



VMware VCAP VMware Cloud Foundation Networking 3V0-25.25 Practice Questions

Shared by Mendoza on 17-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

An administrator has a standalone vSphere 8.0 Update 1a deployment that is running with VMware NSX 4.1.0.2 and has to converge the deployment into a new VMware Cloud Foundation (VCF) instance. How can the administrator accomplish this task?

Options:

- A- Manually upgrade both vSphere and NSX to version 9 prior to converging. Then use the VCF Installer to converge the vSphere 9 and NSX 9 instances into a new VCF management domain.
- B- Manually upgrade vSphere to version 9. Then use the VCF Installer to converge the vSphere 9 environment into a new VCF management domain. Then use the VCF lifecycle management tools to upgrade NSX to version 9.
- C- Use the VCF Installer to converge the existing vSphere 8 and NSX 4 environment into a new VCF management domain. Then use the VCF lifecycle management tools to upgrade to 9.
- D- Manually upgrade vSphere to version 9 and uninstall NSX 4. Then use the VCF Installer to converge the vSphere 9.0 environment into a new VCF management domain at which time NSX 9 will be reinstalled.

Answer:

C

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

The process of bringing existing infrastructure under VCF management is known as 'VCF Import' or 'Convergence.' This is a common path for organizations transitioning from siloed management to the full SDDC stack provided by Cloud Foundation.

According to the VCF 5.x and 9.0 documentation, the VCF Installer (specifically the Cloud Foundation Builder and the Import Tool) is designed to ingest existing environments. The verified best practice is to converge the environment at its current, supported version, provided it meets the minimum baseline requirements for the VCF version you are deploying.

In this scenario, vSphere 8.0 U1 and NSX 4.1 are compatible versions that can be imported into a VCF management framework. By using the VCF Installer to converge the existing environment first (Option C), the SDDC Manager takes ownership of the existing vCenter and NSX Manager. Once the environment is 'VCF-aware,' the administrator gains the benefit of SDDC Manager's Lifecycle Management (LCM).

The SDDC Manager then handles the orchestrated, multi-step upgrade to version 9.0. This

ensures that the automated 'Bill of Materials' (BOM) is strictly followed, ensuring compatibility between vCenter, ESXi, and NSX components. Attempting to manually upgrade components to version 9 before convergence (Options A and B) or uninstalling NSX (Option D) creates a 'Frankenstein' environment that may not align with the VCF BOM, making the automated convergence process fail or resulting in an unsupported configuration. The principle of VCF is to bring the environment in first, then let VCF manage the upgrades.

Question 2

Question Type: MultipleChoice

A cloud service provider runs VPCs with differing traffic patterns:

- * Some VPCs are generating high, large North/South flows.
- * Most of the VPCs generate very little traffic.

The architect needs to optimize Edge dataplane resource consumption while ensuring that noisy VPCs do not impact others.

Which optimization satisfies the requirement?

Options:

- A- Assign one dedicated Edge node per high-traffic VPC.
- B- Reduce the number of VPCs by consolidating VPCs into shared namespaces.
- C- Convert high-traffic VPCs into VLAN-backed segments attached directly to Tier-0 gateways.
- D- Use multiple Edge clusters and distribute VRF-backed VPCs based on traffic profiles.

Answer:

D

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

In a VMware Cloud Foundation (VCF) environment, especially with the architectural evolution in VCF 9.0, the Virtual Private Cloud (VPC) model is the primary way to deliver self-service, isolated networking. The networking performance for North/South traffic---traffic leaving the SDDC for the physical network---is processed by NSX Edge Nodes. These Edge Nodes use DPDK (Data Plane Development Kit) to provide high-performance packet processing, but their

resources (CPU and Memory) are finite.

When dealing with 'noisy neighbors'---tenants or VPCs that consume a disproportionate amount of throughput---it is critical to isolate their data plane impact. According to the VMware Validated Solutions and VCF Design Guides, the most scalable and efficient way to achieve this is through the use of Multiple Edge Clusters. By creating distinct Edge clusters, an architect can physically isolate the compute resources used for routing.

