

VMware.3V0-25.25.v2026-08-16.q20

Exam Code:	3V0-25.25
Exam Name:	Advanced VMware Cloud Foundation 9.0 Networking
Certification Provider:	VMware
Free Question Number:	20
Version:	v2026-08-16
# of views:	114
# of Questions views:	200
https://www.freeqas.com/qa/VMware/3V0-25.25/VMware.3V0-25.25.v2026-08-16.q20.html	

NEW QUESTION: 1

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 (LEAVE A REPLY)

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 the NSX Tier-0 Gateway processes forwarding entries. In an Active/Active configuration, 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 to Uplink-1 only. Because R2 (192.168.101.1) is on a different subnet/VLAN (VLAN 101) that is associated with Uplink-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, the Scope should 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: 2

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

Answer: (SHOW ANSWER)

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.

NEW QUESTION: 3

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?

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

Answer: D (LEAVE A REPLY)

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.

NEW QUESTION: 4

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: ([SHOW ANSWER](#))

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

In the NSX multi-tier routing architecture used by VCF, a Tier-1 Gateway is composed of two primary components: the Distributed Router (DR) and the Service 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 the Router Link. This link is encapsulated in Geneve just 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 the MTU (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 but cannot 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 of 1600 (ideally 9000) 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.

NEW QUESTION: 5

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?

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: (SHOW ANSWER)

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.


NEW QUESTION: 6

An administrator is creating NSX segments in an environment. The NSX segment on an ESX Host is not realized. To troubleshoot the issue, the administrator needs to track the communication of components in the environment.


Drag and drop the component to the appropriate location in the diagram to track the path from desired state to completed state.

Policy	Manager
CCP	APH
NSX-Proxy	nsxt-vdl2

Answer Area



Answer:



Explanation:

Answer Area Placement:

- * NSX Manager Top-Left Box:Policy
- * NSX Manager Top-Middle Box:Manager
- * NSX Manager Top-Right Box:CCP (Central Control Plane)
- * NSX Manager Bottom Box:APH (Asynchronous Proxy Handler)
- * ESXi Host Top Box:NSX-Proxy
- * ESXi Host Bottom Box:nsxt-vdl2

In VMware Cloud Foundation (VCF) and NSX architectures, the realization of a logical object (like a segment) involves a multi-step communication flow across different management and control plane layers.

The Management Plane (NSX Manager)

- * **Policy:** The entry point where the "Desired State" is defined by the user or automation.
- * **Manager:** Receives the policy, validates it, and stores it in the management database.
- * **CCP (Central Control Plane):** Processes the logical configuration and computes the actual instructions needed for the data plane.

* APH (Asynchronous Proxy Handler): Acts as a broker on the NSX Manager, responsible for pushing these instructions down to the transport nodes via NSX RPC TCP 1234 (Management) and NSX RPC TCP 1235 (Control).

The Local Control Plane (ESXi Host)

* NSX-Proxy: A local agent on the ESXi host that maintains a persistent connection to the APH. It receives the instructions and ensures the "Local Control Plane" state matches the "Central Control Plane" intent.

* nsxt-vdl2: The final component in the chain. It interacts directly with the ESXi kernel modules to program the Virtual Distributed Switch (VDS) and realize the segment on the host. Once this step is finished, the segment moves to the "Completed State" and is ready for use.

NEW QUESTION: 7

An administrator is tasked to enable users to configure an individual VPC, but not create subnets. What three NSX roles would the administrator assign to allow access without the ability to create subnets? (Choose three.)

- A. Security Admin
- B. Network Admin
- C. VPC Admin
- D. Security Operator
- E. Network Operator

Answer: C,D,E ([LEAVE A REPLY](#))

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

With the introduction of the Virtual Private Cloud (VPC) consumption model in VCF 9.0 and late 5.x releases, Role-Based Access Control (RBAC) has become more granular to support true multi-tenancy. A VPC is designed to be a self-contained "container" for a department's or user's networking resources.

