Get Free VMware 3V0-42.23 Dumps PDF Questions

Why risk failure? Download updated VMware NSX 4.x Advanced Design exam PDF questions today. Practice with real 3V0-42.23 dumps and verified answers designed to help you ace your certification quickly using PrepBolt 3V0-42.23 exam pdf questions and answers.

Thank you for Downloading 3V0-42.23 exam PDF Demo

https://prepbolt.com/3V0-42.23.html

Question 1

Question Type: MultipleChoice

A Solutions Architect is helping an organization with the Physical Design of an NSX solution.

This information was gathered during the Assessment Phase:

There is a critical application used by the Finance Team.

The critical application has an availability and recoverability SLA of 99.999%.

The critical application is sensitive to network changes.

Which two selections should an architect include in their design? (Choose two.)

Options:

A- Configure Tier-0 gateway for eBGP and ECMP.
B- Enable BFD on Tier-0 gateway.
C- Install and configure hosts with 100Gbps physical NICs.
D- Configure multiple static routes on Tier-1 gateway.
E- Configure Tier-1 gateway for eBGP and ECMP.

Answer:

A, B

Explanation:

1. Ensuring High Availability for Critical Applications

For a 99.999% SLA, the NSX solution must ensure high availability (HA), redundancy, and failover mechanisms.

BGP with ECMP (Equal-Cost Multi-Path) enables multiple active paths for traffic forwarding, improving resiliency.

BFD (Bidirectional Forwarding Detection) ensures sub-second failure detection, minimizing downtime.

2. Why 'BGP with ECMP and BFD' is Correct (A, B)

(A - Configure Tier-0 for eBGP and ECMP)

ECMP allows multiple Tier-0 edges to be active, improving fault tolerance.

BGP dynamically advertises routes, ensuring efficient path selection.

(B - Enable BFD on Tier-0 Gateway)

BFD allows rapid failure detection (sub-second convergence) between NSX Edges and upstream routers.

Reduces packet loss and optimizes failover for North-South traffic.

3. Why Other Options are Incorrect

(C - Install Hosts with 100Gbps NICs):

While high-speed NICs improve performance, they do not ensure application availability.

(D - Configure Multiple Static Routes on Tier-1):

Static routes do not provide dynamic failover, making them unsuitable for high-availability designs.

(E - Configure eBGP on Tier-1):

BGP is typically used on Tier-0 for external routing, not Tier-1.

4. NSX Best Practices for High-Availability Applications

Use Active-Active Tier-0 Gateways with ECMP for redundancy.

Ensure BFD is enabled to provide real-time failure detection.

Implement distributed load balancing and failover testing.

VMware NSX 4.x Reference:

NSX-T BGP and ECMP Deployment Guide

NSX High Availability Design Best Practices

Question 2

Question Type: MultipleChoice

A large multinational company is expanding its data center due to increased demand for online services.

The company is considering shifting from an NSX Edge VM design to a bare-metal NSX Edge design to accommodate new hardware acquisitions and maximize performance.

Which is a potential benefit for the company in shifting from an NSX Edge VM design to a bare-metal NSX Edge design?

Options:

A- It will maximize performance by reducing virtualization overhead.
B- It will allow for the implementation of more VLANs.
C- It will automatically distribute stateful services across Edge nodes.
D- It will eliminate the need for stateful services.

Answer:

A

Explanation:

Performance Benefits of Bare-Metal NSX Edge (Correct Answer - A):

Bare-metal NSX Edge Nodes provide higher performance by eliminating the virtualization overhead associated with Edge VMs running inside ESXi/KVM hosts.

This increases throughput and reduces latency, making it ideal for high-bandwidth applications (e.g., Load Balancing, VPN, and NAT).

Incorrect Options:

(B - More VLANs):

The number of VLANs is not limited by the NSX Edge type. VLAN scalability depends on physical network design.

(C - Automatic Stateful Service Distribution):

Stateful services (NAT, FW, LB, VPN) do not auto-distribute. Stateful HA must be manually configured.