In this scenario, high-traffic VPCs can be backed by specific VRF (Virtual Routing and Forwarding) instances on a Tier-0 gateway that is hosted on a dedicated high-performance Edge Cluster. Meanwhile, the numerous low-traffic VPCs can share a different Edge Cluster. This 'Traffic Profile' based distribution ensures that a spike in traffic within a 'heavy' VPC only consumes the DPDK cycles of its assigned Edge nodes, leaving the resources for the 'quiet' VPCs untouched.

Option A is incorrect because Edge nodes function in clusters for high availability; assigning a single node creates a single point of failure and is administratively heavy. Option B reduces the multi-tenancy benefits and doesn't solve the resource contention at the Edge level. Option C removes the benefits of the software-defined overlay and VPC consumption model. Therefore, distributing VRF-backed VPCs across multiple Edge clusters based on their expected load is the verified design best practice for optimizing resource consumption while maintaining strict performance isolation in a VCF provider environment.

=====

Question 3

Question Type: MultipleChoice

An administrator changed the SFTP server used for scheduled NSX Manager backups. The backup jobs now fail with the error "Host KEY Verification Failed." The connectivity and credentials are correct. How would an administrator resolve the error?

Options:

- A- Turn Off Backup encryption.
- B- Update the SSH fingerprint.
- C- Trust the certificate on the SFTP server.
- D- Use the NSX cluster VIP as the SFTP endpoint.

Answer:

B

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

In VMware Cloud Foundation (VCF), the NSX Manager uses the SFTP protocol to securely transfer configuration backups to an external repository. SFTP is built on top of the SSH protocol, which relies on a 'Trust on First Use' (TOFU) model for verifying the identity of the remote host.

When an NSX Manager first connects to an SFTP server, it retrieves the server's SSH Public Key Fingerprint and stores it in its local known_hosts equivalent database. This fingerprint ensures that future connections are made to the same, verified server, preventing man-in-the-middle attacks.

The error 'Host KEY Verification Failed' occurs when the administrator changes the SFTP server (or if the SFTP server's OS was reinstalled/keys regenerated). Even if the IP address remains the same, the new server presents a different SSH fingerprint than the one currently cached in the NSX Manager configuration. Because the signatures do not match, the NSX Manager aborts the connection for security reasons.

To resolve this issue, the administrator must Update the SSH fingerprint (Option B) within the NSX Manager backup settings. This involves:

Retrieving the new fingerprint from the SFTP server (e.g., via ssh-keyscan).

Navigating to System > Lifecycle > Backup & Restore in the NSX Manager.

Editing the File Server configuration and pasting the new fingerprint into the appropriate field.

Option A is incorrect as it does not address the SSH protocol handshake failure. Option C is incorrect because SFTP/SSH uses fingerprints, not SSL/TLS certificates. Option D is irrelevant as it changes the source/destination of the connection but does not fix the underlying trust mismatch. Therefore, updating the fingerprint is the verified operational step to restore the automated backup workflow in VCF.

Question 4

Question Type: MultipleChoice

When attempting to deploy or expand an edge cluster from an administrator encounters a failure: "Failed to validate the BGP Route Distribution". Prior to calling support, the administrator attempts to troubleshoot the issue. How should the administrator troubleshoot this issue?

Options:

- A- Log into the NSX manager and examine the nsxapi.log for errors.
- B- Log into the Tier-1 router to verify that route distribution is being enabled.
- C- Log into the vCenter and verify there are no errors or warnings from the NSX manager.
- D- Log into the edge node of the Tier-0 being deployed and check the routes being learnt.

Answer:

D

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

In VMware Cloud Foundation (VCF), the SDDC Manager automates the deployment and expansion of NSX Edge Clusters. As part of the automated workflow, particularly in VCF 4.x, 5.x, and 9.0, a 'Verify BGP Route Distribution' task is executed. This task is a validation check designed to ensure that the newly deployed or expanded Edge nodes are successfully peering with the physical Top-of-Rack (ToR) switches and, more importantly, are actually receiving routes.

