VMware.3V0-25.25.v2026-03-09.q25

Exam Code:	3V0-25.25
Exam Name:	Advanced VMware Cloud Foundation 9.0 Networking
Certification Provider:	VMware
Free Question Number:	25
Version:	v2026-03-09
# of views:	145
# of Questions views:	250
https://www.freepdfdumps.com/VMware.3V0-25.25.v2026-03-09.q25.html	

NEW QUESTION: 1
Which of the following statements is true when configuring Remote Tunnel End Points (RTEPs) with NSX Federation?
A. TEP and RTEP networks must use separate physical NICs.
B. RTEP needs to be configured on only one edge node.
C. The default MTU for the RTEP network is 1500.
D. DHCP must be used to assign IP addresses to the RTEP.
Answer: C (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In anNSX Federationdeployment, which is a key component of multi-siteVMware Cloud Foundation (VCF)architectures, theRemote Tunnel End Point (RTEP)is used specifically for inter-site communication.
While standard TEPs (Tunnel End Points) handle overlay traffic within a single site (East-West), RTEPs facilitate the encapsulation of traffic that needs to traverse the Layer 3 network between different geographical locations.
A critical design consideration for RTEP is theMaximum Transmission Unit (MTU). Within a local VCF site, jumbo frames (MTU 1600 or 9000) are highly recommended and often required for the Geneve overlay to account for encapsulation overhead. However, when traffic leaves a site to travel over a WAN or a provider's long-haul network, it often encounters physical infrastructure that only supports the standard internet MTU of1500 bytes.
According to VMware's "NSX Federation Design Guide," the default MTU setting for the RTEP configuration is1500. This ensures that inter-site traffic can pass through standard routers and VPNs without being dropped due to size constraints. If the inter-site physical links support larger frames, this value can be increased, but 1500 remains the baseline compatible default.
Regarding the other options:Ais incorrect because TEP and RTEP can share the same physical N-VDS and physical NICs (pNICs) by using different VLANs or subnets.Bis incorrect because

every Edge node within a cluster that is participating in the Federation must have an RTEP configured to ensure high availability and proper traffic processing for global segments.Dis incorrect as IP addresses for RTEPs are typically assigned viaStatic IP Poolsmanaged within NSX to ensure consistency and ease of tracking across sites, rather than relying on DHCP which is less common in data center backbone configurations.

NEW QUESTION: 2

An administrator was asked to explain the characteristic and requirements of Centralized Connectivity Mode which is planned to be configured in one of the workload domains in VMware Cloud Foundation (VCF) environment.

Drag and drop four options from the Options list on the left and place them into the Centralized Connectivity Mode on the right in any order. (Choose four.)

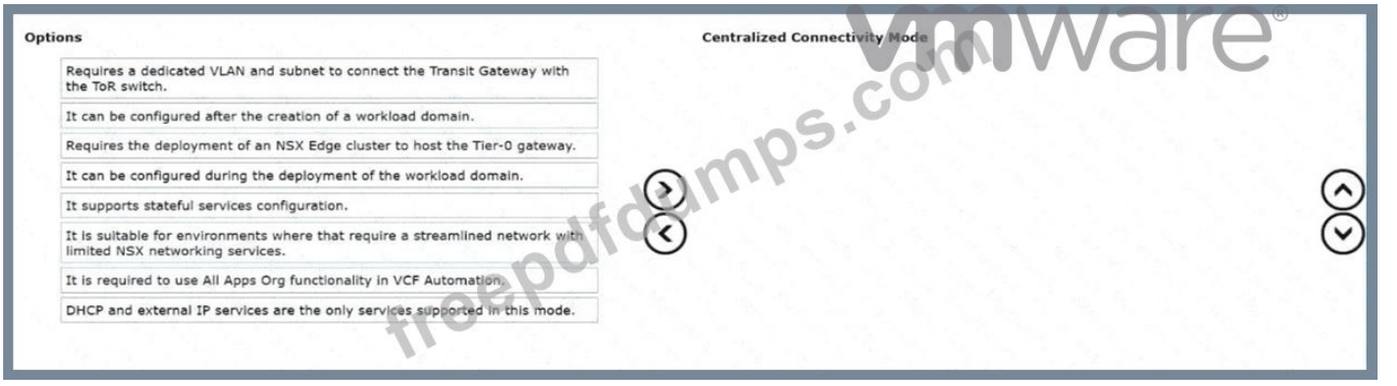

Answer:

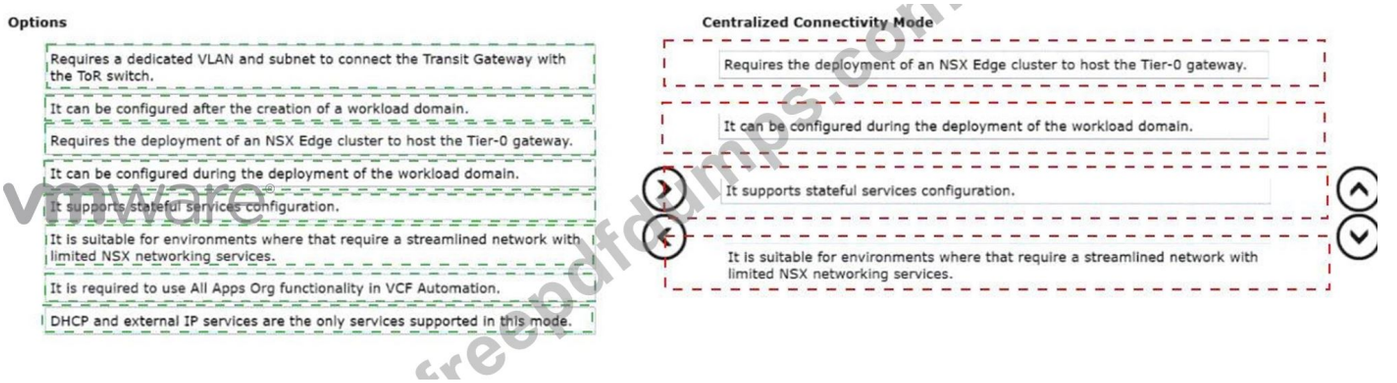

Explanation:

* Requires the deployment of an NSX Edge cluster to host the Tier-0 gateway.
* It can be configured during the deployment of the workload domain.
* It supports stateful services configuration.
* It is suitable for environments that require a streamlined network with limited NSX networking services.

InVMware Cloud Foundation (VCF) 9.0, the networking architecture introduces specialized connectivity modes to cater to different organizational needs, withCentralized Connectivity Modebeing a primary option for streamlined deployments. This mode is fundamentally anchored to the physical infrastructure via localized resources rather than distributed components across the entire cluster.

