Download VMware 3V0-25.25 Exam Dumps Free

Shared by Levine on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An administrator has deployed a new VMware Cloud Foundation (VCF) management domain. To be compliant with company policy, backups must be configured to occur anytime a change is made to the NSX configuration. How can the administrator ensure that complete configuration backups are captured every time a change occurs?

Options:

A- Configure an alarm to detect configuration changes and automatically trigger a complete configuration backup.
B- No action is required as by default NSX will automatically perform a complete backup every time a change is made to the configuration.
C- Configure a cron job on the NSX Manager to automatically perform an incremental backup of the NSX configuration every hour.
D- Create a recurring backup schedule and explicitly indicate that backups should be captured anytime the configuration changes.

Answer:

D

Explanation:

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

In VMware Cloud Foundation (VCF), the protection of the NSX Manager configuration is paramount, as it contains the state of the entire software-defined network, including firewall rules, logical switches, and routing topologies. To meet strict compliance requirements for real-time or change-based protection, NSX offers specific automated backup triggers.

Within the NSX Manager UI (under System > Lifecycle > Backup & Restore), an administrator can configure the backup behavior. While a time-based schedule (e.g., daily at 2:00 AM) is common, it does not satisfy the requirement for backups 'anytime a change is made.' To accomplish this, the administrator must enable the 'Backup on Configuration Change' toggle within the backup scheduling configuration.

When this feature is enabled, the NSX Manager monitors its own management database (DS) for write operations. Once a configuration change is detected (such as adding a segment or modifying a DFW rule), the system initiates an automated backup process. This ensures that the backup repository always contains a near-instantaneous reflection of the current network state,

minimizing data loss in the event of a cluster failure.

Option B is incorrect because this feature is not enabled by default; it requires an external SFTP/FTP server to be configured first. Option C (Cron jobs) is an unsupported manual workaround that bypasses the SDDC-native management tools. Option A is redundant as the functionality is built directly into the NSX backup engine. Consequently, the verified method for compliance is to use the native recurring backup schedule with the 'Detect Configuration Change' option enabled.

Question 2

Question Type: MultipleChoice

Which two requirements are part of the registration process for Local Manager (LM) to a Global Manager (GM) in NSX for centralized management of network and security services across different workload domains deployed in separate locations? (Choose two.)

Options:

A- The LM will validate the GM license to perform the GM registration.
B- The external load balancer VIP is used for NSX Managers without requiring node API certificate updates.
C- The LM Cluster VIP / FQDN is provided for GM-LM communication.
D- The IP / FQDN of any of the 3 LM must be used for registration.
E- The GM-Active requests the LM IP / FQDN and admin credentials for registration.

Answer:

C, E

Explanation:

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

NSX Federation is the architectural framework used within VMware Cloud Foundation (VCF) to provide consistent networking and security across multiple sites. The core of this framework is the relationship between the Global Manager (GM) and one or more Local Managers (LMs).

The registration process is the critical first step in establishing this 'parent-child' relationship. According to the 'NSX-T Data Center Administration Guide' and Federation-specific documentation, the registration is initiated from the Active Global Manager.

Initiation and Credentials (Requirement E): The administrator logs into the Global Manager UI and navigates to the 'System > Fabric > Locations' section. To add a new site, the GM-Active requires the IP address or FQDN of the target Local Manager and the Admin credentials. This allows the GM to authenticate with the LM, exchange security certificates, and establish a secure thumbprint-verified connection.

Stable Communication Endpoint (Requirement C): For the ongoing management and synchronization of 'Global Objects' (like Tier-0s or Security Groups), the GM must communicate with the LM cluster as a whole rather than a single individual node. Therefore, the LM Cluster Virtual IP (VIP) or a FQDN pointing to that VIP is provided. Using the VIP ensures that if the specific LM node that initially handled the registration fails, the GM can continue to communicate with the remaining nodes in the LM cluster without administrative intervention.

Option A is incorrect because the Global Manager typically manages the licensing for the federation, not the LM validating the GM. Option B is incorrect as an external load balancer is not a prerequisite for the native GM-LM registration handshake. Option D is incorrect because providing the IP of an individual node (one of the three) does not provide the high availability required for a production Federation environment. Thus, the use of the Cluster VIP and the GM-Active's request for LM credentials are the verified procedural requirements.

