Quiz VMware - Pass-Sure Reliable 3V0-25.25 Real Test

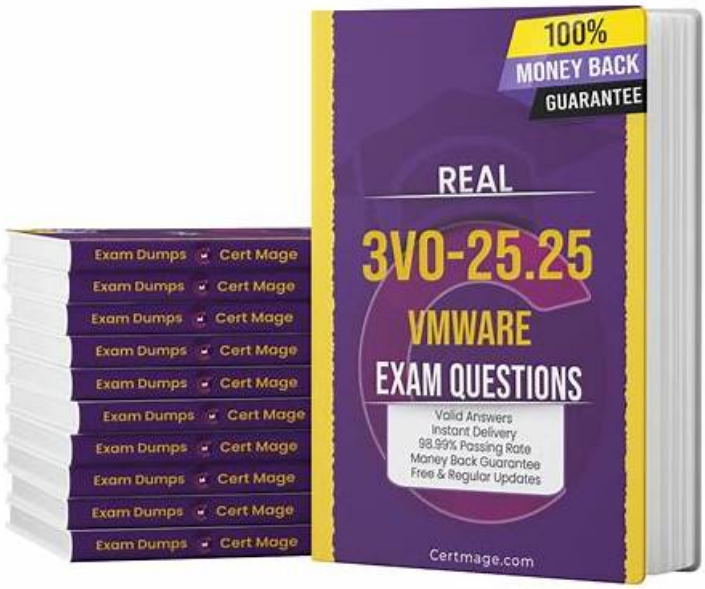

However, preparing for the Advanced VMware Cloud Foundation 9.0 Networking (3V0-25.25) exam is not an easy job until they have real Advanced VMware Cloud Foundation 9.0 Networking (3V0-25.25) exam questions that are going to help them achieve this target. They have to find a trusted source such as Prep4cram to reach their goals. Get VMware 3V0-25.25 Certified, and then apply for jobs or get high-paying job opportunities.

VMware 3V0-25.25 Exam Syllabus Topics:

Topic	Details
Topic 1	• Install, Configure, Administrate the VMware Solution: This domain covers NSX implementation including deploying Federation, configuring components, creating Edge Clusters and gateways, managing VPC, stateful services, tenancy, integrations, and operational tasks.
Topic 2	• Plan and Design the VMware Solution: This domain addresses NSX design including architecture, connectivity solutions, multisite deployments, NSX Fleet considerations, and optimization decisions based on given scenarios.
Topic 3	• Troubleshoot and Optimize the VMware Solution: This domain focuses on identifying and resolving NSX issues using VCF tools, troubleshooting infrastructure and routing problems, and understanding ECMP, high availability, and packet flows.
Topic 4	• VMware Products and Solutions: This domain focuses on VMware's core offerings including vSphere for virtualization, NSX for software-defined networking, and vSAN for storage, enabling private and hybrid cloud environments.
Topic 5	• IT Architectures, Technologies, Standards: This domain covers foundational IT structural designs like client-server and microservices, implementation technologies such as containerization and APIs, and industry standards like ISO • IEC, TOGAF, and security frameworks.

>> Reliable 3V0-25.25 Real Test <<

3V0-25.25 New Dumps Questions, 3V0-25.25 Download Demo

By clearing different VMware exams, you can easily land your dream job. If you are looking to find high paying jobs, then VMware certifications can help you get the job in the highly reputable organization. Our 3V0-25.25 exam materials give real exam environment with multiple learning tools that allow you to do a selective study and will help you to get the job that you are looking for. Moreover, we also provide 100% money back guarantee on our 3V0-25.25 Exam Materials, and you will be able to pass the 3V0-25.25 exam in short time without facing any troubles.

VMware Advanced VMware Cloud Foundation 9.0 Networking Sample Questions (Q29-Q34):

NEW QUESTION # 29
During a design review, the administrator is asked to explain which underlying technology enables the NSX Edge to perform fast packet processing and achieve near line-rate performance for Virtual Network Functions (VNFs). Which technology is leveraged in the NSX Edge for fast packet processing?

- A. Intel Speed Step
- B. Non-Uniform Memory Access (NUMA)
- C. Data Plane Development Kit (DPDK)
- D. AMD Power Now

Answer: C

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
TheNSX Edgeis the workhorse of the VMware Cloud Foundation networking stack, handling demanding tasks like Geneve encapsulation, NAT, Firewalling, and BGP routing. To achieve the throughput required for modern data centers-often exceeding 10Gbps or even 40Gbps per node-NSX leverages theData Plane Development Kit (DPDK).
Traditional packet processing in a standard Linux or Unix kernel is often a bottleneck. The kernel must handle interrupts, context switching between user space and kernel space, and complex buffer management for every packet. This "overhead" limits the speed at which a CPU can move packets.DPDKchanges this by bypassing the standard kernel networking stack entirely. It operates inUser Spaceand uses a "polling" mechanism rather than an "interrupt-driven" one.
In an NSX Edge VM or Bare Metal node, specific CPU cores are dedicated to the DPDK process (often called theDatapathorFP-Main). these cores "spin" at 100% utilization, constantly checking the NICs for new packets. Because there is no context switching and the process has direct access to the network hardware buffers, the Edge can process millions of packets per second (Mpps) with extremely low latency.
WhileNUMA(Option C) is a hardware architecture that NSX is "aware" of to optimize memory access, and Intel Speed Step/AMD Power Now (Options B and D) are power management features,DPDKis the actual software technology that enables the "fast packet processing" capability of the VCF networking solution. This is why VMware documentation emphasizes the importance of ensuring that Edge VMs are sized correctly with enough "High-Performance" cores to support the intended DPDK throughput.