The most critical technical requirement for this mode is that itrequires the deployment of an NSX Edge cluster to host the Tier-0 gateway. Unlike distributed models, centralized connectivity funnels North-South traffic through specific Edge nodes that serve as the demarcation point between the virtual overlay and the physical Top-of-Rack (ToR) switches. This centralization is what enables the next key characteristic: it supports stateful services configuration. Because traffic is anchored to specific Service Routers (SR) on Edge nodes, stateful operations such as NAT, Load Balancing, and stateful firewalls can maintain session persistence, which is not natively possible in a purely distributed Active/Active ECMP environment without specialized configuration.
From a lifecycle perspective, this mode is highly integrated into the SDDC Manager workflows andcan be configured during the deployment of the workload domain. This allows architects to define the networking posture of a new domain at "Day 0," ensuring that the necessary Edge resources and Tier-0/Tier-1 hierarchies are provisioned automatically to meet the domain's specific requirements.
Finally, Centralized Connectivity Modeis suitable for environments that require a streamlined network with limited NSX networking services. It provides a "cloud-lite" approach to networking, offering the necessary isolation and security of NSX without the complexity of managing a full-scale distributed fabric.
This makes it an ideal choice for smaller workload domains, specialized labs, or legacy application environments that do not require the massive scale of a distributed transit gateway but still need robust stateful security and simplified North-South egress.

NEW QUESTION: 3
An administrator has a VMware Cloud Foundation (VCF) instance. A critical NSX security update has been released by Broadcom. How can the administrator install the NSX update?
A. Download the NSX patch to the NSX Manager. Apply it using VCF Operations Fleet Management.
B. Download the NSX patch to VCF Operations. Apply it using NSX Manager.
C. Download the NSX patch to VCF Operations. Apply it using VCF Operations Fleet Management.
D. Download the NSX patch to the NSX Manager. Apply it using NSX Manager.
Answer: (SHOW ANSWER)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In the unified architecture ofVMware Cloud Foundation (VCF) 9.0, the management paradigm has shifted towards a more centralized "Fleet Management" approach. Historically, in VCF 4.x and 5.x, updates were primarily managed via the SDDC Manager using the Lifecycle Management (LCM) engine. However, with the integration advancements in version 9.0,VCF Operations(formerly part of the Aria/vRealize suite) has taken on a more direct role in the orchestration of updates across the entire VCF "Fleet." To comply with the VCF operational model, administrators no longer apply patches directly within the component managers (like NSX

Manager or vCenter) if they wish to remain within the supported, automated framework. Instead, the workflow begins by downloading the bundle or patch toVCF Operations. This ensures that the update is validated against the currentBill of Materials (BOM)and that all dependencies- such as compatibility with the underlying ESXi versions or the management vCenter-are checked before any changes are committed.
Once the patch is available in VCF Operations, the administrator utilizesFleet Managementto apply it. This service orchestrates the update across all NSX Managers and Transport Nodes (Edges and Hosts) in a controlled, non-disruptive manner. If the administrator were to apply the patch directly in NSX Manager (Option D), the SDDC Manager and VCF Operations databases would go out of sync, leading to a
"configuration drift" where the system no longer knows which version is actually running, potentially breaking future automated lifecycle tasks. Therefore, the centralized download and application throughVCF Operations Fleet Managementis the verified procedure for maintaining a healthy and compliant VCF 9.0 environment.

NEW QUESTION: 4
How should the Global Managers (GMs) and Local Managers (LMs) be distributed to ensure high availability and optimal performance in a multi-site NSX Federation deployment comprised of three sites? (Choose two.)
A. Each NSX site must have its own LM cluster that reports to the GM.
B. LMs are only needed on the primary site. Secondary sites can manage their local data plane directly via the GM.
C. LMs should only be deployed as single nodes to reduce overhead.
D. The GM cluster should be deployed across three sites.
E. The GM should be a single appliance placed in a central cloud environment to simplify connectivity, relying on vSphere HA for availability.
Answer: A,D (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In aVMware Cloud Foundation (VCF)Federation deployment across multiple sites, the management architecture is designed to provide "Global Visibility" while maintaining "Local Autonomy." This is achieved through the coordinated distribution ofGlobal Managers (GMs)andLocal Managers (LMs).
For a three-site deployment,NSX Federationbest practices mandate that each site maintains its ownLocal Manager (LM) Cluster(Option A). The LM is responsible for the site-specific control plane, communicating with local Transport Nodes (ESXi and Edges) to program the data plane. If the connection to the GM is lost, the LM ensures the local site continues to function normally. For production environments, these must be clusters (typically 3 nodes) rather than single nodes to ensure local management remains available.
To protect theGlobal Manageritself-which is the source of truth for all global networking and security policies-the GM cluster should bestretched across the three sites(Option D). In a

standard 3-node GM cluster, placing one node at each site ensures that the Federation management plane can survive the complete failure of an entire site. This "stretched" cluster configuration provides a high level of resilience and ensures that an administrator can still manage global policies from any surviving location.
Option B is incorrect because the GM does not communicate directly with the data plane of a site; it must go through an LM. Option C is a risk to availability. Option E is incorrect because vSphere HA cannot protect against a site-wide disaster, and a single appliance represents a significant single point of failure for the entire global network configuration.

NEW QUESTION: 5
An administrator is investigating reports that several Virtual Machines (VMs) deployed on an NSX virtual network segment are dropping packets. To troubleshoot the issue the administrator has attached two test VMs to the virtual network in order to inspect the packets sent between the two test VMs. What tool will allow the administrator to analyze the packet flow?
A. Flows Monitoring in the VCF Operations UI.
B. Traceflow in the NSX Manager UI.
C. Port Mirroring in the NSX Manager UI.
D. Live Traffic Analysis in the NSX Manager UI.
Answer: B (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In aVMware Cloud Foundation (VCF)environment, pinpointing the exact location of packet drops within the software-defined data center requires tools that can see into the logical forwarding pipeline. While traditional networking tools like pings only provide a "binary" up/down status,Traceflowis the definitive diagnostic tool within theNSX Manager UIfor deep packet path analysis.
Traceflow works by injecting a synthetic "trace packet" into the data plane, originating from a source vNIC of a specific VM. This packet is uniquely tagged so that every NSX component it touches-including the Distributed Switch (VDS), Distributed Firewall (DFW) rules, Distributed Routers (DR), and Service Routers (SR) on Edge nodes-reports back an observation.
When an administrator observes packet drops, Traceflow provides a step-by-step visualization of the packet's journey. If the packet is dropped, Traceflow will explicitly identify the component responsible. For example, it might show that the packet was "Dropped by Firewall Rule #102" or "Dropped by SpoofGuard." It can also identify if the packet was lost during Geneve encapsulation or at the physical uplink interface.
Option A (Flows Monitoring) is useful for long-term traffic patterns and session statistics but lacks the packet- level "hop-by-hop" granular detail provided by Traceflow. Option C (Port Mirroring) is used to send a copy of traffic to a physical or virtual appliance (like a Sniffer or IDS), which is more complex to set up and usually reserved for external deep packet inspection (DPI) rather than internal path troubleshooting. Option D (Live Traffic Analysis) is a broader term, but within the context of the NSX troubleshooting toolkit for "packet flow analysis" between two

points,Traceflowis the verified and documented solution for verifying the logical path and identifying drops.