(D - Eliminates Stateful Services):

Stateful services (e.g., NAT, Load Balancer, Firewall) are still required, regardless of Edge deployment mode.

VMware NSX 4.x Reference:

VMware NSX-T Bare-Metal Edge Deployment Guide

NSX-T Edge Node Performance Optimization

Question 3

Question Type: MultipleChoice

Which of the following describes the role of the NSX Gateway Firewall as an inter-tenant firewall within a VMware NSX solution?

Options:

A- It secures communication between on-premises physical servers and virtual machines (VMs) in the cloud.
B- It inspects and filters traffic between virtual machines (VMs) within the same tenant.
C- It isolates different tenants' virtual networks, preventing unauthorized communication between them.
D- It controls access to virtual resources based on user identity and authentication.

Answer:

C

Explanation:

NSX Gateway Firewall for Multi-Tenancy (Correct Answer - C):

The NSX Gateway Firewall acts as an inter-tenant firewall by isolating different tenants' networks to prevent cross-tenant communication.

Ensures multi-tenancy security, per-tenant policy enforcement, and North-South traffic control.

Incorrect Options:

(A - Secures On-Prem to Cloud Communication):

This is handled by IPSec VPN, BGP, or NAT, not the Gateway Firewall.

(B - Filters Intra-Tenant Traffic):

Intra-tenant filtering is the responsibility of the NSX Distributed Firewall (DFW), not the Gateway Firewall.

(D - User-Based Access Control):

Identity-Based Firewall (IDFW) controls access based on user authentication, not network segmentation.

VMware NSX 4.x Reference:

NSX-T Multi-Tenancy and Security Isolation Best Practices

NSX Gateway Firewall Deployment Guide

Question 4

Question Type: MultipleChoice

A Network Architect has been tasked with recommending a solution for traffic management to a client. The client has asked about the differences between IP hash and LACP for link integration.

Which of the following is an accurate description of the differences?

Options:

A- IP hash uses a control protocol to negotiate link aggregation, while LACP uses a hash function to distribute traffic based on the source and destination IP addresses.
B- LACP uses a hash function to distribute traffic based on the source and destination IP addresses, while IP hash uses a control protocol to negotiate link aggregation.
C- IP hash uses a hash function to distribute traffic based on the source and destination IP addresses, while LACP uses a control protocol to negotiate link aggregation.
D- LACP uses a control protocol to negotiate link aggregation, while IP hash uses a hash function to distribute traffic based on the source and destination MAC addresses.

Answer:

C

Explanation:

1. Understanding Link Aggregation in NSX

IP Hash and LACP (Link Aggregation Control Protocol) are methods for link aggregation used in NSX-T networking.

Both techniques allow multiple physical links to be combined into a logical interface for higher bandwidth and redundancy.

2. Why 'IP Hash Uses a Hash Function, LACP Uses a Control Protocol' is Correct (C)

IP Hash:

Uses a hashing function to distribute traffic based on source and destination IP addresses.

It does not negotiate link aggregation dynamically.

LACP:

Uses a control protocol to dynamically negotiate and maintain aggregated links.

Automatically detects and manages failures in aggregated links.

3. Why Other Options are Incorrect

(A - IP Hash Uses Control Protocol):

IP Hash does not use a control protocol; it only applies a hash function.

(B - LACP Uses Hashing Instead of Control Protocol):

LACP does not use a hash function for traffic distribution; it uses a negotiation protocol.

(D - LACP Hashes MAC Instead of IP):

LACP does not perform hashing; it manages link aggregation dynamically.

4. NSX Best Practices for Link Aggregation

LACP is recommended for environments where dynamic link negotiation is required.

IP Hash is used in environments where static load balancing is preferred.

Ensure the correct uplink profile is assigned to NSX Transport Nodes for link aggregation.

VMware NSX 4.x Reference:

NSX-T Link Aggregation and NIC Teaming Best Practices

NSX-T Uplink Profile Design Guide

Question 5

Question Type: MultipleChoice

A customer has two sites and is looking to deploy NSX with stretched security. The customer wants to ensure that only authorized traffic can traverse the stretched security perimeter.