Question 3

Question Type: MultipleChoice

An administrator is troubleshooting BGP flapping in a VMware Cloud Foundation (VCF) 9 environment. A Tier-0 Gateway is running in Active/Active mode with two Edge nodes. BFD is enabled on the eBGP sessions to the upstream routers. Each Edge node uses its own uplink IP for BGP. After some network maintenance, one BGP session starts flapping every few minutes. The other BGP sessions stay stable. On the affected Edge node, the command get bfd-sessions shows:

* State: Down

* Diag: Detect Time Expired

Symptoms:

* The upstream router also shows the BFD session as Down with control Detection Time Expired.

* There are no interface errors, no packet loss for normal traffic, and clearing the BFD session temporarily brings it back up - but it flaps again after few minutes.

What is the root cause?

Options:

A- BFD timers are mismatched between Tier-0 Gateway and the upstream routers.
B- The MTU does not match on the end-to-end between Tier-0 Gateway and upstream routers.
C- BFD is configured in echo mode on the upstream routers.
D- The Edge nodes are undersized and are experiencing high contention on CPU and drops BFD packets.

Answer:

B

Explanation:

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

In a VMware Cloud Foundation (VCF) environment, particularly with the high-performance requirements of North-South routing, BGP and BFD (Bidirectional Forwarding Detection) are used in tandem to ensure rapid failure detection. A common but subtle issue in fresh or modified environments is an MTU (Maximum Transmission Unit) mismatch on the physical or virtual uplinks.

When BGP establishes a neighborship, it initially exchanges small keepalive packets. These small packets easily pass through interfaces even if there is an MTU mismatch (e.g., the Edge is set to 9000 bytes but a physical switch in the path is limited to 1500 bytes). However, once the BGP state reaches 'Established,' the routers begin exchanging full routing tables. These BGP Update packets are often large and will be fragmented or dropped if they exceed the MTU of any hop in the path.

The symptom described---where the session is stable for a few minutes (during the initial handshake) and then flaps---is the hallmark of an MTU issue. The 'Detect Time Expired' diagnostic in BFD occurs because the BGP hold timer expires when it fails to receive the large update packets, or the BFD packets themselves are delayed/lost due to the congestion caused by retrying large, failed transmissions. According to VMware NSX troubleshooting documentation, if pings (small packets) succeed but the BGP session fails specifically when traffic load or route counts increase, the MTU should be the first setting verified.

VCF 9.0 and 5.x designs mandate consistent MTU settings (typically 9000 MTU for the overlay and at least 1500+ for the uplinks) across the entire path, including the virtual switch (VDS), the Edge VM vNICs, and the physical ToR switches. A mismatch here prevents the completion of the BGP state machine's full synchronization, leading to the cyclic 'flapping' observed by the administrator.

==========

Question 4

Question Type: MultipleChoice

An administrator created a new Tier-1 Gateway and is attempting to change the connected gateway for a deployed segment to use the new gateway. In the UI, when the administrator clicks the Connected Gateway dropdown, the new Tier-1 gateway is not shown as an available gateway. What would prevent the new Tier-1 gateway from showing in the list of available gateways?

Options:

A- The Tier-1 Gateway is not connected to an NSX Edge Cluster.
B- The Tier-1 Gateway connectivity policy is set to 'None'.
C- The Tier-1 Gateway and NSX Segment are in different transport zones.
D- The Tier-1 Gateway and NSX Segment are connected to different Tier-0 Gateways.

Answer:

C

Explanation:

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

In VMware Cloud Foundation networking, the relationship between segments and gateways is governed by the underlying Transport Zone (TZ) configuration. A Transport Zone defines the potential span of a virtual network---specifically, which hosts and edges can participate in that network.

When an administrator creates an NSX Segment, they must associate it with a specific Transport Zone (either Overlay or VLAN). Similarly, when a Tier-1 Gateway is created, its reach is determined by the Transport Zones available on the Transport Nodes (Edges and ESXi hosts) where it is instantiated. For a Segment to be attached to a Tier-1 Gateway, both objects must reside within the same Transport Zone.

If the Segment was created in 'Overlay-TZ-01' but the new Tier-1 Gateway is only associated with 'Overlay-TZ-02' (or if one is in a VLAN TZ and the other in an Overlay TZ), the NSX Manager UI will filter out the incompatible gateway to prevent an invalid configuration. The logical switch (Segment) cannot bind to a gateway if they do not share a common broadcast or encapsulation domain defined by the Transport Zone.