NEW QUESTION: 6
An administrator is configuring an NSX segment used by a nested hypervisor deployment where an ESXi VM runs on an ESXi host and multiple VMs run inside the ESXi VM. Which segment profile must be created to satisfy the request?
A. IP Discovery
B. Security
C. MAC Discovery
D. Spoof Guard
Answer: (SHOW ANSWER)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
Nested virtualization-where a hypervisor like ESXi is run as a virtual machine-imposes unique challenges on the virtual networking layer. In a standard VCF environment, an NSX segment port expects to see exactly one MAC address: the MAC address assigned to the VM's vNIC.
When you run anested hypervisor, that single vNIC now acts as an "uplink" for multiple "inner" virtual machines. Consequently, traffic originating from that single nested ESXi VM will contain many different source MAC addresses (one for each nested VM). By default, the NSX/VDS security and switching logic will drop this traffic because it appears asMAC Spoofing-packets are arriving from a port with source MACs that do not match the port's registered ID.
To support this, aMAC Discovery Segment Profilemust be configured and applied to the segment. Within this profile, the administrator must enableMAC Learning. MAC Learning allows the NSX virtual switch to
"learn" and permit multiple MAC addresses on a single logical port. Without this, only the primary MAC of the nested ESXi host would be allowed, and all nested VMs would lose connectivity to the rest of the network.
In VCF 5.x and 9.0 documentation, this is a standard requirement for "Lab-on-a-Lab" designs or development environments. WhileIP Discovery(Option A) andSpoof Guard(Option D) are important for maintaining the IP-to-MAC binding and preventing IP theft, they do not address the fundamental Layer 2 requirement of allowing multiple MAC identities on a single port.
Therefore,MAC Discoverywith MAC learning enabled is the verified profile choice for nested hypervisor support.

NEW QUESTION: 7
An administrator is preparing to deploy a new workload domain that will host vSphere Kubernetes Service (VKS) clusters. Before configuring the network for the Kubernetes clusters, the administrator needs to create a Tier-0 Gateway to handle North/South connectivity. What is the requirement for creating a Tier-0 Gateway for use with a workload domain that is running the vSphere Kubernetes service (VKS) with VPC?

A. The Tier-0 Gateway route map must contain an IP prefix with only a deny rule.
B. The Tier-0 Gateway must be configured in Non-Preemptive failover mode.
C. The Tier-0 Gateway must be configured in Active/Standby mode.
D. The Tier-0 Gateway must have IPv6 enabled.
Answer: C (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
When deployingvSphere Kubernetes Service (VKS)-often referred to as Tanzu with VCF-within a Virtual Private Cloud (VPC)consumption model, the networking requirements are more stringent than a standard VM-only environment. This is because VKS relies on stateful services such asLoad Balancing(via the NSX Advanced Load Balancer or the native NSX LB) andNATto provide ingress and egress for Kubernetes pods and services.
In NSX architecture, any gateway that providesstateful servicesmust be configured inActive/Standbymode.
While an Active/Active Tier-0 gateway is excellent for high-throughput ECMP routing, it cannot support stateful features because return traffic might arrive at the "Standby" (or alternative Active) node which does not share the same session state table, resulting in dropped connections.
Specifically, for VKS clusters integrated with the VPC model in VCF 5.x and 9.0, the Tier-0 gateway acts as the provider-side gateway. To ensure that the KubernetesLoadBalancerservice types andSNAT/DNATfor pods function correctly and maintain session persistence, the gateway must be anchored to a specific Service Router (SR) on an Edge node. This is only possible in anActive/Standbyconfiguration.
Option B (Non-Preemptive) is a failoversettingbut not the primary architectural requirement.
Option D (IPv6) may be used depending on the specific network design, but it is not a mandatory requirement for VKS functionality. Option A is incorrect as route maps usually require "Permit" rules to actually function. Thus, the verified architectural prerequisite for a VKS/VPC-enabled workload domain is anActive/Standby Tier-0 Gateway.

NEW QUESTION: 8
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 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.

NEW QUESTION: 9
A large multinational corporation is seeking proposals for the modernization of a Private Cloud environment.
The proposed solution must meet the following requirements:
* Support multiple data centers located in different geographic regions.
* Provide a secure and scalable solution that ensures seamless connectivity between data centers and different departments.
Which three NSX features or capabilities must be included in the proposed solution? (Choose three.)
A. NSX Edge
B. AVI Load Balancer
C. vDefend
D. Virtual Private Cloud (VPC)
E. Centralized Network Connectivity
F. NSX L2 Bridging
Answer: (SHOW ANSWER)

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In a modern VMware Cloud Foundation (VCF) architecture, particularly when addressing the needs of a multinational corporation with geographically dispersed data centers, the solution must prioritize multi- tenancy, security, and consistent delivery. The integration ofNSXwithin VCF provides these core pillars.
First, theNSX Edgeis a foundational requirement for any multi-site or modern cloud environment. It serves as the bridge between the virtual overlay network and the physical world. In a multi-region deployment, NSX Edges facilitate North-South traffic and are essential for supporting features like Global Server Load Balancing (GSLB) or site-to-site connectivity. Without the Edge, the software-defined data center (SDDC) cannot communicate with external networks or peer via BGP with physical routers.
Second,vDefend(formerly known as NSX Security) provides the advanced security framework required for a
"secure and scalable" environment. This includes Distributed Firewalling (DFW), Distributed IDS/IPS, and Malware Prevention. For a corporation with different departments, vDefend allows for micro-segmentation, ensuring that a security breach in one department's segment cannot move laterally to another. This is critical for meeting compliance and isolation requirements across global regions.
Third, theVirtual Private Cloud (VPC)model is the cornerstone of the latest VCF 9.0 and 5.x architectures.
It enables the "scalable solution" for different departments by providing a self-service consumption model.
Each department can manage its own isolated network space, including subnets and security policies, without needing deep networking expertise or constant tickets for the central IT team. This abstraction simplifies management across multiple data centers and allows for consistent application of policies regardless of the physical location.
While AVI Load Balancer and Centralized Network Connectivity are valuable, they are often considered add- ons or outcomes rather than the core architectural features that define the multi-tenant, secure, and geographically distributed nature of a modern VCF private cloud modernization project.