What is the VMware recommended approach for implementing micro-segmentation in this scenario?

Options:

A- Use Distributed Firewall rules to enforce micro-segmentation across the stretched security perimeter.
B- Use Service Composer policies to enforce micro-segmentation across the stretched security perimeter.
C- Use Identity Firewall policies to enforce micro-segmentation across the stretched security perimeter.
D- Use Group Firewall policies to enforce micro-segmentation across the stretched security perimeter.

Answer:

A

Explanation:

Micro-Segmentation Across Stretched Security (Correct Answer - A):

NSX Distributed Firewall (DFW) enforces security at the workload level across both sites.

DFW provides East-West traffic control, preventing unauthorized lateral movement.

Enforcement remains consistent across sites, maintaining Zero Trust Security.

Incorrect Options:

(B - Service Composer Policies):

Service Composer is deprecated in NSX-T and not used for micro-segmentation.

(C - Identity Firewalling):

Identity-Based Firewall (IDFW) applies user-based security, not network segmentation.

(D - Group Firewall Policies):

Group-based policies work with DFW, but DFW is the primary enforcement mechanism.

VMware NSX 4.x Reference:

NSX-T Micro-Segmentation Security Best Practices

Distributed Firewall Design Guide for Stretched Security

Question 6

Question Type: MultipleChoice

A Network Solutions Architect is tasked with designing an optimized and high-performing NSX solution, keeping in mind the need for DPU-based acceleration. The architect needs to consider the use of Geneve Offload, Receive Side Scaling (RSS), Geneve Rx Filters, SSL Offload, and the effects of Multi-TEP, MTU size, and NIC speed on throughput. Furthermore, the architect also needs to consider the key performance factors for compute nodes and NSX Edge nodes.

As the company's traffic continues to surge, there's a requirement to ensure NSX Edge nodes can handle the increasing load.

Which of the following factors should primarily be considered for performance optimization?

Options:

A- The NSX Edge VM node size
B- The available storage for the cluster
C- The number of ESXi hosts
D- The number of NSX Edge Node uplinks

Answer:

A

Explanation:

NSX Edge VM Node Size for Performance Optimization (Correct Answer - A):

NSX Edge VM size determines CPU, memory, and throughput capacity.

Larger Edge nodes (Large/Extra Large) support higher bandwidth, more services, and faster packet processing.

NSX Advanced Load Balancer and Firewall policies consume Edge CPU cycles, requiring proper sizing.

Incorrect Options:

(B - Available Storage):

Storage capacity does not directly impact NSX Edge performance.

(C - Number of ESXi Hosts):

More hosts improve NSX resiliency, but do not increase Edge performance.

(D - Number of NSX Edge Uplinks):

Multi-TEP and high-speed NICs improve performance, but Edge node size is the primary factor.

VMware NSX 4.x Reference:

NSX Edge Node Performance Optimization Guide

DPU-Based Acceleration Best Practices in NSX-T

Question 7

Question Type: MultipleChoice

Which of the following is a requirement for using NSX Federation for disaster recovery?

Options:

A- All sites must have the same NSX version and build.
B- All sites must have the same physical hardware.
C- All sites must be located in the same geographical region.
D- All sites must have the same IP address space.

Answer:

A

Explanation:

NSX Federation Requirements (Correct Answer - A):

NSX Federation allows managing multiple NSX-T Data Center instances centrally across multiple locations.

To ensure seamless disaster recovery, all sites must run the same NSX version and build to support:

Global Policies & Rules Consistency

Inter-Site Transport Zone Communication

Seamless Failover & Policy Replication

Incorrect Options:

(B - Same Physical Hardware Required):

NSX Federation does not require identical hardware. However, each site should meet the minimum hardware specifications for compatibility.

(C - Must Be in the Same Region):

Federation supports multi-region deployments, allowing disaster recovery across different geographical locations.

(D - Must Have the Same IP Address Space):

Each NSX site can have different IP address spaces, as NSX Federation supports routing between sites using Tier-0 Gateways and BGP.

