PassTestking

- 24/7 customer support, Secure shopping site
- Free One year updates to match real exam scenarios
- If you failed your exam after buying our products we will refund the full amount back to you.

Download Demo

- Online Tool, Convenient, easy to study.
- Instant Online Access
- Supports All Web Browsers
- Practice Online Anytime
- Test History and Performance Review
- Supports Windows / Mac / Android / iOS, etc.

- Installable Software Application
- Simulates Real Exam Environment
- Builds Exam Confidence
- Supports MS Operating System
- Two Modes For Practice
- Practice Offline Anytime

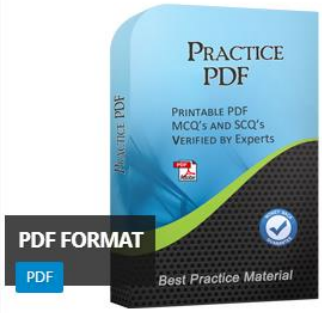

- Printable PDF Format
- Prepared by IT Experts
- Instant Access to Download
- Study Anywhere, Anytime
- 365 Days Free Updates
- Free PDF Demo Available

48923+ Happy Clients

48923+ Shares

97846+ Downloads

9999+ Years in Business

http://www.passtestking.com/

Newest Exam Practice Questions Dumps added from PassTestking

Exam : **3V0-25.25**

Title : Advanced VMware Cloud Foundation 9.0 Networking

Vendor : VMware

Version : DEMO

NO.1 An administrator is investigating packet loss reported by workloads connected to VLAN segments in an NSX environment. Initial checks confirm:
* All VMs are powered on
* VLAN segment IDs are consistent across transport nodes
* Physical switch configurations are correct.
Which two NSX tools can be used to troubleshoot packet loss on VLAN Segments? (Choose two.)
A. Flow Monitoring
B. Traceflow
C. Packet Capture
D. Activity Monitoring
E. Live Flow
Answer: B C
Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In a VMware Cloud Foundation (VCF) environment, troubleshooting packet loss requires tools that can provide visibility into both the logical and physical paths of a packet. When dealing specifically withVLAN segments(as opposed to Overlay segments), the traffic does not leave the host encapsulated in Geneve; instead, it is tagged with a standard 802.1Q header.
Traceflowis the primary diagnostic tool within NSX for identifying where a packet is being dropped. It allows an administrator to inject a synthetic packet into the data plane from a source (such as a VM vNIC) to a destination. The tool then reports back every "observation point" along the path, including switching, routing, and firewalling. If a packet is dropped by a Distributed Firewall (DFW) rule or a physical misconfiguration that wasn't caught initially, Traceflow will explicitly state at which stage the packet was lost.
Packet Captureis the second essential tool. NSX provides a robust, distributed packet capture utility that can be executed from the NSX Manager CLI or UI. This tool allows administrators to capture traffic at various points, such as the vNIC, the switch port, or the physical uplink (vmnic) of the ESXi Transport Node. By comparing captures from different points, an administrator can determine if a packet is reaching the virtual switch but failing to exit the physical NIC, or if return traffic is reaching the host but not the VM.
Options likeFlow MonitoringandLive Floware excellent for observing traffic patterns and session statistics (IPFIX), but they are less effective for pinpointing the exact cause of "packet loss" compared to the granular, packet-level analysis provided by Traceflow and Packet Capture.Activity Monitoringis typically used for endpoint introspection and user-level activity, which is irrelevant to Layer 2/3 packet loss troubleshooting.

NO.2 A sovereign cloud provider has a VMware Cloud Foundation (VCF) stretched Workload Domain across two data centers (AZ1 and AZ2), where site connectivity via Layer 3 is provided by the underlay. The following NSX details are included in the design:
* Each site must host its own local NSX Edge Cluster for availability zones.
* Tier-0 gateways must be configured in active/active mode with BGP ECMP to local top-of-rack switches.
* Inter-site Edge TEP traffic must not cross the inter-DC link.
* SDDC Manager is used to automate NSX deployment.