NEW QUESTION: 10
An administrator has noticed an issue in a freshly deployed VMware Cloud Foundation (VCF) environment where the BGP neighborship between the Tier-0 gateway and a physical router remains in the Idle state. Pings between the uplink IPs are successful. What is the issue?
A. Autonomous System number mismatch.
B. Distributed Firewall blocking traffic.
C. Geneve tunnel down.
D. Overlay MTU too low.
Answer: A (LEAVE A REPLY)

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In the context ofVMware Cloud Foundation (VCF), particularly versions 5.x and the architectural advancements inVCF 9.0, the establishment of North-South routing via theNSX Tier-0 Gatewayis a critical post-deployment or bring-up task. The Tier-0 gateway usesBorder Gateway Protocol (BGP)to peer with physical Top-of-Rack (ToR) switches to exchange reachability information for the overlay networks.
When a BGP session is reported in the"Idle"state, it indicates that the BGP Finite State Machine (FSM) is at its first stage and is not yet attempting a TCP connection, or it has encountered an error that forced it back to this state. According to VMware VCF documentation and NSX troubleshooting guides, if the administrator can successfully ping between the Tier-0 uplink IP and the physical router interface,Layer 3 reachability is confirmed. This eliminates issues related to physical cabling, VLAN tagging on the trunk ports, or basic IP interface configuration.
The primary reason a BGP session remainsIdledespite successful ICMP reachability is a configuration mismatch. Specifically, anAutonomous System (AS) number mismatchis the most frequent culprit. BGP requires that the "Remote AS" configured on the Tier-0 gateway matches the "Local AS" of the physical peer.
If the SDDC Manager automated workflow or the manual configuration in NSX Manager contains a typo in these values, the protocol handshake will fail immediately.
While aDistributed Firewall (DFW)could technically block port 179, it is not common in a "freshly deployed" environment for the default rules to block the Edge Node's control plane traffic.Geneve tunnelsand MTU issues(Option C and D) typically affect the data plane-causing packet loss for encapsulated guest VM traffic-but they do not prevent the BGP control plane (running over standard TCP) from moving beyond the Idle state. Therefore, verifying the AS numbers in the VCF Planning and Preparation Workbook against the physical switch configuration is the verified resolution path.

NEW QUESTION: 11
An administrator has observed an NSX Local Manager (LM) outage at the secondary Site. However, the NSX Global Manager (GM) in secondary Site remains operational. What happens to data plane operations and policy enforcement at the secondary site?
A. All traffic is blocked until secondary site LM recovers.
B. Only local policies work; global policies cease to apply on the secondary site.
C. The data plane operates normally until LM recovery and reconnection.
D. Secondary site must failover all workloads to Primary site.
Answer: C (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
The architecture ofNSX Federationwithin a VCF Multi-Site design is built upon a separation of theControl Planeand theData Plane. This "decoupled" architecture ensures high availability and resiliency even when management components become unavailable.

In NSX Federation, theGlobal Manager (GM)handles the configuration of objects that span multiple locations, while theLocal Manager (LM)is responsible for pushing those configurations down to the local Transport Nodes (ESXi hosts and Edges) within its specific site. When a configuration is pushed, the Local Manager communicates with theCentral Control Plane (CCP)and subsequently theLocal Control Plane (LCP)on the hosts.
If an NSX Local Manager goes offline, the "Management Plane" for that site is lost. This means no new segments, routers, or firewall rules can be created or modified at that site. However, the existing configuration is already programmed into theData Plane(the kernels of the ESXi hosts and the DPDK process of the Edge nodes).
According to VMware's "NSX Multi-Location Design Guide," the data plane remains fully operational during a Management Plane outage. Existing VMs will continue to communicate, BGP sessions on the Edges will remain established, and Distributed Firewall (DFW) rules will continue to be enforced based on the last known good configuration state cached on the hosts. The data plane does not require constant heartbeats from the Local Manager to forward traffic. Therefore, operations continue normally "headless" until the LM is restored and can resume synchronization with the Global Manager and local hosts. Failover to a primary site (Option D) is only necessary if the actual data plane (hosts/storage) fails, not just the management components.

NEW QUESTION: 12
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 (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
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, theVCF 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 theVCF Installerto 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 ofSDDC 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 9beforeconvergence (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 tobring the environment in first, then let VCF manage the upgrades.

NEW QUESTION: 13
An administrator must prevent a new VPC from exporting any of its prefixes to the datacenter while still receiving a default route. Where should the routing policy be applied?
A. On the VPC default route advertiser
B. On the VPC's Transit Gateway
C. On the providers' BGP peer template
D. On the VPC Gateway Firewall
Answer: B (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In the advanced networking architecture ofVMware Cloud Foundation (VCF) 9.0and the evolution ofNSX VPCs, the control of route propagation is managed through the relationship between the consumer (the VPC) and the provider (the Tier-0 or Tier-1 Gateway). When a VPC is created, it is logically connected to the provider's infrastructure via aTransit Gateway(or a Provider-side logical router acting as a transit point).
To control the flow of routing information-specifically to prevent the data center's physical network from learning about internal VPC subnets (prefixes) while ensuring the VPC can still reach the outside world via a default route-the routing policy must be applied at the point of intersection.
TheTransit Gatewayserves as this demarcation point. By applying a route filter or prefix list on the Transit Gateway, the administrator can explicitly deny the advertisement of internal VPC prefixes "upstream" to the provider's BGP process.
Simultaneously, the provider can still inject or "advertise" a default route ($0.0.0.0/0$) "downstream" into the VPC.

Applying the policy on theVPC Gateway Firewall(Option D) would impact the data plane (blocking traffic) but would not prevent the routing table from being populated. TheBGP peer template(Option C) is too broad, as it would likely affect all VPCs connected to that provider, rather than just the "new VPC" in question. Thedefault route advertiser(Option A) only controls the egress of the default route, not the suppression of internal prefixes. Therefore, the Transit Gateway is the verified location for granular route control in a multi-tenant VCF VPC environment.

NEW QUESTION: 14
The administrator is working to ascertain the encapsulation of GENEVE by reviewing the capture on Wireshark.
The administrator instructed VM-1 to send a continuous ICMP request directed at VM-2.
Click to highlight where the administrator should observe the GENEVE encapsulated packet.

