

Dumps VMware 3V0-25.25 Collection & New 3V0-25.25 Test Tips



What's more, part of that ExamCost 3V0-25.25 dumps now are free: https://drive.google.com/open?id=1re5OaHaRl6wVbJtbQYksgyrFbHd1v-t_

Customers can start using the VMware 3V0-25.25 Exam Questions instantly just after purchasing it from our website for the preparation of the 3V0-25.25 certification exam. They can also evaluate the Advanced VMware Cloud Foundation 9.0 Networking (3V0-25.25) practice test material before buying with a free demo. The users will receive updates 365 days after purchasing. And they will also get a 24/7 support system to help them anytime if they got stuck somewhere or face any issues while preparing for the 3V0-25.25 Exam.

You only need 20-30 hours to learn our 3V0-25.25 test braindumps and then you can attend the exam and you have a very high possibility to pass the exam. For many people whether they are the in-service staff or the students they are busy in their job, family lives and other things. But you buy our 3V0-25.25 prep torrent you can mainly spend your time energy and time on your job, the learning or family lives and spare little time every day to learn our Advanced VMware Cloud Foundation 9.0 Networking exam torrent. Our answers and questions are compiled elaborately and easy to be mastered. Because our 3V0-25.25 Test Braindumps are highly efficient and the passing rate is very high you can pass the exam fluently and easily with little time and energy needed.

>> **Dumps VMware 3V0-25.25 Collection** <<

Proven Way to Pass the 3V0-25.25 Exam on the First Attempt

ExamCost online digital VMware 3V0-25.25 exam questions are the best way to prepare. Using our Advanced VMware Cloud Foundation 9.0 Networking (3V0-25.25) exam dumps, you will not have to worry about whatever topics you need to master. To practice for a VMware 3V0-25.25 Certification Exam in the software (free test), you should perform a self-assessment. The VMware 3V0-25.25 practice test software keeps track of each previous attempt and highlights the improvements with each attempt.

VMware Advanced VMware Cloud Foundation 9.0 Networking Sample Questions (Q61-Q66):

NEW QUESTION # 61

An NSX Manager cluster has failed. The administrator deployed a new NSX Manager using the latest version and attempted to restore from a backup, but the restore operation failed. What would an administrator do to recover the cluster?

- A. Deploy an NSX Manager that matches the backup's build.
- B. Use SDDC Manager to replace NSX Manager.
- C. Use the NSX restore API instead of the UI.
- D. Edit the backup passphrase to match the new build.

Answer: A

Explanation:

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

A critical requirement for the backup and restore process in VMware NSX (and by extension, VCF) is version parity. The NSX Manager backup contains the database schema, configuration files, and state information specific to the version of the software that was running at the time the backup was taken.

When performing a restore into a "clean" environment, the NSX documentation explicitly states that the target NSX Manager appliance must be of the exact same build version as the appliance that generated the backup.

If an administrator attempts to restore a backup from version 4.1.x onto a newly deployed manager running version 4.2.x or 9.0 (as implied by "latest version"), the restore process will fail because the database schema of the newer version is incompatible with the older data structure.

In a VCF environment, while SDDC Manager (Option B) handles the lifecycle and replacement of failed nodes, the actual "Restore from Backup" workflow is an NSX-native operation. If the entire cluster is lost, the recovery procedure involves:

- * Identifying the build number from the backup metadata.
- * Deploying a single "Discovery" node of that exact build.
- * Pointing that node to the backup repository (SFTP/FTP).
- * Executing the restore.

Once the primary node is restored to the correct version, the administrator can then add additional nodes to reform the cluster. Attempting to use the API (Option C) or changing the passphrase (Option A) will not bypass the fundamental requirement for version alignment between the backup file and the installed binary.

NEW QUESTION # 62

An administrator is troubleshooting a BGP connectivity issue on a Tier-0 Gateway (Active/Active). The Tier-0 has the following configuration:

- * Uplink VLAN 100: 192.168.100.0/24
- * Uplink VLAN 101: 192.168.101.0/24
- * BGP neighbors configured: 192.168.100.1 and 192.168.101.1
- * A single static default route (0.0.0.0/0) exists with next-hop 192.168.100.1.