To meet the specific requirement where a user can configure aspects of an individual VPC but is restricted from creating new subnets (which involves modifying the underlying network CIDR blocks and IPAM), a combination of specific roles is required.

* VPC Admin: This is the primary role for the user within their assigned VPC. It allows the user to manage the overall VPC environment, including high-level settings and monitoring. However, the VPC Admin's power is often limited by the specific quotas and policies set by the Enterprise Admin.

* Security Operator: This role allows the user to view security configurations and policies without having the permission to modify the network fabric or create new infrastructure components like subnets. It provides the "read-only" visibility into the security posture of the VPC.

* Network Operator: Similar to the Security Operator, the Network Operator role provides visibility into the networking state—such as routing tables, segment status, and connectivity—without granting the

"Write" permissions required to provision new subnets or alter the network topology.

Assigning Network Admin (Option B) or Security Admin (Option A) would grant too much privilege, as these roles typically include the ability to create, delete, and modify subnets and firewall policies at a structural level. By combining the VPC Admin role with Operator-level roles, the administrator ensures the user has the necessary context to manage their assigned resources while strictly adhering to the restriction against creating new network subnets.

NEW QUESTION: 8

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?

- 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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 9

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?

- 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 (LEAVE A REPLY)

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.

NEW QUESTION: 10

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 (LEAVE A REPLY)

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 with VLAN 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.

Traceflow is 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 Capture is 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 (vnic) 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 like Flow Monitoring and Live Flow are 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 Monitoring is typically used for endpoint introspection and user-level activity, which is irrelevant to Layer 2/3 packet loss troubleshooting.

NEW QUESTION: 11

An administrator is troubleshooting intermittent connectivity failures between two workloads connected to NSX VLAN segments using Traceflow. In-band Network Telemetry (INT) has been enabled in the NSX Global Configuration. How does Traceflow identify issues in a VLAN network?

- A. Injects ICMP traffic into the data plane and observes the results in the control plane.
- B. Injects synthetic traffic into the data plane and observes the results in the control plane.
- C. Traceflow cannot be enabled to analyze VLAN network segments in NSX.
- D. Compares intended network state in the control plane with Tunnel End Point (TEP) keepalives in the data plane.

Answer: B (LEAVE A REPLY)

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

In VMware Cloud Foundation (VCF) and NSX, Traceflow is a powerful diagnostic tool designed to provide visibility into the logical and physical path of a packet as it traverses the SDDC. Unlike standard ping or traceroute utilities that use real ICMP traffic from the Guest OS, Traceflow operates by injecting synthetic traffic directly into the data plane at the source point (usually the vNIC of a Virtual Machine).

When Traceflow is initiated, the NSX Manager creates a "trace packet" that mimics the characteristics of the traffic being investigated (such as TCP, UDP, or ICMP with specific headers). This synthetic packet is marked with a special metadata tag. As the packet moves through the virtual switches (VDS), logical routers (DR/SR), and distributed firewalls (DFW) on the ESXi Transport Nodes, each component recognizes the tag and reports an "observation" back to the Central Control Plane (CCP). The CCP then aggregates these observations and presents them in the NSX Manager UI.

For VLAN-backed segments, Traceflow functions similarly to how it works on Overlay segments. It tracks the packet as it is switched at Layer 2 and processed by any applicable distributed services. The inclusion of In-band Network Telemetry (INT) in modern VCF versions (5.x and 9.0) enhances this by allowing the synthetic packet to collect telemetry data from INT-capable physical switches in the fabric. This provides a

"hop-by-hop" view that includes both the virtual and physical segments of the journey.

Option A is incorrect because Traceflow is not limited to ICMP; it can simulate various protocols. Option C is incorrect as Traceflow fully supports VLAN segments. Option D is incorrect as it describes a state- comparison mechanism rather than the active injection process that defines Traceflow. Therefore, the injection of synthetic traffic to observe data plane behavior via the control plane is the verified mechanism.