During deployment of the Edge Cluster for AZ2, the SDDC Manager workflow fails because the Edge transport nodes' TEP IPs are not reachable from the ESXi transport nodes. Which step ensures correct Edge Cluster deployment in multi-site stretched domains?

A. Disable the liveness check during Edge deployment in SDDC Manager.

B. Configure BGP neighbors before deploying the Edge Cluster.

C. Reuse the TEP IP pool from AZ1.

D. Create an AZ2-specific Edge TEP IP pool and map it to the AZ2 uplink profile before deploying the Edge Cluster.

Answer: D

Explanation:

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

In aVMware Cloud Foundation (VCF)stretched cluster or Multi-Availability Zone (Multi-AZ) architecture, the networking design must account for the fact that AZ1 and AZ2 typically reside in different Layer 3 subnets. While the NSX Overlay provides Layer 2 adjacency for virtual machines across sites, the underlying Tunnel Endpoints (TEPs)must be able to communicate over the physical Layer 3 network.

According to the VCF Design Guide for Multi-AZ deployments, when stretching a workload domain, each availability zone should have its own dedicatedTEP IP Pool. This is because TEP traffic is encapsulated (Geneve) and routed via the physical underlay. If the Edge nodes in AZ2 were to use the same IP pool as AZ1 (Option C), the physical routers would likely encounter routing conflicts or reachability issues, as the subnet for AZ1 would not be natively routable or "local" to the AZ2 Top-of-Rack (ToR) switches.

The failure during the SDDC Manager workflow occurs because the automated "Liveness Check" or "Pre- validation" step attempts to verify that the newly assigned TEP IPs in AZ2 can reach the existing TEPs in the environment. To resolve this and ensure a successful deployment, the administrator must define a uniqueAZ2- specific IP Poolin NSX. Furthermore, this pool must be associated with anUplink Profile(or a Sub-Transport Node Profile in VCF 5.x/9.0) that uses the specific VLAN tagged for TEP traffic in the second data center.

This ensures that the Edge Nodes in AZ2 are assigned IPs that are valid and routable within the AZ2 underlay, allowing Geneve tunnels to establish correctly to the ESXi hosts in both sites without requiring a stretched Layer 2 physical network for the TEP infrastructure.

NO.3 An administrator is responsible for the management of a VMware Cloud Foundation (VCF) Fleet that consists of two VCF instances that are located in different physical locations. The administrator has been tasked with configuring a VPN between the two locations and has been tasked with identifying the two supported NSX Gateway configurations for an IPSec VPN. Drag and drop two items from the list of Possible Configurations into the list of Supported Configurations in any order.(Choose two.)

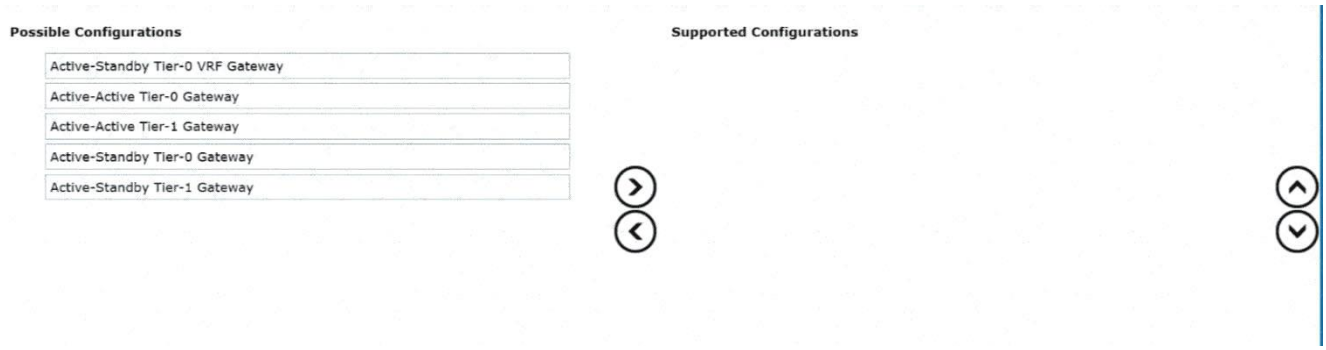

Answer:

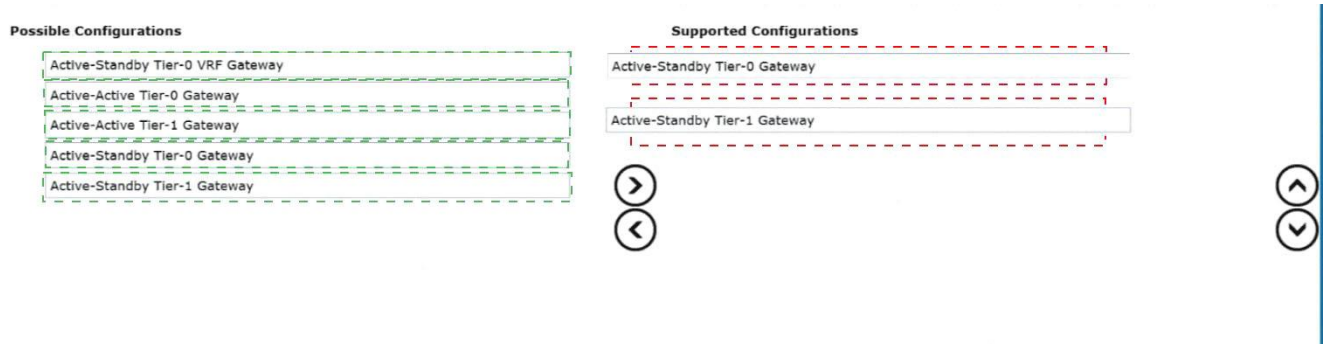

Explanation:
* Active-Standby Tier-0 Gateway
* Active-Standby Tier-1 Gateway
In aVMware Cloud Foundation (VCF)multi-site or multi-instance architecture, established viaNSX Federation, secure connectivity between sites is often achieved throughIPSec VPN. IPSec VPN is considered a stateful service within the NSX networking stack.
Stateful services-which also include NAT and Load Balancing-require a centralized point of processing to maintain the security association (SA) and session state tables. In the NSX gateway architecture, this necessitates the presence of aService Router (SR)component. For stateful consistency and to avoid session disruption that would occur if asymmetric traffic were processed by different nodes, these gateways must operate in anActive-Standbyhigh-availability mode.
According to the "NSX-T Data Center VPN Configuration Guide," IPSec VPN services can be deployed on either the provider tier (Tier-0 Gateway) or the tenant tier (Tier-1 Gateway). When configured on a Tier-0 gateway, the VPN typically provides broad connectivity between the physical infrastructure of two sites.
When configured on a Tier-1 gateway, it often provides targeted connectivity for a specific project or department's workload segments.
Configurations involvingActive-Activegateways (whether Tier-0 or Tier-1) do not support the native NSX IPSec VPN service because the ECMP (Equal Cost Multi-Pathing) nature of Active-Active mode could lead to packets belonging to the same VPN tunnel being processed by different Edge nodes, which cannot share the real-time encryption state. Therefore, for an administrator to successfully implement a cross-location VPN in a VCF Fleet, they must ensure the target gateway-be it Tier-0 or Tier-1-is deployed inActive-Standby mode.

NO.4 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?
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:
InVMware Cloud Foundation (VCF), theSDDC Managerautomates the deployment and expansion ofNSX 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 (specificallyKB 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 tolog into the CLI of the Edge nodeidentified 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.

NO.5 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?
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 theNSXmulti-tier routing architecture used by VCF, aTier-1 Gatewayis composed of two primary

components: theDistributed Router (DR)and theService 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 theRouter Link. This link is encapsulated inGenevejust 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 theMTU (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 butcannot 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 of1600(ideally9000) 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.

NO.6 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. Static routes do not support Equal Cost Multi-Pathing (ECMP) in NSX.

B. The static route Scope is set to only one uplink interface, but the next-hops are on two different VLANs.

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: B
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.