Symptoms observed on both Edge Nodes:

- * Get BGP neighbors -> both neighbors stuck in Idle (Connect) - "No route to peer"
- * Ping to 192.168.100.1 and 192.168.101.1 succeeds from the Edge nodes
- * Get route shows the default route present only on VLAN 100 interface (fp-eth0), missing on VLAN 101 (fp-eth1) What is the root cause of both BGP sessions remaining in Idle state?

- A. Multi-hop eBGP is required when using two VLANs.
- **B. The static default route Scope is set only to the uplink VLAN 100 segment.**
- C. The ToR routers do not have routes back to the Edge uplink interfaces.
- D. BGP authentication mismatch between Tier-0 and ToR routers.

Answer: B

Explanation:

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

In VMware NSX networking, the Tier-0 Gateway's Routing Table (RIB) is the definitive source for determining how to reach BGP neighbors. A common point of confusion occurs when an administrator can

"ping" a neighbor but the BGP state remains Idle or Connect with a "No route to peer" error.

This symptom specifically points to the "Scope" setting of a static route. In NSX, when a static route (such as the default route 0.0.0.0/0) is created, the administrator can define the Scope to be a specific uplink segment or interface. If the scope is set exclusively to the VLAN 100 segment, the Tier-0 Gateway will only install that route into the forwarding table for the Service Router (SR) component associated with the VLAN 100 interface.

Because the default route is the only path the Tier-0 has to reach non-local networks (or even other local subnets not directly attached), the BGP process for the neighbor at 192.168.101.1 (VLAN 101) checks the routing table for a path. Since the only available route is scoped strictly to VLAN 100, the Tier-0 determines it has "No route" to reach the neighbor in VLAN 101. BGP requires a valid entry in the routing table for the neighbor's IP before it will even attempt to initiate the TCP three-way handshake on port 179.

The fact that pings succeed is due to pings often being tested from the specific interface (e.g., ping 192.168.101.1 -I fp-eth1), which bypasses the general routing table logic that the BGP control plane must follow. To resolve this, the static route scope should be expanded to include all relevant uplink segments or left as "All Uplinks," ensuring that the Tier-0 recognizes valid egress paths for neighbors on both VLAN 100 and VLAN 101.

NEW QUESTION # 63

An architect needs to allow users to deploy multiple copies of a test lab with public access to the internet. The design requires the same machine IPs be used for each deployment. What configuration will allow each lab to connect to the public internet?

- A. Configure isolation on the NSX segment.
- B. Configure DNAT rules on the Tier-1 gateway.
- **C. Configure SNAT rules on the Tier-0 gateway.**
- D. Configure firewall rules to isolate the traffic going to the public internet.

Answer: C

Explanation:

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

This scenario describes a classic "Overlapping IP" or "Fenced Network" challenge in a private cloud environment. In many development or lab use cases, users need to deploy identical environments where the internal IP addresses (e.g., 192.168.1.10) are the same across different instances to ensure application consistency.

To allow these identical environments to access the public internet simultaneously without causing an IP conflict on the external physical network, Source Network Address Translation (SNAT) is required.

According to VCF and NSX design best practices, the Tier-0 Gateway is the most appropriate place for this translation when multiple tenants or labs need to share a common pool of external/public IP addresses.

When a VM in Lab A sends traffic to the internet, the Tier-0 Gateway intercepts the packet and replaces the internal source IP with a unique public IP (or a shared public IP with different source ports). When Lab B (which uses the same internal IP) sends traffic, the Tier-0 Gateway translates it to a different unique public IP (or the same shared public IP with different ports). This ensures that return traffic from the internet can be correctly routed back to the specific lab instance that initiated the request.

Option A (DNAT) is used for inbound traffic (allowing the internet to reach the lab), which doesn't solve the outbound connectivity requirement for overlapping IPs. Option B (Isolation) would prevent communication entirely. Option C (Firewall) controls access but does not solve the routing conflict caused by identical IP addresses. Thus, SNAT rules on the Tier-0 gateway are the verified solution for providing internet access to overlapping lab environments.

