

Cisco 300-540


Designing and Implementing Cisco Service Provider Cloud Network Infrastructure v1.0 Exam

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/300-540>

Latest Version: 6.0

Question: 1

What is a benefit of using VXLANs in a cloud-scale environment?

- A. extends Layer 2 segments across the underlying Layer 3 infrastructure
- B. extends Layer 3 segments across the underlying Layer 2 infrastructure
- C. reduces spanning-tree complexity across the Layer 2 infrastructure
- D. eliminates the need for a Layer 3 underlay in the service provider infrastructure

Answer: A

Explanation:

In a cloud-scale or data center-scale environment, Virtual Extensible LAN (VXLAN) is used as an overlay technology to transport Layer 2 segments over a Layer 3 underlay network. VXLAN encapsulates Layer 2 Ethernet frames inside UDP/IP packets, allowing broadcast, unknown unicast, and multicast (BUM) traffic and tenant Layer 2 domains to be extended across a routed IP fabric.

Key points aligned with Cisco Service Provider Cloud Infrastructure design principles:

VXLAN creates a Layer 2 overlay on top of a Layer 3 underlay.

The VXLAN Network Identifier (VNI) provides a much larger segmentation space than traditional VLANs, enabling multi-tenancy at cloud scale.

Because the underlay is pure Layer 3 (IP routed fabric), VXLAN allows you to interconnect Layer 2 segments between leaf switches or data centers over an IP/MPLS backbone without relying on large Layer 2 domains in the physical network.

Why the options evaluate as follows:

Option A: extends Layer 2 segments across the underlying Layer 3 infrastructure ☑

This is the core benefit of VXLAN in cloud-scale designs. VXLAN encapsulates Layer 2 frames into IP/UDP headers, allowing isolated Layer 2 segments (per VNI) to be stretched across a routed IP network. This enables:

Multi-tenant Layer 2 connectivity across a distributed cloud fabric

Mobility of virtual machines or containers while keeping same IP/MAC addressing

Use of an IP-based leaf-spine or service provider underlay for scalability and resiliency

Option B: extends Layer 3 segments across the underlying Layer 2 infrastructure ☐

This is the opposite of what VXLAN does. VXLAN is explicitly L2-over-L3, not L3-over-L2. Extending pure Layer 3 segments over Layer 2 is not the VXLAN use case.

Option C: reduces spanning-tree complexity across the Layer 2 infrastructure ☐☐ (Partially related but not the primary or direct benefit)

In modern designs, the underlay is Layer 3 routed, and VXLAN overlays provide logical Layer 2 segments. This design avoids dependence on spanning tree in the fabric, which indirectly reduces STP complexity. However, the fundamental, exam-relevant benefit is L2 extension over L3, so C is not the best or most accurate answer compared to A.

Option D: eliminates the need for a Layer 3 underlay in the service provider infrastructure ☐

VXLAN absolutely requires an IP (Layer 3) underlay for transport. VXLAN tunnels are built over a routed

infrastructure (leaf–spine, MPLS/IP core, etc.). It does not remove the need for Layer 3; it depends on it.

Question: 2

An engineer must configure NTP servers in Cisco Enterprise NFVIS. The primary NTP server has an IP address of 192.168.1.1 and the backup NTP server has an IP address of 192.168.2.1. Which two commands must be run to complete the configuration? (Choose two.)

- A. `system time ntp preferred_server 192.168.1.1`
- B. `utils ntp server add 192.168.2.1 backup`
- C. `system set-manual-time 192.168.1.1 192.168.2.1`
- D. `utils ntp server add 192.168.1.1 primary`
- E. `system time ntp backup_server 192.168.2.1`

Answer: A, E

Explanation:

In Cisco Enterprise NFVIS, time synchronization is configured using the `system time ntp` command structure. NFVIS requires a primary and optionally a backup NTP server to maintain accurate system time for the hypervisor and guest VMs.

Correct NFVIS command syntax for NTP configuration:

`system time ntp preferred_server <IP>`

This command configures the preferred (primary) NTP server used for system clock synchronization.

`system time ntp backup_server <IP>`

This command configures the backup NTP server, which the system uses if the primary becomes unreachable.

These two commands match Cisco NFVIS time-configuration behavior described in NFV infrastructure design and implementation guidelines.

Why the Correct Answers Are A and E

Option A: `system time ntp preferred_server 192.168.1.1`

This properly configures the primary NTP server in NFVIS. The preferred server is always the first choice for time synchronization.

Option E: `system time ntp backup_server 192.168.2.1`

This correctly configures the backup NTP server. If the preferred server fails, NFVIS automatically falls back to the backup server.

Both commands directly match NFVIS's NTP command hierarchy and are the only ones that correctly apply to NFVIS.

Why the Other Options Are Not Correct


Option B uses `utils ntp`, which is not an NFVIS command.

Option C sets manual time and does not configure NTP servers.

Option D also uses the `utils ntp` syntax, which applies to other Cisco platforms but not NFVIS.

Question: 3

Refer to the exhibit.



An engineer must connect switch TOR1 and switch TOR2 to switch LEAF1 and switch LEAF2 by using double-sided vPCs. LEAF1 and LEAF2 are already configured as vPC peers. Which action must be taken next to complete the configuration?

- A. Add all the switches to the fabric.
- B. Configure peering between LEAF1 and LEAF2 and TOR1 and TOR2.
- C. Configure MSTP between TOR1 and TOR2.
- D. Configure a vPC between TOR1 and TOR2.