According to VMware/Broadcom technical documentation (specifically KB 388351), the workflow expects to see at least one route (often the default route or specific physical prefixes) learned via BGP from the northbound peer. If the Edge nodes establish a BGP session but the physical switches are not advertising any routes (or are only advertising routes that the Edge ignores due to filters), the SDDC Manager validation fails with the error 'Failed to validate the BGP Route Distribution'.

The verified troubleshooting step is to log into the CLI of the Edge node identified in the failure. Using the command `get route bgp` from within the Tier-0 Service Router (SR) VRF context allows the administrator to see the current Routing Information Base (RIB). If the table is empty or only contains internal 'ISR' (Inter-SR) routes, it confirms that the physical network is not providing the expected advertisements. This allows the administrator to correct the BGP advertisement settings on the physical ToR switches---such as enabling default-originate---and then simply 'Resume' the task in SDDC Manager without needing to redeploy the entire cluster.

Question 5

Question Type: MultipleChoice

An administrator must prevent a new VPC from exporting any of its prefixes to the datacenter while still receiving a default route. Where should the routing policy be applied?

Options:

- A- On the VPC default route advertiser
- B- On the VPC's Transit Gateway
- C- On the providers' BGP peer template
- D- On the VPC Gateway Firewall

Answer:

B

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

In the advanced networking architecture of VMware Cloud Foundation (VCF) 9.0 and the evolution of NSX VPCs, the control of route propagation is managed through the relationship between the consumer (the VPC) and the provider (the Tier-0 or Tier-1 Gateway). When a VPC is created, it is logically connected to the provider's infrastructure via a Transit Gateway (or a Provider-side logical router acting as a transit point).

To control the flow of routing information---specifically to prevent the data center's physical network from learning about internal VPC subnets (prefixes) while ensuring the VPC can still reach the outside world via a default route---the routing policy must be applied at the point of intersection. The Transit Gateway serves as this demarcation point. By applying a route filter or prefix list on the Transit Gateway, the administrator can explicitly deny the advertisement of internal VPC prefixes 'upstream' to the provider's BGP process. Simultaneously, the provider can still inject or 'advertise' a default route (0.0.0.0/0) 'downstream' into the VPC.

Applying the policy on the VPC Gateway Firewall (Option D) would impact the data plane (blocking traffic) but would not prevent the routing table from being populated. The BGP peer template (Option C) is too broad, as it would likely affect all VPCs connected to that provider, rather than just the 'new VPC' in question. The default route advertiser (Option A) only controls the egress of the default route, not the suppression of internal prefixes. Therefore, the Transit Gateway is the verified location for granular route control in a multi-tenant VCF VPC environment.

=====

Question 6

Question Type: MultipleChoice

An administrator needs to prevent the datacenter from advertising any internal prefixes toward

a new VPC, while still ensuring the VPC receives a default route learned from the datacenter's upstream network. Where should the routing policy be applied?

Options:

- A- On each segment default gateway.
- B- On the Tier-1 gateway.
- C- On the VPC transit gateway.
- D- On the provider Tier-0 neighbor.

Answer:

C

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

In the VMware Cloud Foundation (VCF) 9.0 and NSX VPC architecture, the Transit Gateway (TGW) is the central routing element that interconnects VPCs to each other and to the provider's infrastructure (Tier-0 or VRF gateways). It acts as the 'Project-level' gateway that aggregates North-South traffic.

To control the visibility of routes within a specific VPC, the administrator must utilize Route Filtering at the VPC's boundary. When a VPC is attached to a Transit Gateway, a logical interface is created. To prevent the data center's internal prefixes (such as management networks or other tenant subnets) from being seen by the VPC while still providing a path to the internet, a prefix list or route map should be applied to the VPC Transit Gateway. This policy will explicitly 'Deny' specific internal CIDR ranges while 'Permitting' the \$0.0.0.0/0\$ default route advertisement from the provider.