Option A is incorrect because a Tier-1 Gateway does not strictly require an Edge Cluster unless it

is providing stateful services (like NAT, LB, or Firewall). It can exist purely as a distributed component on the hypervisors. Option B (Connectivity Policy) determines if the T1 advertises routes to the T0, but it doesn't prevent a segment from connecting to it. Option D is also incorrect, as a Tier-1 Gateway can be moved between Tier-0s, or even exist without a Tier-0 connection initially. Therefore, the Transport Zone mismatch is the fundamental architectural barrier preventing the gateway from appearing in the selection list.

==========

Question 5

Question Type: MultipleChoice

An administrator is configuring Border Gateway Protocol (BGP) routing on a Tier-0 Gateway to optimize north---south traffic flow between the NSX environment and multiple upstream physical routers. The environment includes two external connections that advertise overlapping routes to the same destination networks. To ensure predictable and efficient routing behavior, the administrator decides to manipulate specific BGP attributes on outbound advertisements and inbound route updates. What are two valid BGP Attributes that can be used to influence the route path traffic will take? (Choose two.)

Options:

A- BFD
B- Cost
C- AS-Path Prepend
D- MED

Answer:

C, D

Explanation:

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

In a VMware Cloud Foundation (VCF) architecture, the Tier-0 Gateway is the primary point of integration between the virtualized network and the physical world. When dealing with multiple upstream routers (multi-homing), administrators must influence the BGP path selection process to ensure traffic follows the desired path and avoids suboptimal routing or asymmetric flows.

AS-Path Prepend is a common technique used to influence inbound traffic (traffic coming from the physical network into the NSX environment). By repeating its own Autonomous System (AS) number multiple times in the BGP advertisement, the Tier-0 Gateway makes a specific path look 'longer' and therefore less desirable to the upstream physical routers. Since BGP prefers the shortest AS-Path, the routers will favor the alternate link that does not have the prepended AS numbers. This is a critical tool in VCF designs to ensure that a primary link is utilized unless a failure occurs.

MED (Multi-Exit Discriminator) is an attribute that suggests to an adjacent external AS which path to take among multiple entry points to the same AS. Like AS-Path Prepend, it influences inbound traffic. A lower MED value is preferred over a higher one. In a VCF environment with multiple Edge Nodes or multiple Tier-0 uplinks, setting different MED values allows the administrator to prioritize specific entry points for traffic entering the SDDC.

BFD (Bidirectional Forwarding Detection) is not a BGP attribute; it is a detection protocol used to provide fast failure detection of the link between BGP neighbors. While it triggers faster convergence, it does not influence path selection based on attributes. Cost is an OSPF attribute, not a native BGP attribute. Therefore, in the context of NSX Tier-0 BGP configuration, AS-Path Prepend and MED are the verified methods for path manipulation.

==========

Question 6

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 7

Question Type: MultipleChoice

An administrator is troubleshooting east---west network performance between several virtual machines connected to the same logical segment. The administrator inspects the internal forwarding tables used by ESXi and notices that different tables exist for MAC and IP mapping. Which table on an ESXi host is used to determine the location of a particular workload for frame forwarding?

Options:

A- ARP Table
B- FIP Table
C- TEP Table
D- MAC Table

Answer:

D

Explanation:

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

In the context of VMware Cloud Foundation (VCF) networking, understanding how an ESXi host (acting as a Transport Node) handles East-West traffic is fundamental. East-West traffic refers to communication between workloads within the same data center, often on the same logical segment.

When a Virtual Machine sends a frame to another VM on the same logical segment, the ESXi host's virtual switch must determine the 'location' of the destination MAC address to perform frame forwarding. The MAC Table (also known as the Forwarding Table or L2 Table) is the primary structure used for this decision. For each logical segment, the host maintains a MAC table that maps the MAC addresses of virtual machines to their specific 'locations.'

If the destination VM is residing on the same host, the MAC table points the frame toward a specific internal port (vUUID) associated with that VM's vNIC. If the destination VM is on a different host (in an overlay environment), the MAC table entry for that remote MAC address will point to the Tunnel End Point (TEP) IP of the remote ESXi host. While the TEP table (Option C) contains the list of known Tunnel Endpoints and the ARP table (Option A) maps IP addresses to MAC addresses, neither is the primary table used for the final frame forwarding decision.