NEW QUESTION: 12

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.)

- 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 (LEAVE A REPLY)

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.

NEW QUESTION: 13

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?

- 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 (LEAVE A REPLY)

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.

NEW QUESTION: 14

An administrator has been tasked with enabling OSPF as the routing protocol for a Tier-0 Gateway. Which two items must be configured to enable OSPF for a Tier-0 Gateway? Mark two answers by clicking the two correct locations on the image. (Choose two.)

Answer Area

vmware®

freeqas.com

Name	HA Mode ⓘ	Linked Tier-1 Gateways	Linked Segments
Max 30 allowed. Click (+) to add			

> INTERFACES

> ROUTING

> BGP Disabled

> OSPF Enabled

OSPF Enabled

Graceful Restart Helper Only

Route Summarization Set

OSPF Configured Interface Set ⓘ

Area Definition Set

ECMP Enabled

Default Route Originate Disabled ⓘ

OSPF Neighbor View

Answer:

Answer Area



Explanation:

To enable OSPF on a Tier-0 Gateway within a VMware Cloud Foundation (VCF) or NSX environment, an administrator must define the areas where the gateway will participate and the specific interfaces it will use for peering.

Based on the NSX Manager configuration interface, the two required items to be configured are:

* Area Definition: This is necessary to define the OSPF area (e.g., Area 0) the Tier-0 gateway will participate in.

* OSPF Configured Interfaces: OSPF must be explicitly configured on the relevant uplink interfaces to establish neighbor relationships and exchange routing information with physical routers.

NEW QUESTION: 15

When using a DHCP Relay on a segment, which design restriction must be considered?

- A. DHCP settings, DHCP options, and static bindings cannot be configured on the segment.
- B. DHCP client requests cannot be relayed to the external DHCP servers.
- C. DHCP settings, DHCP options, and static bindings can be configured on the segment.
- D. DHCP Relay service is available to all the other segments in the network.

Answer: A (LEAVE A REPLY)

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

In VMware Cloud Foundation (VCF) networking, IP address management within an NSX segment can be handled by either the native NSX DHCP server or by an external DHCP server. When an administrator chooses to use an existing external corporate DHCP infrastructure, they must configure a DHCP Relay on the logical segment.

The DHCP Relay works by intercepting the initial DHCP Discover broadcast from a workload VM and forwarding it (as a unicast packet) to the specified IP address of the external DHCP server. However, NSX enforces a strict mutual exclusivity in its configuration logic to prevent conflicts and unpredictable address assignments.

According to the "NSX-T Data Center Administration Guide," once a segment is configured to use a DHCP Relay profile, the native NSX DHCP capabilities for that specific segment are disabled. This means that DHCP settings, DHCP options, and static bindings cannot be configured on that segment (Option A). All such configurations, including IP reservations and scope options (like DNS or NTP), must be managed centrally on the external DHCP server.

Option C is incorrect because the UI will physically grey out or prevent the entry of native DHCP parameters once the Relay is selected. Option B is incorrect as the primary purpose of a Relay is precisely to forward requests to external servers. Option D is incorrect because a DHCP Relay is configured on a per-segment or per-gateway basis; it is not a "global" service that automatically covers all other segments in the network.

Therefore, the architectural trade-off when choosing a Relay is the shift of all management and binding logic to the external physical or virtual DHCP appliance.

NEW QUESTION: 16

An administrator is enabling IPv6-to-IPv4 communication for workloads hosted in an NSX environment. The workloads use IPv6-only addressing, but the external systems they must reach are IPv4-only. To provide this translation service, the administrator decides to configure NAT64. Which two following characteristics about NAT64 are true? (Choose two.)

