling "Local Disk Mirroring"
b) Configuring an "FC-NVMe" or "NVMe over RoCE" interface in the Storage Policy
c) Setting the disk drive speed to 15k RPM
d) Enabling "SATA AHCI" mode in the BIOS

Answer: b

Question: 8

A 32-node AI training job has slowed down significantly. An engineer checks the Nexus switch counters and sees a high number of PFC Rx Pause frames on the ports connected to the leaf switches, but zero drops. What is the most likely cause?

a) The fabric is dropping packets due to a hardware failure.
b) The spine switches are congested and are pushing "back-pressure" toward the leaves.
c) The GPUs have stopped sending data.
d) The VLANs are incorrectly configured.

Answer: b

Question: 9

An organization is using Cisco Intersight to manage a globally distributed AI infrastructure. Which feature allows them to automate the deployment of OS-level drivers and NVIDIA libraries across all nodes?

a) Intersight Mobile App
b) Intersight Search
c) Intersight Virtual Appliance
d) Intersight Cloud Orchestrator (ICO)

Answer: d

Question: 10

You are designing the IP addressing for a large AI cluster using RoCEv2 across a Layer 3 leaf-spine. Why is it recommended to use a small MTU for the Control Plane but a Jumbo MTU (9000+) for the Data Plane?

a) There is no such recommendation; MTU should always be 1500.
b) Because GPUs can only read large packets.
c) Small packets are better for routing updates, while large packets maximize throughput and reduce CPU overhead for the massive AI data transfers.
d) To save money on IP addresses.

Answer: c

Study Guide to Crack Cisco CCNP Data Center 300-640 Exam:

- Getting details of the 300-640 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the 300-640 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the Cisco provided training for 300-640 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the 300-640 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on 300-640 practice tests is must. Continuous practice will make you an expert in all syllabus areas.

Reliable Online Practice Test for 300-640 Certification

Make NWExam.com your best friend during your Implementing Cisco Data Center AI Infrastructure exam preparation. We provide authentic practice tests for the 300-640 exam. Experts design these online practice tests, so we can offer you an exclusive experience of taking the actual 300-640 exam. We guarantee you 100% success in your first exam attempt if you continue practicing regularly. Don't bother if you don't get 100% marks in initial practice exam attempts. Just utilize the result section to know your strengths and weaknesses and prepare according to that until you get 100% with our practice tests. Our evaluation makes you confident, and you can score high in the 300-640 exam.

Start Online practice of 300-640 Exam by visiting URL
https://www.nwexam.com/cisco/300-640-implementing-cisco-data-center-ai-infrastructure-dcai