Answer Area

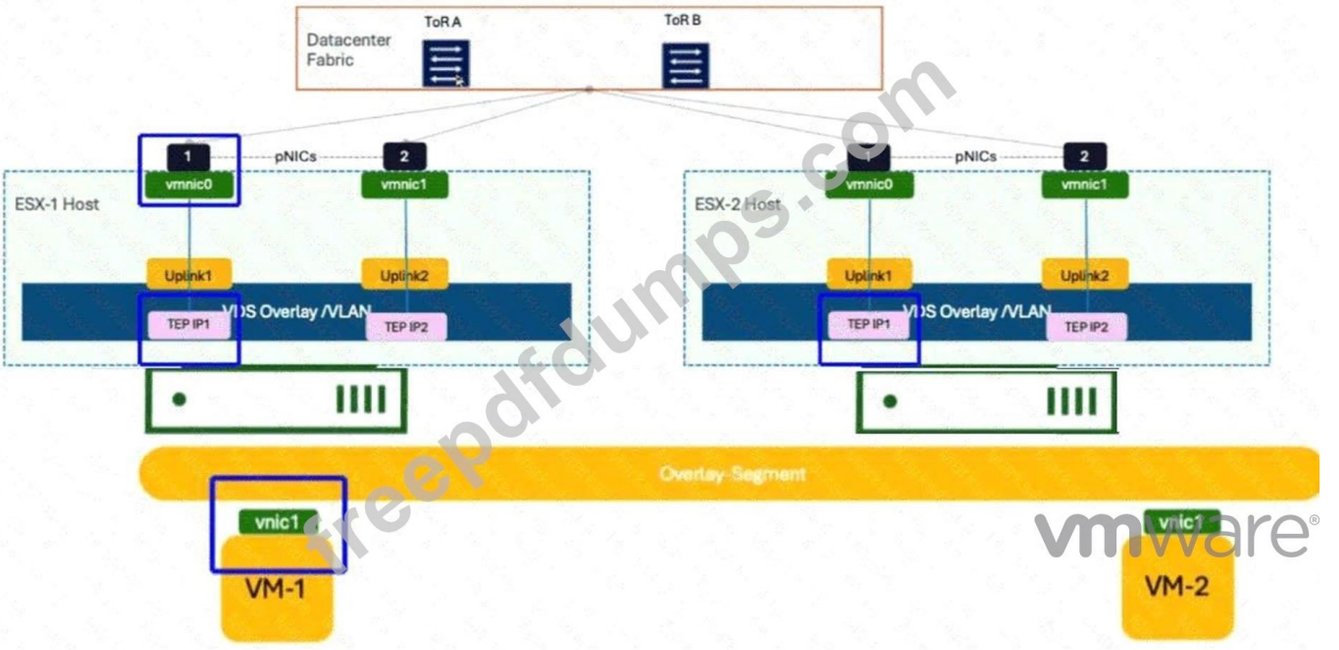

Answer:

Answer Area

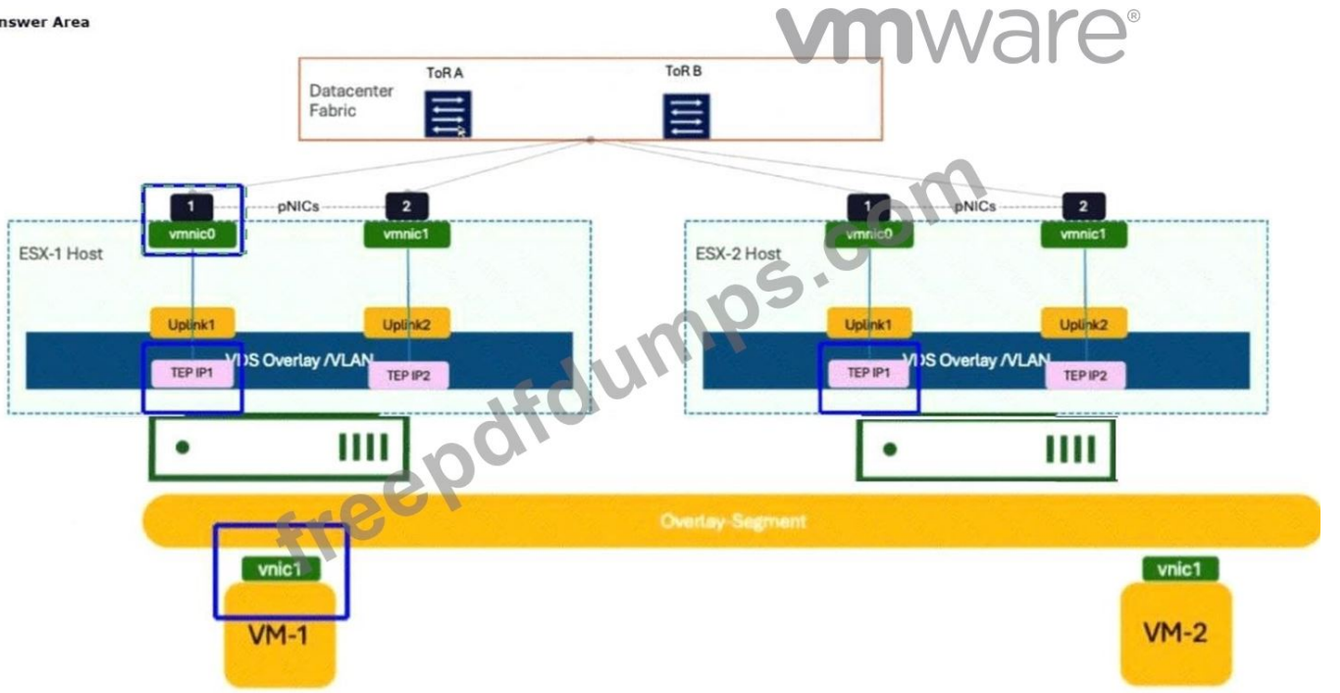

Explanation:
The administrator should click thevmnic0interface on theESX-1 Host.
In aVMware Cloud Foundation (VCF)environment, theGENEVE (Generic Network Virtualization Encapsulation)protocol is the industry-standard tunnel format used by NSX to create an overlay network.
This protocol allows Layer 2 traffic from virtual machines to be "tunneled" over a Layer 3 physical IP fabric, enabling workloads to communicate as if they were on the same segment even when separated by physical routers.
When VM-1 on ESX-1 sends an ICMP request to VM-2 on ESX-2, the packet starts as a standard Ethernet frame at the virtual machine'svnic1. At this stage, the packet contains no encapsulation.
As the frame enters theVirtual Distributed Switch (VDS)and hits theTunnel End Point (TEP), the host's kernel performs the encapsulation process. The TEP adds a GENEVE header, a UDP header (port 6081), and an outer IP header.
Thevmnic0(physical NIC) on the source host (ESX-1) is the specific "egress" point where this transformation is complete. A packet capture taken at this physical interface will show the "Outer IP" address of the source TEP and destination TEP, with the original ICMP packet hidden inside the GENEVE payload. If the administrator were to click on the VM's vnic, they would only see standard ICMP. By selecting thevmnic0, the administrator captures the traffic as it is placed onto the physical wire, which is the verified location to troubleshoot MTU issues, encapsulation errors, or physical fabric connectivity in a VCF environment.

NEW QUESTION: 15
An administrator must provide North/South connectivity for a VPC. The fabric exposes a distributed external VLAN across all ESX hosts. But, the only BGP peer to the core is on a VLAN only accessible on the Edge Cluster. Which design is required?
A. Use a VPC Tier-0 Gateway in active/active mode with distributed eBGP peering.