- A. NAT64 is stateless and requires gateways to be deployed in active-standby mode.
- B. NAT64 requires the Tier-1 gateway to be configured in active-standby mode.
- C. NAT64 is supported on Tier-1 gateways only.
- D. NAT64 is supported on Tier-0 and Tier-1 gateways.
- E. NAT64 requires the Tier-1 gateway to be configured in active-active mode.

Answer: ([SHOW ANSWER](#))

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

As organizations modernize their infrastructure with VCF 5.x and 9.0, IPv6 adoption becomes more prevalent.

NAT64 is a critical transition technology that allows IPv6-only hosts to communicate with IPv4-only resources by translating the packet headers.

In NSX, NAT64 is a stateful service. Stateful services in the NSX architecture require a centralized point of processing to maintain the session state table. Because of this requirement, any gateway (Tier-0 or Tier-1) providing NAT64 services must be configured in Active-Standby high availability mode. In Active-Active mode, asymmetric return traffic could hit a different Edge node that does not have the session information, causing the translation to fail. This is a fundamental design constraint for stateful NAT in NSX.

Furthermore, VMware NSX documentation specifies that NAT64 is a flexible service that can be implemented at multiple tiers of the logical routing hierarchy. It is supported on both Tier-0 and Tier-1 gateways. The choice of where to place the NAT64 service depends on the design requirements: placing it on the Tier-1 gateway allows for tenant-specific translation and offloads the Tier-0, while placing it on the Tier-0 provides a centralized translation point for all connected segments.

Option A is incorrect because NAT64 in NSX is stateful, not stateless. Option C is incorrect because it is not limited to Tier-1. Option E is incorrect because Active-Active mode does not support the stateful nature of the NAT64 engine. Consequently, the correct architecture requires an Active-Standby configuration on either a Tier-0 or Tier-1 gateway to properly facilitate the translation between the IPv6 workloads and the IPv4 external world.

Valid 3V0-25.25 Dumps shared by PrepPdf.com for Helping Passing 3V0-25.25 Exam! PrepPdf.com now offer the **newest 3V0-25.25 exam dumps**, the PrepPdf.com 3V0-25.25 exam questions have been updated and answers have been corrected get the **newest** PrepPdf.com 3V0-25.25 dumps with Test Engine here: <https://www.preppdf.com/VMware/3V0-25.25-prepaway-exam-dumps.html> (64 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 17

An administrator has been tasked with providing a networking solution including a Source and Destination NAT for a single Tenant. The tenant is using Centralized Connectivity with a Tier-0 Gateway named Ten-A- Tier-0 supported by an Edge cluster in Active-Active mode. The NAT solution must be available for multiple subnets within the Tenant space. The administrator chooses to deploy a Tier-1 Gateway to implement the NAT solution. How would the administrator complete the task?

- A. Change Ten-A-Tier-0 to Active-Standby to support the stateful NAT.
- B. Create a new Tier-0 Gateway in Active-Standby mode and attach another Tier-1 Gateway.
- C. Create a Tier-1 Gateway in Distributed Routing mode only and do not attach it to Ten-A-Tier-0.
- D. Create a new Tier-1 Gateway in Active-Standby mode and attach it to Ten-A-Tier-0.

Answer: D (LEAVE A REPLY)

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

In a VMware Cloud Foundation (VCF) environment, the implementation of stateful services-such as Source NAT (SNAT) and Destination NAT (DNAT)-requires a specific architectural configuration within the NSX component. This is because stateful services need a centralized point of processing (a Service Router or SR) to maintain the session state tables and ensure that return traffic is processed by the same node that initiated the session.

The scenario describes a provider-level Tier-0 Gateway running in Active-Active mode. While Active-Active provides high-performance North-South throughput via ECMP (Equal Cost Multi-Pathing), it does not support stateful NAT services because asymmetric traffic flows would break the session tracking. Rather than changing the Tier-0 to Active-Standby (which would reduce overall throughput for the entire environment), the architecturally sound approach is to offload the stateful services to a Tier-1 Gateway.

