PassQueen
a.k.a dumpsqueen.com

Advanced Design VMware NSX-T Data Center

VMware 3V0-42.20

Version Demo

Total Demo Questions: 7

Total Premium Questions: 79

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architectures and Technologies	6
Topic 2, VMware Products and Solutions	7
Topic 3, Planning and Designing	65
Topic 4, Performance-tuning, Optimization and Upgrades	1
Total	79

QUESTION NO: 1

An architect is helping an organization with the Logical Design of an NSX-T Data Center solution.

This information was gathered during the Assessment Phase:

Which selection should the architect include in their design? (Choose the best answer.)

A. Deploy a Tier-0 gateway in Active/Standby mode. Configure policy-based IPSec VPN with SHA512 with RSA as the hash algorithm.

B. Deploy a Tier-0 gateway in Active/Active mode. Configure route-based IPSec VPN with SHA512 with RSA as the hash algorithm.

C. Deploy a Tier-0 gateway in Active/Standby mode. Configure route-based IPSec VPN with SHA512 with RSA as the hash algorithm.

D. Deploy a Tier-0 gateway in Active/Active mode. Configure policy-based IPSec VPN with SHA512 with RSA as the hash algorithm.

ANSWER: C

Explanation:

Deploy a Tier-0 gateway in Active/Standby mode. Configure route-based IPSec VPN with SHA512 with RSA as the hash algorithm. This design aligns the gateway high-availability mode with the required stateful Tier-0 services, including IPSec VPN. An Active/Standby Tier-0 has one active service router at a time, allowing NSX to maintain stateful services and provide failover to the standby service router when required. This is the appropriate Tier-0 mode when the design requires VPN connectivity together with stateful gateway functionality.

Route-based IPSec VPN is also the suitable VPN model where protected traffic is selected through routing. NSX creates a tunnel interface for the VPN session, and traffic reaches the tunnel according to static routes or routes learned through dynamic routing. This approach supports a scalable logical design because networks can be advertised and withdrawn through the routing architecture rather than requiring a separate policy traffic selector for every protected subnet pair. SHA-512/RSA provides the specified cryptographic authentication and integrity choice for the VPN configuration. VMware documents Tier-0 high-availability modes and supported services in its Tier-0 gateway documentation, and describes route-based VPN routing behavior in its IPSec VPN documentation.

QUESTION NO: 2

An architect is designing the physical underlay for an NSX overlay environment. The organization requires a scalable Layer 3 fabric that provides equal-cost paths between transport nodes.

Which three selections should the architect recommend? (Choose three.)

A. Use a routed Layer 3 underlay.

B. Configure dynamic routing between the physical switches.

C. Enable ECMP in the physical network.

D. Extend a single Layer 2 VLAN across all racks for TEP connectivity.

E. Disable routing between top-of-rack switches.

F. Use spanning tree convergence as the primary path-failover mechanism.

ANSWER: A B C

Explanation:

Use a routed Layer 3 underlay, configure dynamic routing between the physical switches, and enable ECMP in the physical network. A routed underlay provides IP reachability between NSX transport-node tunnel endpoints without extending Layer 2

domains across the data center. Dynamic routing allows the fabric to learn and withdraw paths automatically as links or devices change state.

ECMP permits the physical fabric to use multiple equal-cost routed paths between tunnel endpoints. This provides aggregate bandwidth and resilient forwarding for overlay traffic. The underlay must also provide end-to-end MTU support for Geneve encapsulation and should be designed with redundant paths. See the VMware NSX Installation Guide.

QUESTION NO: 3

A customer deploying NSX-T Data Center requires role based access controls be enforced in NSX Manager with these requirements:

Which identity deployments would meet the customer's requirements? (Choose the best answer.)

A. NSX Manager OAuth 2.0 registered to a 3-node Active Directory Federation Services cluster.

B. NSX Manager OAuth 2.0 registered to a 2-node VMware Identity Manager cluster.

C. NSX Manager OAuth 2.0 registered to a 2-node Active Directory Federation Services cluster

D. NSX Manager OAuth 2.0 registered to a 3-node VMware Identity Manager cluster.

ANSWER: D

Explanation:

NSX Manager OAuth 2.0 registered to a 3-node VMware Identity Manager cluster is the appropriate identity deployment for centralized authentication and role-based access control in this NSX-T Data Center design. VMware Identity Manager (now commonly associated with Workspace ONE Access) provides the supported identity-provider integration model for NSX Manager in this version. It can authenticate users against enterprise directories and expose users and groups for assignment to NSX roles. NSX administrators can then grant built-in NSX privileges to the appropriate synchronized identity groups, rather than creating and managing separate local NSX accounts for each administrator.

A three-node VMware Identity Manager deployment also provides the resilient, highly available architecture expected for a management-plane service that controls administrative access to NSX. The nodes operate behind a load balancer and avoid making NSX access dependent on a single identity appliance. This aligns identity-service availability with the availability expectations of the NSX management environment while allowing consistent group-based authorization.

VMware's Workspace ONE Access installation guidance describes clustered deployments and load-balancer considerations at VMware Workspace ONE Access documentation. NSX administrative and authorization concepts are documented in the NSX-T Data Center Administration Guide.

QUESTION NO: 4

An architect is helping an organization with the Physical Design of an NSX-T Data Center solution.This information was gathered during a workshop about ESXi Host networking:

A total of 50 ESXi hosts to be configured as Transport Nodes.
All ESXi hosts have a dedicated 2 × Intel 10Gbps Physical Network adapter for the Overlay Traffic.

To achieve low latency, high throughput, redundancy, and performance, which two NIC teaming policies should the architect recommend? (Choose two.)