Applying the policy at the Tier-1 gateway (Option B) is technically similar but in the VPC model, the 'Tier-1' is often an obscured or automated component of the VPC itself; the Transit Gateway is the designed administrative point for inter-project and North-South policy enforcement. Applying it at the provider Tier-0 neighbor (Option D) would be too global, affecting all VPCs or projects connected to that Tier-0, rather than the 'new VPC' specifically. Therefore, the Transit Gateway provides the necessary granular control for multi-tenant isolation and routing optimization as per the VCF 9.0 networking model.

=====

Question 7

Question Type: MultipleChoice

Which of the following statements is true when configuring Remote Tunnel End Points (RTEPs) with NSX Federation?

Options:

- A- TEP and RTEP networks must use separate physical NICs.
- B- RTEP needs to be configured on only one edge node.
- C- The default MTU for the RTEP network is 1500.
- D- DHCP must be used to assign IP addresses to the RTEP.

Answer:

C

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

In an NSX Federation deployment, which is a key component of multi-site VMware Cloud Foundation (VCF) architectures, the Remote Tunnel End Point (RTEP) is used specifically for inter-site communication. While standard TEPs (Tunnel End Points) handle overlay traffic within a single site (East-West), RTEPs facilitate the encapsulation of traffic that needs to traverse the Layer 3 network between different geographical locations.

A critical design consideration for RTEP is the Maximum Transmission Unit (MTU). Within a local VCF site, jumbo frames (MTU 1600 or 9000) are highly recommended and often required for the Geneve overlay to account for encapsulation overhead. However, when traffic leaves a site to travel over a WAN or a provider's long-haul network, it often encounters physical infrastructure that only supports the standard internet MTU of 1500 bytes.

According to VMware's 'NSX Federation Design Guide,' the default MTU setting for the RTEP configuration is 1500. This ensures that inter-site traffic can pass through standard routers and VPNs without being dropped due to size constraints. If the inter-site physical links support larger frames, this value can be increased, but 1500 remains the baseline compatible default.

Regarding the other options: A is incorrect because TEP and RTEP can share the same physical N-VDS and physical NICs (pNICs) by using different VLANs or subnets. B is incorrect because every Edge node within a cluster that is participating in the Federation must have an RTEP configured to ensure high availability and proper traffic processing for global segments. D is incorrect as IP addresses for RTEPs are typically assigned via Static IP Pools managed within NSX to ensure consistency and ease of tracking across sites, rather than relying on DHCP which is less common in data center backbone configurations.

=====

Question 8

Question Type: MultipleChoice

An administrator is configuring NSX resource sharing to allow shared access to multiple resources in the default space.

By default, which user role owns the shared resources for the default space?

Options:

- A- Network Admin
- B- Security Admin
- C- Project Admin
- D- Enterprise Admin

Answer:

D

Explanation:

In NSX Multi-Tenancy (Projects), the Enterprise Admin acts as the provider-level administrator who owns global objects in the default space. This ensures central control over resources that are shared across different projects.

Question 9

Question Type: MultipleChoice

An administrator is troubleshooting an issue where workloads connected to a Tier-1 Gateway named T1-App can no longer reach external North/South destinations.

* The Tier-1 is connected to an Active/Standby Tier-0 Gateway named T0-Prod.

Symptoms observed:

- * VMs on segments attached to T1-App can ping each other.
- * VMs on T1-App cannot reach any external IP outside T0-Prod.

- * From a VM on the segment, ping to the T1-App Distributed Router (DR) IP succeeds.
- * Ping from the VM to the T1-App Service Router (SR) fails.
- * The Edge cluster hosting the T1-App SR shows both Edge nodes Up and Healthy.
- * No failover has occurred --- the same Edge node is still shown as Active for T1-App.

What is the most likely cause of this issue?