According to VCF design guides, when a Tier-1 Gateway is required to perform NAT for multiple subnets, it must be configured as a Stateful Tier-1. This involves associating the Tier-1 with an Edge Cluster and setting its high-availability mode to Active-Standby. Once the Tier-1 is created in this mode, it creates a Service Router (SR) component on the selected Edge Nodes. By attaching this Active-Standby Tier-1 to the existing Active-Active Tier-0 (Ten-A-Tier-0), the tenant's subnets can enjoy the benefits of localized stateful NAT while the environment maintains high-performance, non-stateful routing at the Tier-0 layer.

Option A is inefficient as it impacts the entire Tier-0. Option B is redundant. Option C is incorrect because a "Distributed Routing only" Tier-1 (one without an Edge Cluster association) cannot perform stateful NAT.

Therefore, creating an Active-Standby Tier-1 and linking it to the provider Tier-0 is the verified VCF multi-tenant design pattern.

NEW QUESTION: 18

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?

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 (LEAVE A REPLY)

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.

NEW QUESTION: 19

An administrator has noticed that both the active and standby Global Managers have gone offline. What is the correct sequence of events to restore the Global Managers?

Delete both the active and standby Global Managers. Ensure there are no Global Manager appliances up in any other clusters.

Deploy a new Global Manager with the same IP address/FQDN as the old active Global Manager.

Restore the active Global Manager from backup.

Deploy an additional new Global Manager on another site and onboard it to the restored Global Manager.

vmware®

Answer:

Delete both the active and standby Global Managers. Ensure there are no Global Manager appliances up in any other clusters.

Deploy a new Global Manager with the same IP address/FQDN as the old active Global Manager.

Restore the active Global Manager from backup.

Deploy an additional new Global Manager on another site and onboard it to the restored Global Manager.

vmware®

Explanation:

- * Step 1: Delete both the active and standby Global Managers. Ensure there are no Global Manager appliances up in any other clusters.
- * Step 2: Deploy a new Global Manager with the same IP address/FQDN as the old active Global Manager.
- * Step 3: Restore the active Global Manager from backup.
- * Step 4: Deploy an additional new Global Manager on another site and onboard it to the restored Global Manager.

In a VMware Cloud Foundation multi-site deployment using NSX Federation, the Global Manager (GM) manages the global networking configuration across multiple sites. If the entire GM cluster (Active and Standby) fails, the following architectural principles apply:

- * Cleanup (Step 1): Before initiating a restore, the environment must be "cleaned." If old, failed VMs remain in the inventory or on the hosts, they can cause IP address conflicts or UUID mismatches during the deployment of the new appliance. You must ensure the management plane is clear of the original failed nodes.
- * Identity Consistency (Step 2): When restoring an NSX appliance (Local or Global) from backup, the new appliance must be deployed with the exact same IP address and FQDN as the original active node.

This is critical because the existing Local Managers (LMs) at each site already have established thumbprints and communication channels tied to that specific identity.

- * The Restore Operation (Step 3): Once the "seed" appliance is deployed, the restore process is triggered through the NSX Manager UI/API. This process re-populates the database with the global segments, firewall rules, and Tier-0/Tier-1 configurations.
- * Restoring Redundancy (Step 4): The backup only contains the configuration of the cluster. It does not "restore" the standby VM itself. High Availability (HA) must be manually re-established by deploying a second GM appliance at the secondary site and joining it to the newly restored Global Manager cluster to act as the standby.

NEW QUESTION: 20

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: (SHOW ANSWER)

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

In a VMware 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 dedicated TEP 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.

Valid 3V0-25.25 Dumps shared by PrepPdf.com for Helping Passing 3V0-25.25 Exam! PrepPdf.com now offer the **newest 3V0-25.25 exam dumps**, the PrepPdf.com 3V0-25.25 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com 3V0-25.25 dumps with Test Engine here: <https://www.preppdf.com/VMware/3V0-25.25-prepaway-exam-dumps.html> (**64** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)