A. Load Balance Source MAC

B. Load Balance Port ID

C. Load Balance Source

D. Load Balance Source Port ID

E. Failover Order

ANSWER: A C

Explanation:

Load Balance Source MAC and **Load Balance Source** are the appropriate NSX-T uplink teaming policies for dedicated overlay uplinks where both throughput and resiliency are required. Load Balance Source uses the transport node's interface assignment to pin traffic to active uplinks. This enables multiple uplinks to be actively used, allowing the two dedicated 10-Gbps adapters to carry overlay traffic concurrently while retaining redundancy if an uplink becomes unavailable. It is especially suitable when uplink utilization can be planned through explicit transport-node interface-to-uplink mappings.

Load Balance Source MAC selects an active uplink by hashing the source Ethernet MAC address. Because different source MAC addresses can map to different uplinks, traffic can be distributed across the available physical adapters while preserving a consistent uplink selection for a given source. With two active physical NICs, this provides active-active use of the overlay network capacity and continued connectivity through the surviving uplink during a failure. These policies are supported teaming modes for NSX transport-node uplink profiles and align with the requirement for high throughput and redundant overlay connectivity. See the VMware NSX uplink profiles documentation and the NSX transport-node networking guidance.

QUESTION NO: 5

An architect is helping an organization with the Logical Design of an NSX-T Data Center solution.

This information was gathered during the Assessment Phase:

Which selection should the architect include in their design? (Choose the best answer.)

A. Review average North/South traffic from the core switches and firewall.

B. Include a segment QoS profile and review the impact of utilizing this feature.

C. Meet with the organization's application team to get additional information.

D. Monitor East-West traffic throughout normal business cycles.

ANSWER: B

Explanation:

Include a segment QoS profile and review the impact of utilizing this feature is the appropriate logical-design selection when assessment findings identify application flows that need differentiated treatment, such as guaranteed bandwidth, controlled latency or jitter, or protection from packet loss during congestion. NSX-T Quality of Service is applied through QoS profiles associated with segments, enabling the design to define traffic handling policy close to the logical network where the affected workloads connect. A segment QoS profile can specify shaping and bandwidth behavior, allowing the architect to translate an application's service requirements into an enforceable NSX-T logical-network construct.

The logical design should also explicitly review the operational and capacity impact of the selected QoS settings. QoS does not create additional physical bandwidth; rather, it manages how available bandwidth is allocated and how traffic is treated when limits or contention occur. Consequently, documenting the profile, affected segments, expected traffic classes, and validation criteria makes the design implementable and testable. VMware's NSX-T documentation describes QoS profiles and their application to segments in the NSX-T Quality of Service documentation. Additional NSX-T design guidance is available from the VMware NSX Reference Design Guide.

QUESTION NO: 6

An architect is creating a logical design for distributed firewall policy. The customer wants firewall rules to remain effective when application virtual machines are moved or when their IP addresses change.

Which three criteria can the architect use to create dynamic NSX groups? (Choose three.)

A. VM tags

B. Virtual machine names

C. IP addresses

D. Physical switch port numbers

E. Edge node form factors

F. BGP autonomous system numbers

ANSWER: A B C

Explanation:

VM tags, **virtual machine names**, and **IP addresses** can be used as membership criteria for NSX groups. Groups provide reusable policy objects that can be referenced as sources and destinations in distributed firewall rules. Dynamic membership enables policy to follow workload identity rather than requiring administrators to update individual firewall rules whenever a workload changes.

Tag-based grouping is commonly preferred for application-centric designs because tags can represent application, environment, ownership, or compliance classifications. IP-based criteria can be appropriate for workloads that are not managed through vCenter inventory, while VM-name criteria can support environments that use established naming standards. Physical switch port numbers, Edge node form factors, and BGP autonomous system numbers are not NSX group membership criteria. See the VMware NSX technical documentation.

QUESTION NO: 7

A customer has a requirement to implement a next generation firewall (NGFW) to improve security network introspection. The customer wants to apply the NGFW to all workloads exposed both internally and externally. The customer wants the NGFW to work seamlessly with NSX-T Data Center and vSphere.

Which solution should be recommended to the customer? (Choose the best answer.)

A. Use network introspection only on the external workloads and use NSX DFW for internal workloads.

B. Apply the NGFW on bare metal hosts which will offer better performance of inline network introspection.

C. Apply the NGFW to internal and external workloads for increased protection and use NSX-T Data Center with Federation to set network policies.

D. Use NSX-T Data Center leveraged with NSX Intelligence to protect all workloads at the network inspection level.

E. Deploy an NSX Partner NGFW through NSX service insertion for both east-west and north-south workload traffic.

ANSWER: E

Explanation:

Deploy an NSX Partner NGFW through NSX service insertion for both east-west and north-south workload traffic. This approach directly addresses the requirement for next-generation firewall inspection across all workloads, regardless of whether their communication is internal or externally exposed. NSX service insertion provides the integration framework through which certified partner security services can be deployed and consumed in an NSX environment. It enables security policy to steer applicable traffic to the NGFW service while retaining NSX-T Data Center's centralized networking, policy, and vSphere operational model.

This design is appropriate because the NGFW supplies advanced inspection capabilities, such as application-aware controls and threat prevention, while NSX provides distributed enforcement context and automated policy consumption for virtualized workloads. The integration is designed to make partner services operationally compatible with NSX and vSphere rather than treating the firewall as an isolated, manually managed appliance. Service insertion can be used to apply the security service where policy requires it, supporting consistent protection for workload traffic flows. VMware describes the service-insertion architecture and partner-service integration model in its NSX Service Insertion overview and provides further NSX security design guidance in the NSX Security Reference Architecture.