Options:

- A- The overlay network between DR and SR has an MTU mismatch.
- B- Route advertisement from T1-App to T0-Prod for 100.64.x.x/31 is disabled.
- C- Static default route is missing on the Tier-1 DR component.
- D- Localized control plane is enabled on the Tier-1 causing the SR to remain admin-down.

Answer:

A

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

In the NSX multi-tier routing architecture used by VCF, a Tier-1 Gateway is composed of two primary components: the Distributed Router (DR) and the Service Router (SR). The DR runs as a kernel module on every ESXi host in the transport zone, facilitating East-West traffic. The SR resides on the NSX Edge nodes and provides centralized services like North-South connectivity and stateful services.

Communication between the DR (on the ESXi host) and the SR (on the Edge node) occurs over a hidden internal segment known as the Router Link. This link is encapsulated in Geneve just like VM-to-VM traffic. When a VM attempts to reach an external destination, the packet is first routed by the DR on the local host. The DR then encapsulates the packet and sends it across the overlay to the TEP (Tunnel Endpoint) of the Edge node hosting the SR.

If the MTU (Maximum Transmission Unit) is misconfigured on the physical network or the virtual switches, large encapsulated packets will be dropped. However, small packets (like pings between VMs on the same host) might still succeed. In this scenario, the fact that the VM can ping the local DR but cannot reach the SR---and therefore cannot reach external networks---points to a failure in the transport between the host and the Edge.

If the Geneve-encapsulated packet containing the ping request to the SR's internal interface exceeds the physical network's MTU, it will fail. Since VCF 5.x/9.0 requires a minimum MTU of 1600 (ideally 9000) for the overlay to account for the Geneve overhead, a mismatch anywhere in the fabric will break the DR-to-SR 'backplane' communication. This prevents the Tier-1 from passing any traffic to its Tier-0 uplink, effectively isolating the workloads from North-South

traffic.

Question 10

Question Type: MultipleChoice

An administrator is tasked to enable users to configure an individual VPC, but not create subnets. What three NSX roles would the administrator assign to allow access without the ability to create subnets? (Select three.)

Options:

- A- Security Admin
- B- Network Admin
- C- VPC Admin
- D- Security Operator
- E- Network Operator

Answer:

C, D, E

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:

With the introduction of the Virtual Private Cloud (VPC) consumption model in VCF 9.0 and late 5.x releases, Role-Based Access Control (RBAC) has become more granular to support true multi-tenancy. A VPC is designed to be a self-contained 'container' for a department's or user's networking resources.

To meet the specific requirement where a user can configure aspects of an individual VPC but is restricted from creating new subnets (which involves modifying the underlying network CIDR blocks and IPAM), a combination of specific roles is required.

VPC Admin: This is the primary role for the user within their assigned VPC. It allows the user to manage the overall VPC environment, including high-level settings and monitoring. However, the VPC Admin's power is often limited by the specific quotas and policies set by the Enterprise Admin.

Security Operator: This role allows the user to view security configurations and policies without having the permission to modify the network fabric or create new infrastructure components

like subnets. It provides the 'read-only' visibility into the security posture of the VPC.

Network Operator: Similar to the Security Operator, the Network Operator role provides visibility into the networking state---such as routing tables, segment status, and connectivity---without granting the 'Write' permissions required to provision new subnets or alter the network topology.

Assigning Network Admin (Option B) or Security Admin (Option A) would grant too much privilege, as these roles typically include the ability to create, delete, and modify subnets and firewall policies at a structural level. By combining the VPC Admin role with Operator-level roles, the administrator ensures the user has the necessary context to manage their assigned resources while strictly adhering to the restriction against creating new network subnets.



To Get Premium Files for 3V0-25.25 Visit

<https://www.p2pexams.com/products/3v0-25.25>

For More Free Questions Visit

<https://www.p2pexams.com/vmware/pdf/3v0-25.25>

20%
DISCOUNT

P2P
exams