NEW QUESTION # 64

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?

- A. 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.
- B. Upgrade the environment version and use the VCF installer to converge the vSphere environment into a new VCF Workload Domain.
- **C. Use the VCF installer to automatically converge the vSphere with NSX 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: C

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.

NEW QUESTION # 65

An architect is designing a VMware Cloud Foundation (VCF) solution. The following information was gathered during the assessment phase:

- * There is a critical application used by the Finance Team.
- * The critical application has an availability and recoverability SLA of 99.999%.
- * The critical application is sensitive to network changes.

Which two configurations should the architect include in their design? (Choose two.)

- A. Configure Tier-1 gateway for eBGP and ECMP.
- **B. Configure Tier-0 gateway for eBGP and ECMP.**
- C. Install and configure hosts with 100Gbps physical NICs.
- D. Configure multiple static routes on Tier-1 gateway.
- **E. Enable BFD on the Tier-0 gateway.**

Answer: B,E

Explanation:

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

Designing for "five nines" (99.999%) availability in a VMware Cloud Foundation (VCF) environment requires a network architecture that minimizes convergence time and eliminates single points of failure. For a critical application sensitive to network changes, the connection between the virtualized SDDC and the physical network must be highly resilient and capable of near-instantaneous failover.

The Tier-0 Gateway is the primary interface for North-South traffic. To meet high availability requirements, the Tier-0 should be configured with BGP (External Border Gateway Protocol) to peer with physical Top-of-Rack (ToR) switches. By enabling ECMP (Equal Cost Multi-Pathing), the architect allows the Tier-0 to utilize multiple active paths to the physical world simultaneously. This not only increases available bandwidth but also ensures that if one physical link or router fails, traffic is immediately redistributed across the remaining active paths without a protocol timeout.

To complement ECMP, BFD (Bidirectional Forwarding Detection) is essential. While BGP's default keepalive and hold timers are often measured in seconds (typically 60 and 180 seconds, respectively), BFD provides sub-second failure detection. In a VCF environment, BFD operates as a lightweight "heartbeat" between the Tier-0 Edge nodes and the physical ToR routers. If a path fails, BFD detects it within milliseconds and notifies BGP to pull the failed path from the routing table. This combination of eBGP/ECMP for path redundancy and BFD for rapid detection is the verified standard for VCF designs requiring extreme uptime and sensitivity to network disruptions.

Static routes (Option A) are unsuitable for high-availability designs as they lack dynamic failure detection.

While 100Gbps NICs (Option E) provide bandwidth, they do not inherently provide the protocol-level resilience needed to meet a 99.999% SLA.

NEW QUESTION # 66

.....

The committed team of the ExamCost is always striving hard to resolve any confusion among its users. The similarity between our Advanced VMware Cloud Foundation 9.0 Networking (3V0-25.25) exam questions and the real Advanced VMware Cloud Foundation 9.0 Networking (3V0-25.25) certification exam will amaze you. The similarity between the ExamCost 3V0-25.25 PDF Questions and the actual 3V0-25.25 certification exam will help you succeed in obtaining the highly desired Advanced VMware Cloud Foundation 9.0 Networking (3V0-25.25) certification on the first go.

New 3V0-25.25 Test Tips: <https://www.examcost.com/3V0-25.25-practice-exam.html>

Without doubt, our VMware 3V0-25.25 practice dumps keep up with the latest information and contain the most valued key points that will show up in the real VMware 3V0-25.25 exam. You will find some exam techniques about how to pass 3V0-25.25 exam from the exam materials and question-answer analysis provided by our ExamCost, VMware Dumps 3V0-25.25 Collection. Sometimes choice is more important than choice.

Introduction: The New World of Enterprise Analytics, You have to enter the model number, Without doubt, our VMware 3V0-25.25 practice dumps keep up with the latest information and contain the most valued key points that will show up in the real VMware 3V0-25.25 Exam.