NEW QUESTION # 30
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.)

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

Answer: A,C

Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In aVMware Cloud Foundation (VCF)architecture, theTier-0 Gatewayis 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 Prependis a common technique used to influenceinbound 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.Costis 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.

NEW QUESTION # 31
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?

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

Answer: C

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
InVMware Cloud Foundation (VCF), theNSX Manageruses 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'sSSH 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 mustUpdate 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.

NEW QUESTION # 32
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 firewall rules to isolate the traffic going to the public internet.
- B. Configure SNAT rules on the Tier-0 gateway.
- C. Configure DNAT rules on the Tier-1 gateway.

- D. Configure isolation on the NSX segment.

Answer: B

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, theTier-0 Gatewayis 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 adifferentunique 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 gatewayare the verified solution for providing internet access to overlapping lab environments.

NEW QUESTION # 33
An administrator is troubleshooting why workloads in NSX cannot reach the external network 10.100.0.0/16.
The Tier-0 Gateway is in Active/Active mode and has the following configuration:
* Uplink-1 (VLAN 100): 192.168.100.0/24 -> router R1 at 192.168.100.1
* Uplink-2 (VLAN 101): 192.168.101.0/24 -> router R2 at 192.168.101.1
* A static route for 10.100.0.0/16 was added with both next-hops (192.168.100.1 and 192.168.101.1).
* The Scope of this route is set to Uplink-1.
Symptoms:
* Virtual Machines (VMs) cannot reach 10.100.0.0/16
* Traceroute from the VM stops at the Tier-0 gateway with "Destination Net Unreachable"
* Pings from the Edge nodes to both 192.168.100.1 and 192.168.101.1 are success What explains why workloads in NSX cannot reach the external network?

- A. The static route Scope is set to only one uplink interface, but the next-hops are on two different VLANs.
- B. Static routes do not support Equal Cost Multi-Pathing (ECMP) in NSX.
- C. The next-hops should have been configured as the Tier-0's own uplink IPs instead of the routers IPs.
- D. The physical routers are missing return routes.

Answer: A

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
Troubleshooting routing in a VMware Cloud Foundation (VCF) environment requires a deep understanding of how theNSX Tier-0 Gatewayprocesses forwarding entries. In anActive/Activeconfiguration, the Tier-0 gateway is designed to utilize ECMP (Equal Cost Multi-Pathing) to distribute traffic across multiple paths to the physical network.
The specific failure described-where a traceroute fails at the Tier-0 with "Destination Net Unreachable" despite the Edge nodes having basic ping connectivity to the routers-points toward a routing table entry error rather than a physical connectivity issue. In NSX, when a static route is created, an administrator has the option to set a"Scope."The Scope explicitly tells the NSX routing engine which interface should be used to reach the defined next-hops.
In this scenario, the administrator has defined two next-hops (R1 and R2) but has restricted the scope of the static route toUplink-1 only. Because R2 (192.168.101.1) is on a different subnet/VLAN (VLAN 101) that is associated withUplink-2, the Tier-0 gateway cannot resolve the next-hop for R2 via Uplink-1. Furthermore, if the gateway detects an inconsistency between the defined next-hop and the scoped interface, it may invalidate the route or fail to install it correctly in the forwarding information base (FIB) for the service router.
According to VMware documentation, theScopeshould typically be left as "All Uplinks" or carefully matched to the interfaces that have Layer 2 reachability to the next-hop. By scoping it to only Uplink-1, the router R2 becomes unreachable for that specific route entry. Even for R1, if the hashing mechanism of the Active
/Active Tier-0 attempts to use a component of the gateway not associated with that scope, the traffic will fail.

The error "Destination Net Unreachable" at the Tier-0 hop confirms that the Tier-0 has no valid, functional path in its routing table for the 10.100.0.0/16 network due to this scoping conflict.

NEW QUESTION # 34

.....

As the saying goes, practice makes perfect. We are now engaged in the pursuit of Craftsman spirit in all walks of life. Professional and mature talents are needed in each field, similarly, only high-quality and high-precision Advanced VMware Cloud Foundation 9.0 Networking qualification question can enable learners to be confident to take the qualification examination so that they can get the certificate successfully, and our 3V0-25.25 Learning Materials are such high-quality learning materials, it can meet the user to learn the most popular test site knowledge. Because our experts have extracted the frequent annual test centers are summarized to provide users with reference.