

# cisco 300-640

## Implementing Cisco Data Center AI Infrastructure (DCAI)

For More Information – Visit link below:

<https://www.examsempire.com/>

**Product Version**

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/300-640>

# Latest Version: 6.0

## Question: 1

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: 2

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: 3

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: 4

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**

### Question: 5

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.
- B. AI nodes don't use IP addresses.
- C. GPUs don't support power redundancy.
- D. AI traffic is only one-way.

**Answer: A**

### Question: 6

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: 7

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: 8

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: 9

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: 10

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**

**Thank You for Trying Our Product**  
**Special 16 USD Discount Coupon: NSZUBG3X**  
**Email: [support@examsempire.com](mailto:support@examsempire.com)**

**Check our Customer Testimonials and ratings  
available on every product page.**

**Visit our website.**

**<https://examsempire.com/>**