B. Distributed Transit Gateway with an EVPN route reflector on the transport nodes.
C. Centralized Transit Gateway on the Edge Cluster.
D. Deploy a Provider Tier-1 with BGP and connect the VPC Transit Gateway via route leaking.
Answer: C (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In aVMware Cloud Foundation (VCF)environment utilizing theVirtual Private Cloud (VPC)model, North
/South connectivity is managed by theTransit Gateway (TGW). The TGW acts as the bridge between the VPC-internal networks and the provider-level physical network.
The scenario presents a specific constraint: while an external VLAN exists across all hosts, the actual BGP peering point (the interface to the physical core routers) is restricted to theNSX Edge Cluster. In NSX terminology, when a gateway or service must be anchored to specific Edge Nodes to access physical network services-such as BGP peering, NAT, or stateful firewalls-it must be configured as aCentralizedcomponent.
ACentralized Transit Gateway(Option C) is instantiated on the Edge nodes. This allows the TGW to participate in the BGP session with the core routers on the VLAN that is only accessible to those Edges. The TGW then handles the routing for the VPC's internal segments. Traffic from the ESXi transport nodes (East- West) travels via the Geneve overlay to the Edge nodes, where it is then routed North-South by the Centralized TGW using the physical BGP peer.
Option A is incorrect because "distributed eBGP peering" would require every ESXi host to have peering capabilities, which contradicts the constraint. Option B involves EVPN, which is a significantly more complex and different architecture than what is required for standard VPC North/South access. Option D is an unnecessarily complex routing design that is not the standard VCF/VPC implementation pattern. Thus, the use of a Centralized Transit Gateway on the Edge cluster is the verified design requirement to bridge the gap between the overlay VPC and the localized BGP peering point.

NEW QUESTION: 16
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 identical Tier-0 gateway BGP autonomous system numbers.
B. All sites must be managed by the same VCF instance.
C. All Sites must use the same VTEP VLAN and IP pools.
D. All sites must have the same NSX version and build.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 17
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. Data Plane Development Kit (DPDK)
B. AMD Power Now
C. Non-Uniform Memory Access (NUMA)
D. Intel Speed Step
Answer: A (LEAVE A REPLY)
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: 18
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 theVMware Cloud Foundation (VCF) 9.0andNSX VPCarchitecture, theTransit 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 utilizeRoute Filteringat 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 theVPC 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 theTier-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 theprovider 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: 19
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 ofVMware 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 performframe 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 thesame 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 adifferent host(in an overlay environment), the MAC table entry for that remote MAC address will point to theTunnel 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 forwardingdecision.
TheMAC Tableis the authoritative source for Layer 2 forwarding. In an NSX-managed VCF environment, these tables are dynamically populated and synchronized via theLocal 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: 20
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:
InVMware Cloud Foundation (VCF)and NSX,Traceflowis 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 byinjecting synthetic trafficdirectly 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 theCentral Control Plane (CCP). The CCP then aggregates these observations and presents them in the NSX Manager UI.
ForVLAN-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: 21
An NSX Manager cluster has failed. The administrator deployed a new NSX Manager using the latest version and attempted to restore from a backup, but the restore operation failed. What would an administrator do to recover the cluster?
A. Edit the backup passphrase to match the new build.
B. Use SDDC Manager to replace NSX Manager.
C. Use the NSX restore API instead of the UI.
D. Deploy an NSX Manager that matches the backup's build.
Answer: D (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
A critical requirement for the backup and restore process inVMware NSX(and by extension, VCF) is version parity. The NSX Manager backup contains the database schema, configuration files, and state information specific to the version of the software that was running at the time the backup was taken.
When performing a restore into a "clean" environment, the NSX documentation explicitly states that the target NSX Manager appliancemust be of the exact same build versionas the appliance that generated the backup.

If an administrator attempts to restore a backup from version 4.1.x onto a newly deployed manager running version 4.2.x or 9.0 (as implies by "latest version"), the restore process will fail because the database schema of the newer version is incompatible with the older data structure. In aVCF environment, whileSDDC Manager(Option B) handles the lifecycle and replacement of failed nodes, the actual "Restore from Backup" workflow is an NSX-native operation. If the entire cluster is lost, the recovery procedure involves:

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

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

NEW QUESTION: 22

An administrator is responsible for a VMware Cloud Foundation (VCF) Private Cloud. The administrator has been tasked with identifying why there is no data ingress into a workload domain.

The workload domain has been configured with:

. A dedicated NSX Edge Cluster.
. A Tier 0 gateway.
. A Tier-1 gateway that is configured for Distributed Routing only.
. An NSX segment where a test virtual machine is located.

As part of the exercise, the administrator must map the traffic flow for data ingress into the workload domain to identify the steps that external network traffic will take to ingress into the workload domain and reach the virtual machine.

Drag and drop the six steps from the Steps list on the right and place them in order in the Solution Steps.

(Choose six.)

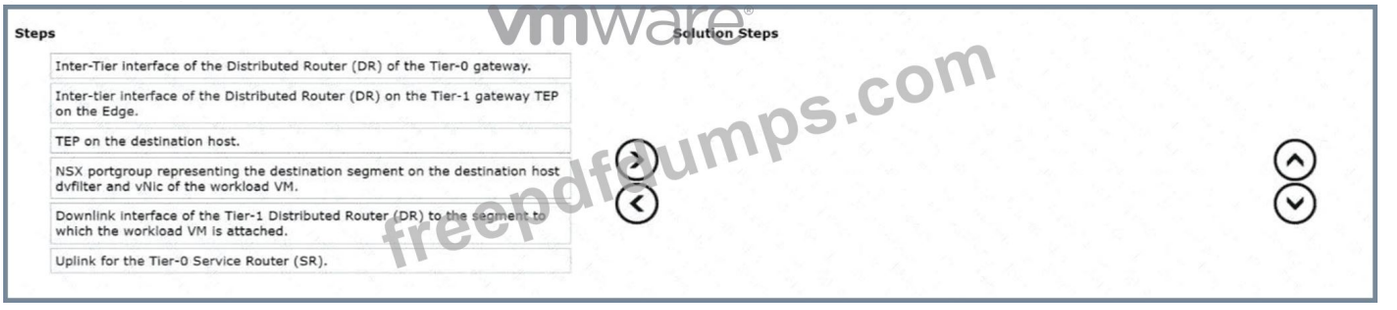

Answer:

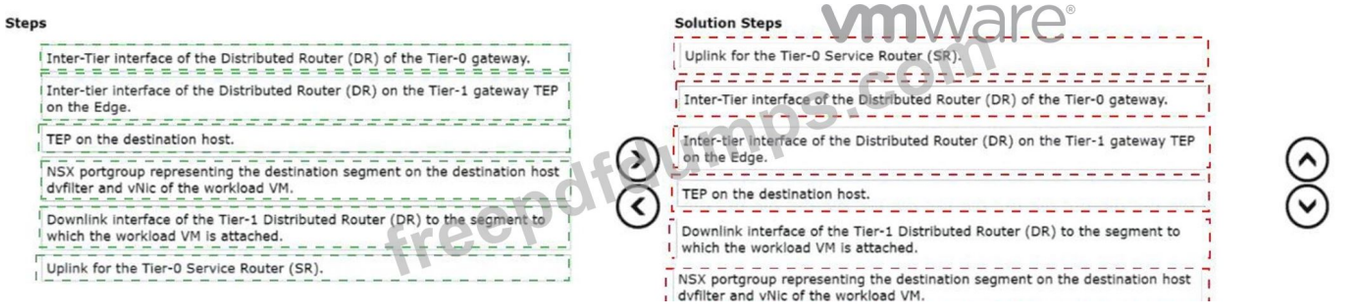

Explanation:
To identify why there is no data ingress into a workload domain, an administrator must understand the specific path external traffic takes. For a workload domain configured with a Tier-0 gateway and a Tier-1 gateway (Distributed Routing only), the ingress traffic flow follows a hierarchical path from the physical network through the NSX logical components to the virtual machine.
Ingress Traffic Flow Sequence
The correct sequence of steps for external network traffic to ingress the workload domain and reach the virtual machine is as follows:
* Uplink for the Tier-0 Service Router (SR): Traffic enters the NSX environment from the physical network through the physical-to-logical interface on the Edge node.
* Inter-Tier interface of the Distributed Router (DR) of the Tier-0 gateway: After being received by the Service Router, the packet is routed internally within the Tier-0 gateway to its distributed component.
* Inter-tier interface of the Distributed Router (DR) on the Tier-1 gateway TEP on the Edge: The Tier-0 gateway routes the packet to the Tier-1 gateway. In this specific scenario, since the Tier-1 is
"Distributed Routing only," this logical transition occurs on the Edge node participating in the transport zone.
* TEP on the destination host: The Edge node encapsulates the packet (typically via Geneve) and tunnels it across the physical fabric to the specific ESXi host where the target virtual machine is currently residing.
* Downlink interface of the Tier-1 Distributed Router (DR) to the segment to which the workload VM is attached: On the destination host, the packet is de-encapsulated. The local Tier-1 DR instance identifies the correct logical segment (VNI) for the destination IP.
* NSX portgroup representing the destination segment on the destination host dvfilter and vNic of the workload VM: The packet is delivered to the virtual switch port, passes through any applied Distributed Firewall (dvfilter) rules, and finally reaches the virtual machine's network interface card (vNIC).

NEW QUESTION: 23
An architect has just deployed a new NSX Edge cluster in a VMware Cloud Foundation (VCF) fleet. The BGP peer between the NSX Tier-0 gateway and the top-of-rack routers is successfully up and stable.

* BGP Connection is established, but the NSX Tier-0 is not receiving a default route from the top-of-rack routers.
* Workloads inside NSX have no Internet access.
What could be the solution?
A. Tier-0 gateway community settings are missing on the top-of-rack router configuration.
B. The top-of-rack router receives a default route from Tier-0 gateway.
C. Tier-0 gateway has a limit set too low for how many routes it can accept.
D. There is no default route configured on the top-of-rack router for the Tier-0 gateway.
Answer: D (LEAVE A REPLY)
Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:
In aVMware Cloud Foundation (VCF)deployment, establishing a stable BGP neighborship between the Tier-0 Gatewayand the physicalTop-of-Rack (ToR)switches is only the first step in enabling North-South connectivity. While the BGP state may show as "Established," this only confirms that the control plane handshake is complete and the peers are ready to exchange prefixes.
The primary reason for a lack of external connectivity in this scenario is that norouting informationis being shared. For workloads within the SDDC to reach the internet, the Tier-0 Gateway must have a path to external networks. In most enterprise VCF designs, the physical network (ToR) is expected to provide adefault route (0.0.0.0/0)to the Tier-0 Gateway.
If the Tier-0 is not receiving this route, the issue typically lies in the physical router's configuration. BGP does not automatically "originate" or "redistribute" a default route unless explicitly commanded to do so. On most physical network platforms (like Cisco, Arista, or Juniper), the administrator must specifically configure a
"default-originate" command or ensure a static default route exists in the physical RIB and is allowed to be advertised into the BGP session with the NSX Edge nodes.
Options A and C are unlikely to be the primary cause of a completely missing default route in a fresh deployment. Option B describes the inverse-where the virtual network tells the physical network how to find the internet-which is incorrect for a standard VCF consumer model.
Therefore, verifying and enabling the default route advertisement on the physical ToR switchesis the verified solution to provide the Tier-0 with the necessary egress path for internet-bound workload traffic.

NEW QUESTION: 24
An administrator is troubleshooting the packet flow of an incoming response to an ICMP Reply payload destined for 10.1.1.10 in the diagram.
The packet arrived at the Tier-0 SR at 172.16.215.100/29.
Which highlighted location identifies the next hop in the path to the destination?