The MAC Table is the authoritative source for Layer 2 forwarding. In an NSX-managed VCF environment, these tables are dynamically populated and synchronized via the Local Control Plane (LCP), which receives updates from the Central Control Plane. This ensures that even as VMs move via vMotion, the MAC table remains updated across all transport nodes, allowing for seamless East-West connectivity without the need for traditional MAC learning (flooding) in the physical fabric.

Question 8

Question Type: MultipleChoice

An administrator is tasked to configure NSX Federation between separate VMware Cloud Foundation (VCF) Fleets. Which requirement must all sites meet before being added to a Global Manager (GM) for NSX Federation?

Options:

A- All Sites must use the same VTEP VLAN and IP pools.
B- All sites must use identical Tier-0 gateway BGP autonomous system numbers.
C- All sites must be managed by the same VCF instance.
D- All sites must have the same NSX version and build.

Answer:

D

Explanation:

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

NSX Federation, a core component of large-scale VCF deployments across multiple sites or 'fleets,' introduces a hierarchical management model where a Global Manager (GM) orchestrates security policies and networking objects across multiple Local Managers (LMs).

To ensure stability and compatibility in the communication between the Global Manager and the Local Managers, VMware documentation specifies strict version parity requirements. When onboarding a site into a Federation, the Local Manager at that site must be running the same NSX version and build as the other sites in the Federation and must be compatible with the Global Manager's version. Discrepancies in versions can lead to synchronization failures, as the API schemas and internal database structures for Global Objects (like Global Segments or Groups) may differ between builds.

While Federation allows for geographic and administrative separation, the underlying software-defined networking stack must be synchronized. Option A is incorrect; in fact, VTEP/TEP VLANs and IP pools should be unique to each site to avoid IP conflicts in the transport network, though they must have Layer 3 reachability to one another. Option B is incorrect; unique BGP AS numbers are often preferred for multi-site routing to prevent loops. Option C is also incorrect, as Federation is specifically designed to link different VCF instances (sites) together into a single manageable entity.

In a VCF 5.x or 9.0 context, the SDDC Manager helps maintain this requirement by ensuring that the 'Bill of Materials' (BOM) is consistent across sites intended for Federation. Before the GM can successfully register and 'push' configuration to an LM, the handshake process validates the build version to prevent the corruption of the global intended state.

Question 9

Question Type: MultipleChoice

An administrator must provide North/South connectivity for a VPC. The fabric exposes a distributed external VLAN across all ESX hosts. But, the only BGP peer to the core is on a VLAN only accessible on the Edge Cluster. Which design is required?

Options:

A- Use a VPC Tier-0 Gateway in active/active mode with distributed eBGP peering.
B- Distributed Transit Gateway with an EVPN route reflector on the transport nodes.
C- Centralized Transit Gateway on the Edge Cluster.
D- Deploy a Provider Tier-1 with BGP and connect the VPC Transit Gateway via route leaking.

Answer:

C

Explanation:

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

In a VMware Cloud Foundation (VCF) environment utilizing the Virtual Private Cloud (VPC) model, North/South connectivity is managed by the Transit Gateway (TGW). The TGW acts as the bridge between the VPC-internal networks and the provider-level physical network.

The scenario presents a specific constraint: while an external VLAN exists across all hosts, the actual BGP peering point (the interface to the physical core routers) is restricted to the NSX Edge Cluster. In NSX terminology, when a gateway or service must be anchored to specific Edge Nodes to access physical network services---such as BGP peering, NAT, or stateful firewalls---it must be configured as a Centralized component.

A Centralized Transit Gateway (Option C) is instantiated on the Edge nodes. This allows the TGW to participate in the BGP session with the core routers on the VLAN that is only accessible to those Edges. The TGW then handles the routing for the VPC's internal segments. Traffic from the ESXi transport nodes (East-West) travels via the Geneve overlay to the Edge nodes, where it is then routed North-South by the Centralized TGW using the physical BGP peer.

Option A is incorrect because 'distributed eBGP peering' would require every ESXi host to have peering capabilities, which contradicts the constraint. Option B involves EVPN, which is a significantly more complex and different architecture than what is required for standard VPC

North/South access. Option D is an unnecessarily complex routing design that is not the standard VCF/VPC implementation pattern. Thus, the use of a Centralized Transit Gateway on the Edge cluster is the verified design requirement to bridge the gap between the overlay VPC and the localized BGP peering point.

P2P exams