



DEMO VERSION

Cisco

300-540 Exam

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



Exam Latest Version: 6.0



Question 1. (Single Select)

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

Correct 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

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

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 (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

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. (Multi Select)

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

Correct 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. (Single Select)

An engineer must add VNF implementation definitions and VNF service definitions to an OpenStack deployment data model to deploy virtual routers and firewalls to an enterprise network. After the virtual machine resources are prepared, which action must be taken next?

- A: Specify the operational characteristics.
- B: Create the initial configuration for day zero.
- C: Configure key performance indicator monitoring.
- D: Define the NFV network.

Correct Answer: D

Explanation:

In an NFV environment following Cisco NFV Infrastructure and ETSI MANO architecture, the onboarding workflow for VNF deployment follows this high-level order:

Define VNF descriptors (VNFD and NSD)

- Includes implementation details and service definitions.

Prepare virtual machine resources

- Flavor, image, CPU/memory/storage, and placement.

Define the NFV network

- This step builds the underlay/overlay connectivity needed for the VNF.
- It includes virtual network creation, interfaces, VLs (Virtual Links), connection points, and mapping into OpenStack (Neutron networks, subnets, ports).
- No VNF can be deployed until the NFV service networks (management, control, data, HA) are defined.

Apply Day-0 configuration

- Bootstrapping/initial config.

Apply Day-1 operational characteristics

– Routing, services, policies.

Configure operational monitoring (KPIs)

– Performance and fault measurements.

Since the question states “after the VM resources are prepared”, the next mandatory stage is defining the virtual network connectivity on which the VNFs will operate.

Question 4. (Single Select)

A large company's legacy network is set up with equipment from multiple vendors. The company engaged a network architect to optimize the network for virtualization. The architect must ensure robust and efficient operation, considering the company's immediate needs but also anticipating future network complexities and scalability requirements. The chosen strategy must be capable of integrating seamlessly with existing systems, while providing a pathway for innovation and growth. The solution must facilitate end-to-end service automation throughout the entire lifecycle, and the implementation must ensure the validation, execution, and abstraction of network configurations and services. Which action must be taken to meet the requirements?

A: Implement a service life-cycle approach with simplified monitoring that plans for post-deployment adjustments to be incorporated into the automation CI/CD pipeline.

B: Implement a configuration-management approach that allows for configuring each network device individually to optimize its performance.

C: Implement a flexible service-modeling approach that leverages automation for ongoing management and refinement as demands on the network evolve.

D: Implement a service-modeling approach with a static YANG one-size-fits-all model that includes the unique requirements of each different network element.

Correct Answer: C

Explanation:

Cisco NSO-based orchestration principles in a multi-vendor environment require:

Service modeling using flexible, reusable YANG models

Abstraction of vendor-specific device differences

Transaction-safe configuration validation and execution

End-to-end automation across lifecycle stages (Day-0, Day-1, Day-N)

Scalability and adaptability for evolving requirements

Option C aligns perfectly with NSO service-modeling approaches:

Service models must be flexible, not rigid, enabling changes as technologies and needs evolve.

The architecture must support continuous refinement, enabling multi-vendor abstraction and lifecycle automation.

This ensures the network evolves seamlessly while remaining stable and automated.

Why the Other Options Are Incorrect

A – Simplified monitoring and post-deployment adjustments do not meet the core need for full lifecycle service modeling and abstraction.

B – Configuring devices individually contradicts the entire purpose of orchestration and abstraction.

D – A static YANG model cannot accommodate multi-vendor environments or future scalability.

Thus, only Option C matches full NSO-capable service modeling requirements.

Question 5. (Single Select)

Which cloud provider connection permits BGP peering?

A: Azure S2S VPN

B: Azure Bastion

C: AWS Direct Connect

D: AWS-managed VPN

Correct Answer: C

Explanation:

Comprehensive and Detailed Explanation

Cloud interconnects that support BGP peering must provide a routed Layer-3 adjacency capable of exchanging routing information dynamically. In major cloud architectures:

AWS Direct Connect supports private virtual interfaces (VIFs) where BGP is used between the customer router and AWS to exchange routes.

Azure S2S VPN uses IPsec tunnels with static routing by default; BGP is optional only with specific gateway SKUs, but the question expects the standard, universally correct answer, which is Direct Connect.

Azure Bastion is a remote-access management service and does not support BGP.

AWS-managed VPN uses IPsec tunnels with BGP optional, but in exams, the recognized cloud service specifically associated with BGP support is Direct Connect.

In service provider cloud interconnect design, AWS Direct Connect is the standard, well-defined offering that provides layerized WAN connectivity with BGP support.

ExamsIndex

Demo PDF Complete

Your 300-540 Demo (5 Questions)

Get the Complete Version

Full Questions with Detailed Explanations

Interactive Web-Based Exams Available

To get 30% off, use Coupon Code: NEWYEAR30

<https://examsindex.com/exam/300-540>