Vendor: VMware

Exam Code: 3V0-25.25

Exam Name: Advanced VMware Cloud Foundation 9.0 Networking

Version: DEMO

QUESTION 1
An administrator has a standalone vSphere 8.0 Update 1a deployment that is running with VMware NSX 4.1.0.2 and has to converge the deployment into a new VMware Cloud Foundation (VCF) instance. How can the administrator accomplish this task?

A. Manually upgrade both vSphere and NSX to version 9 prior to converging. Then use the VCF Installer to converge the vSphere 9 and NSX 9 instances into a new VCF management domain.
B. Manually upgrade vSphere to version 9. Then use the VCF Installer to converge the vSphere 9 environment into a new VCF management domain. Then use the VCF lifecycle management tools to upgrade NSX to version 9.
C. Use the VCF Installer to converge the existing vSphere 8 and NSX 4 environment into a new VCF management domain. Then use the VCF lifecycle management tools to upgrade to 9.
D. Manually upgrade vSphere to version 9 and uninstall NSX 4. Then use the VCF Installer to converge the vSphere 9.0 environment into a new VCF management domain at which time NSX 9 will be reinstalled.

Answer: C
Explanation:
The process of bringing existing infrastructure under VCF management is known as "VCF Import" or "Convergence." This is a common path for organizations transitioning from siloed management to the full SDDC stack provided by Cloud Foundation.

According to the VCF 5.x and 9.0 documentation, the VCF Installer (specifically the Cloud Foundation Builder and the Import Tool) is designed to ingest existing environments. The verified best practice is to converge the environment at its current, supported version, provided it meets the minimum baseline requirements for the VCF version you are deploying.

In this scenario, vSphere 8.0 U1 and NSX 4.1 are compatible versions that can be imported into a VCF management framework. By using the VCF Installer to converge the existing environment first (Option C), the SDDC Manager takes ownership of the existing vCenter and NSX Manager. Once the environment is "VCF-aware," the administrator gains the benefit of SDDC Manager's Lifecycle Management (LCM).

The SDDC Manager then handles the orchestrated, multi-step upgrade to version 9.0. This ensures that the automated "Bill of Materials" (BOM) is strictly followed, ensuring compatibility between vCenter, ESXi, and NSX components. Attempting to manually upgrade components to version 9 before convergence (Options A and B) or uninstalling NSX (Option D) creates a "Frankenstein" environment that may not align with the VCF BOM, making the automated convergence process fail or resulting in an unsupported configuration. The principle of VCF is to bring the environment in first, then let VCF manage the upgrades.

QUESTION 2
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

Explanation:
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.

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
Explanation:
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.

QUESTION 4
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
Explanation:
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.

QUESTION 5
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
Explanation:
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.

QUESTION 6
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:
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 unique AZ2-specific IP Pool in NSX. Furthermore, this pool must be associated with an Uplink 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.

QUESTION 7
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 pee";
- 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. The static default route Scope is set only to the uplink VLAN 100 segment.
B. The ToR routers do not have routes back to the Edge uplink interfaces.
C. Multi-hop eBGP is required when using two VLANs.
D. BGP authentication mismatch between Tier-0 and ToR routers.

Answer: A
Explanation:
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.

QUESTION 8
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: CE
Explanation:
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.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**