prep BOLT

Get Free VMware 3V0-25.25 Dumps PDF Questions

Why risk failure? Download updated VMware Cloud Foundation 9.0 Networking exam PDF questions today. Practice with real 3V0-25.25 dumps and verified answers designed to help you ace your certification quickly using PrepBolt 3V0-25.25 exam pdf questions and answers.

Thank you for Downloading 3V0-25.25 exam PDF Demo

https://prepbolt.com/3V0-25.25.html

QUESTIONS & ANSWERS

DEMO VERSION

(LIMITED CONTENT)

Question 1

Question Type: MultipleChoice

An administrator has a vSphere 8 Update 1a with NSX 4.1.0.2 environment. What option can the administrator use to converge this vSphere with NSX environment into a VMware Cloud Foundation (VCF) Workload Domain?

Options:

A- Use the VCF installer to automatically converge the vSphere with NSX environment into a new VCF Workload Domain.

B- Upgrade NSX to version 9 into the vSphere 8 environment and use the VCF installer to converge the vSphere 8 with NSX environment into a new VCF Workload Domain.

C- Upgrade the environment version and use the VCF installer to converge the vSphere environment into a new VCF Workload Domain.

D- Upgrade the environment and use VCF Operations to converge the vSphere environment into a new VCF Workload Domain.

Answer:

A

Explanation:

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

The process of transforming an existing, 'brownfield' environment into a VCF-managed infrastructure is known as Convergence. In VCF 5.x and the advancements found in VCF 9.0, VMware provides the VCF Import Tool (often bundled or utilized alongside the VCF Installer/Cloud Builder) specifically for this purpose.

An environment running vSphere 8 Update 1a and NSX 4.1.0.2 is within the supported compatibility matrix for VCF 5.x convergence. The most direct and verified method (Option A) is to use the VCF Installer to 'ingest' the existing vCenter and NSX Manager. During this process, the installer validates the current configuration, ensures the hosts are compatible, and then brings them under the management of a newly deployed SDDC Manager.

One of the significant advantages of this approach is that it avoids the need for a 'rip and replace' of the existing networking. The VCF Installer identifies the existing NSX Manager and the logical networking constructs. Once the convergence is successful, the environment is treated as a standard VCF Workload Domain.

Options B and C are incorrect because VCF's design principle is to perform the convergence at a known stable and compatible version before using the SDDC Manager's Lifecycle Management (LCM) to perform upgrades. Manually upgrading to version 9 prior to convergence can introduce configuration drifts that the VCF Installer may not be able to reconcile. Option D is incorrect as VCF Operations (formerly vRealize Operations) is a monitoring and optimization tool; it does not have the administrative capability to perform the structural convergence of the SDDC stack. Therefore, the automated convergence via the VCF Installer is the correct architectural path.

===========

Question 2

Question Type: MultipleChoice

An administrator encountered a failure with one of the NSX Managers in a VCF Fleet. The administrator has successfully re-deployed an NSX Manager from SFTP backups. However, after replacing the failed manager node, the new node joins successfully, but the cluster status remains "Degraded".

* The get cluster status command on the leader still shows the old UUID with state "REMOVED".

What is the command to resolve the issue?

Options:

A- detach node <new-uuid>

B- delete node <old-uuid>

C- detach node <old-uuid> then delete node <old-uuid>

D- detach node <old-uuid>

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, the NSX Management Cluster consists of three nodes to ensure high availability and quorum. When a single node fails and is subsequently replaced---either through a manual deployment or an orchestrated recovery via SDDC Manager---the internal database (Corfu) and the cluster manager must be updated to reflect the current members of the cluster.

When a node is lost or manually deleted from vCenter without being properly decommissioned through the NSX API or CLI, the remaining 'Leader' node retains the metadata and the UUID of that missing member. Even after a new node joins the cluster and synchronizes data, the cluster state often remains in a 'Degraded' status because the control plane still expects a response from the original, failed UUID.

According to NSX troubleshooting and recovery guides, the specific command to purge a stale or defunct member from the cluster configuration is detach node <UUID>. This command must be executed from the CLI of the current Cluster Leader. By running detach node <old-uuid>, the administrator instructs the cluster manager to permanently remove the record of the failed node from the management plane's membership list.

Option B and C are incorrect because 'delete node' is not the primary CLI command used for cluster membership cleanup; 'detach' is the specific primitive required to break the logical association. Option A would remove the healthy new node, worsening the situation. Once the stale UUID is detached, the cluster status should transition from 'Degraded' to 'Stable' as it no longer tries to communicate with the non-existent entity. This process is essential in VCF operations to maintain a healthy 'green' status in both the NSX Manager and the SDDC Manager dashboard.

===========

Question 3

Question Type: MultipleChoice

In an NSX environment, an administrator is observing low throughput and intermittent congestion between the Tier-0 Gateway and the upstream physical routers. The environment was designed for high availability and load balancing, using two Edge Nodes deployed in Active/Active mode. The administrator enables ECMP on the Tier-0 gateway, but the issues persist. Which action would address low throughput and congestion?

Options:

A- Convert Tier-1 gateways to be edgeless.
B- Disable NAT on the Tier-0 gateway.
C- Add an additional vNIC to the NSX Edge node.
D- Deploy additional Edge nodes.

Answer:

D

Explanation:

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

When a VMware Cloud Foundation (VCF) environment experiences North-South congestion at the Tier-0 Gateway, it typically indicates that the processing capacity of the existing NSX Edge Nodes has been reached. In an Active/Active configuration, the Tier-0 gateway utilizes Equal Cost Multi-Pathing (ECMP) to distribute traffic across all available Edge nodes in the cluster.

If a two-node Edge cluster is saturated despite ECMP being enabled, the standard 'Scale-Out' procedure is to deploy additional Edge nodes (Option D). NSX supports up to 8 Edge nodes in a single cluster for a Tier-0 gateway. By adding more nodes, the administrator increases the total number of CPU cores dedicated to the DPDK (Data Plane Development Kit) packet processing engine. Each additional node provides more 'bandwidth lanes' for the ECMP hash to utilize, effectively multiplying the aggregate throughput capability of the North-South exit point.

Option A is incorrect because 'edgeless' Tier-1 gateways (Distributed Routers only) improve East-West performance by keeping traffic on the ESXi hosts, but they do not help with North-South traffic that must eventually hit a Tier-0 Service Router on an Edge. Option B (Disabling NAT) might reduce CPU overhead slightly, but it doesn't solve a fundamental capacity bottleneck and is often not an option due to architectural requirements. Option C (Adding a vNIC) does not increase the underlying compute/DPDK processing power of the Edge VM and can sometimes complicate the load-balancing hash.

In VCF operations, this expansion is handled via the SDDC Manager, which can automate the addition of new Edge nodes to an existing cluster, ensuring they are configured symmetrically with the correct uplink profiles and BGP peering sessions. This horizontal scaling is the verified method for resolving congestion in high-demand VCF networking environments.

Question 4

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 5

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 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.

===========