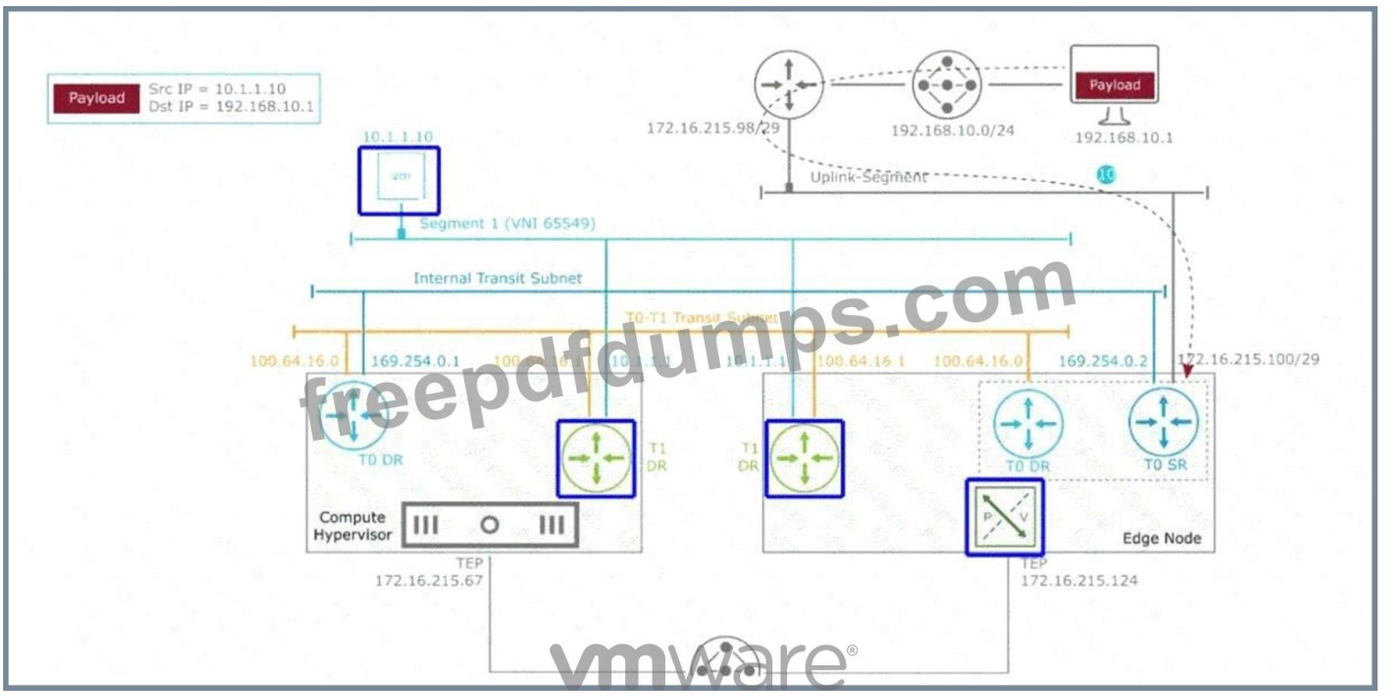

Answer:

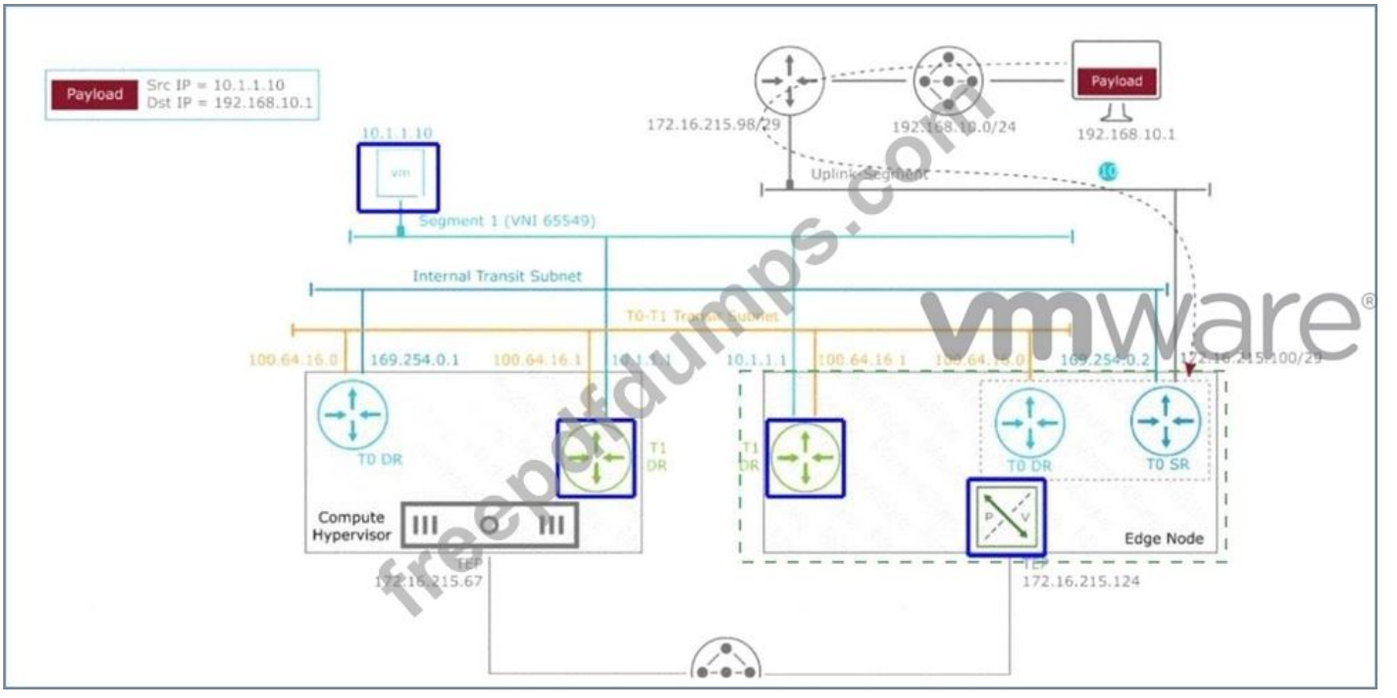

Explanation:

the administrator should click theTier-1 DR iconlocated within theEdge Node.

Comprehensive and Detailed 250 to 350 words of Explanation From VMware Cloud Foundation (VCF) documents:In aVMware Cloud Foundation (VCF)environment, North-South traffic flows through a hierarchical routing structure composed ofTier-0andTier-1 Gateways. Each gateway is further divided into a Distributed Router (DR)component, which runs as a kernel module on all Transport Nodes (ESXi and Edges), and aService Router (SR), which provides centralized services and resides on the Edge Nodes.

According to the packet walk logic for an incoming (North-to-South) packet, once the traffic arrives from the physical router at theTier-0 Service Router (SR)on the Edge Node, it must be routed toward the destination virtual machine (10.1.1.10). In a multi-tier NSX architecture, the

Tier-0 SR identifies that the destination subnet belongs to a connectedTier-1 Gateway. The communication between the Tier-0 and Tier-1 gateways occurs over an internal transit subnet, often referred to as theRouter Link(in this diagram, represented by the 100.64.16.0/31 subnet).
The "Next Hop" for the packet currently residing at the Tier-0 SR on the Edge Node is theTier-1 Distributed Router (DR)instance located on that same Edge Node. This is because the Edge Node participates as a Transport Node in the overlay and maintains local instances of all Distributed Routers to ensure efficient path processing. After the packet is processed by the local Tier-1 DR on the Edge Node, it determines that the destination VM is residing on a remote host (Compute Hypervisor). Only then is the packet encapsulated in a Geneveheader and sent via theTunnel Endpoints (TEP)from the Edge Node (172.16.215.124) to the Compute Hypervisor (172.16.215.67). Therefore, the Tier-1 DR on the Edge Node is the immediate logical next step in the routing pipeline before any host-to-host encapsulation occurs.

NEW QUESTION: 25
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 aVMware Cloud Foundation (VCF)environment, particularly with the high-performance requirements of North-South routing,BGPandBFD (Bidirectional Forwarding Detection)are used in tandem to ensure rapid failure detection. A common but subtle issue in fresh or modified environments is anMTU (Maximum Transmission Unit) mismatchon 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. TheseBGP Updatepackets 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 (typically9000 MTUfor 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.