VMware NSX 4.x Reference:

NSX-T Federation Deployment Guide

NSX-T Multi-Location Disaster Recovery Architecture

Question 8

Question Type: MultipleChoice

Which of the following should be taken into account when designing the uplink profile and transport node profile?

Options:

A- The type of network interface cards (NICs) used in the ESXi hosts.
B- The amount of available CPU and memory resources on the ESXi hosts.
C- The number of virtual machines running on each ESXi host.
D- The physical location of the ESXi hosts within the data center.

Answer:

A

Explanation:

NIC Type Selection for Uplink & Transport Node Profile (Correct Answer - A):

The performance and capacity of the physical NICs impact the overlay and VLAN transport traffic.

High-performance NICs (25G, 40G, 100G) enhance throughput and reduce latency.

DPU-based NICs (Data Processing Units) improve performance by offloading packet processing.

Incorrect Options:

(B - CPU & Memory Considerations):

While CPU/memory impact overall NSX performance, they do not determine uplink/transport profile design.

(C - Number of VMs Per Host):

VM density affects overlay traffic, but uplink profile design depends on NIC configuration.

(D - Physical Location of ESXi Hosts):

Location is important for high availability, but it does not directly define uplink profiles.

VMware NSX 4.x Reference:

NSX-T Uplink Profile & Transport Zone Design Guide

Question 9

Question Type: MultipleChoice

A rapidly growing e-commerce company, with a global customer base, is seeking to enhance their current network infrastructure to ensure a seamless and secure user experience. They have opted for VMware NSX to leverage software-defined networking (SDN) capabilities, and are particularly interested in employing NSX Edge to maximize their network performance.

A solutions architect is tasked with designing an effective and efficient solution using NSX Edge that meets the customer's requirements. The design should incorporate North-South routing to handle traffic to and from the internet.

To meet the company's requirements, what optimal solution should the solutions architect recommend, utilizing NSX Edge?

Options:

A- Deploy a single NSX Edge node for North-South routing services.
B- Deploy a single NSX Edge node and use a separate physical router for East-West routing.
C- Deploy multiple NSX Edge nodes and configure a Tier-0 Gateway in Active-Standby mode for North-South routing services.
D- Deploy multiple NSX Edge nodes and configure a Tier-0 Gateway in Active-Active mode for North-South routing services.

Answer:

D

Explanation:

1. Importance of NSX Edge for North-South Traffic

NSX Edge nodes provide routing, NAT, firewall, and load balancing services for North-South traffic (external connectivity).

Active-Active Tier-0 Gateway provides maximum performance and resiliency for high traffic volume.

2. Why Active-Active Tier-0 with Multiple Edge Nodes is the Best Choice (D)

Supports Equal-Cost Multi-Path (ECMP) routing, distributing North-South traffic across multiple paths.

Provides better scalability and performance than Active-Standby mode.

Ideal for high-volume applications like e-commerce sites that require low-latency, high-throughput connections.

3. Why Other Options are Incorrect

(A - Single NSX Edge Node):

Single Edge Nodes introduce a single point of failure.

(B - Using a Physical Router for East-West Routing):

NSX handles East-West traffic internally using Distributed Routing.

(C - Active-Standby Tier-0 Gateway):

Active-Standby mode does not provide load balancing across multiple nodes.

4. NSX Edge and Tier-0 Gateway Design Considerations

Ensure sufficient bandwidth allocation for North-South traffic.

Use BGP or OSPF for dynamic route advertisement.

Configure ECMP for efficient multi-path forwarding.

VMware NSX 4.x Reference:

NSX-T Edge Node Scaling and Performance Best Practices

Tier-0 Gateway Active-Active vs. Active-Standby Deployment Guide

Thank You for trying 3V0-42.23 PDF Demo

To try our 3V0-42.23 practice exam software visit link below

https://prepbolt.com/3V0-42.23.html

Start Your 3V0-42.23 Preparation

Use Coupon “SAVE50” for extra 50% discount on the purchase of Practice Test Software. Test your 3V0-42.23 preparation with actual exam questions.