Answer: D

Explanation:

In Cisco data center and cloud-scale designs, a double-sided vPC (also called vPC-to-vPC) is used when both ends of a Layer 2 port channel are formed by a pair of switches that operate as vPC peers. In this model:

On the aggregation or leaf side, two switches (in this case, LEAF1 and LEAF2) form a vPC domain with a vPC peer-link and keepalive.

On the access or ToR side, two switches (in this case, TOR1 and TOR2) must also form their own vPC domain with a peer-link and vPC keepalive.

The port-channel that interconnects the two vPC domains is then configured as a vPC on both sides, creating a vPC-to-vPC topology.

The problem statement specifies that LEAF1 and LEAF2 are already configured as vPC peers. For a double-sided vPC to work, the other side (TOR1 and TOR2) must also behave as a single logical entity for the downstream Cisco UCS server and for the upstream vPC connection towards LEAF1 and LEAF2. This is only achieved when TOR1 and TOR2 are configured as vPC peers with:

A vPC domain ID

A vPC peer-link between TOR1 and TOR2

vPC member port-channels towards LEAF1 and LEAF2 and towards the Cisco UCS server

Therefore, the next required step is to configure a vPC between TOR1 and TOR2.

Evaluation of the options:

Option A, "Add all the switches to the fabric," is generic and not specific to vPC configuration. It does not address the technical requirement to form a vPC domain on the ToR side.

Option B, "Configure peering between LEAF1 and LEAF2 and TOR1 and TOR2," is incorrect because vPC

peering is only configured between the two switches that form each vPC domain (LEAF1–LEAF2 and TOR1–TOR2), not across all four switches together.

Option C, “Configure MSTP between TOR1 and TOR2,” is not required for establishing a double-sided vPC. vPC designs rely on the vPC control plane and the peer-link, not on spanning-tree between the vPC peers for normal operation.

Option D, “Configure a vPC between TOR1 and TOR2,” correctly describes configuring TOR1 and TOR2 as a vPC pair (vPC domain with peer-link), which is the mandatory step to create a double-sided vPC topology with LEAF1 and LEAF2.

Question: 4

What is a valid connection method between carrier-neutral facilities within the same metro area?

- A. OSPF backbone area adjacency
- B. private wireless connection
- C. DWDM ring
- D. CAT6e connection

Answer: C

Explanation:

Comprehensive and Detailed Explanation Based on Designing and Implementing Cisco Service Provider Cloud Network Infrastructure Knowledge

When connecting carrier-neutral facilities (CNFs) or data centers within the same metropolitan area, service providers typically use high-bandwidth, low-latency optical transport methods. The most appropriate and commonly deployed interconnection technology is:

DWDM (Dense Wavelength Division Multiplexing) ring, which provides:

High capacity (10G, 40G, 100G, 400G)

Low latency Redundancy through ring or mesh topologies

Multi-wavelength multiplexing for cost efficiency

Carrier-grade reliability for metro interconnect services

This aligns with cloud interconnect and metro transport design used in service provider environments.

Evaluation of the Options

A . OSPF backbone area adjacency

This is a routing protocol adjacency, not a physical connection method. It requires a transport link underneath but does not represent the physical interconnect itself.

B . Private wireless connection

Not suitable for CNF or metro DC interconnect because it lacks the bandwidth, reliability, and deterministic performance required for large-scale carrier-grade interconnects.

C . DWDM ring


This is the correct method. DWDM-based metro fiber rings are the standard for connecting carrier neutral facilities in the same metro region.

D . CAT6e connection

This is limited to short-distance copper Ethernet (tens of meters). It is not used for metro-scale interconnects or between CNFs.

Question: 5

Refer to the exhibit.



Refer to the exhibit. The indicated configuration was applied to a Cisco switch Switch_A located in the Los Angeles DC data center; however, Switch_A fails to establish OTV connectivity to Cisco switch Switch_C. Which overlay interface command must be run on Switch_A to resolve the issue?

- A. otv extend-vlan 101-111
- B. otv isis authentication-type md5
- C. otv isis authentication-check
- D. otv join-interface vlan 101-111

Answer: A

Explanation:

Overlay Transport Virtualization (OTV) allows Layer 2 extension across Layer 3 infrastructures. To operate, OTV requires three fundamental components on the overlay interface:

Join interface – used to reach the OTV control plane over L3 (already configured: otv join-interface g1/0).

Control-group multicast address – for control-plane advertisement (already configured: otv controlgroup 224.1.1.1).

Extended VLAN list – specifies which VLANs will be transported through the OTV overlay.

The configuration shown in the exhibit includes the join-interface, control-group, and data-group, but it

does NOT specify which VLANs should be extended. Without the `otv extend-vlan` command, OTV will form the overlay interface but will not forward any Layer 2 information, preventing adjacency and MAC distribution between sites.

In OTV, the command required to activate VLANs for transport is:

`otv extend-vlan <vlan-range>`

This enables the VLANs (such as 101–111) to be carried across the OTV overlay, completing the configuration and establishing connectivity.

Why the Other Options Are Incorrect

B . `otv isis authentication-type md5`

This is optional and only required if ISIS authentication is enabled on both edges. It does not resolve the absence of VLAN extension.

C . `otv isis authentication-check`

This command enforces authentication verification but does not fix connectivity when VLANs are not extended.

D . `otv join-interface vlan 101-111`

This is not a valid OTV command. The join-interface must be a routed interface, not a VLAN list.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

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

Visit our website.

<https://examsempire.com/>