

DUMPSBOSS.

VMware NSX 4.x Advanced Design

VMware 3V0-42.23

Version Demo

Total Demo Questions: 5

Total Premium Questions: 52

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, VMware NSX Architecture and Technologies	10
Topic 2, VMware NSX Design	42
Total	52

QUESTION NO: 1

A customer has a physical VLAN-based network for a group of legacy appliances. The appliances cannot be migrated immediately, but virtual machines managed by NSX must connect to the same Layer 2 network. The customer does not require an overlay network for this workload group.

Which design should the Solutions Architect recommend?

- A. Create a VLAN-backed segment mapped to the existing physical VLAN.
- B. Create an overlay segment and require the legacy appliances to become transport nodes.
- C. Deploy a Tier-1 gateway without creating a segment.
- D. Deploy Layer 2 bridging between two VLAN-backed segments.

ANSWER: A

Explanation:

A VLAN-backed segment provides the appropriate connection for NSX-managed workloads that must attach directly to an existing physical VLAN. The segment maps to the VLAN transport network and allows the virtual machines and legacy appliances to remain in the same Layer 2 broadcast domain without introducing an overlay requirement for that workload group.

An overlay segment is appropriate when the design requires a Geneve-based logical network independent of a physical VLAN. A Tier-1 gateway provides Layer 3 connectivity but does not by itself connect workloads into the existing physical Layer 2 VLAN. Layer 2 bridging is useful when an overlay segment must be extended to a VLAN network, but it adds a bridge design where a direct VLAN-backed segment satisfies the stated requirement.

QUESTION NO: 2

What is the function of the data plane in NSX?

- A. It controls the behavior of the NSX environment.
- B. It manages the data traffic in the NSX environment.
- C. It provides the NSX APIs for automation and integration.
- D. It handles the configuration of the NSX environment.

ANSWER: B

Explanation:

It manages the data traffic in the NSX environment. In NSX, the data plane is the component that performs the actual packet processing and forwarding for workload traffic. It operates on transport nodes, including ESXi hosts and supported bare-metal or KVM hosts, where it applies the forwarding state supplied by NSX control-plane services. This enables traffic to be switched and routed across logical networks, including overlay segments that use Geneve encapsulation between tunnel endpoints.

The data plane also enforces network and security services directly in the traffic path. For example, distributed firewall rules are enforced close to the source workload, which allows NSX to apply policy to east-west traffic without requiring that traffic first traverse a centralized physical appliance. The plane therefore turns the intended connectivity and security policy into packet-level handling: forwarding permitted flows according to the relevant logical switching and routing state, while applying configured security enforcement as traffic is processed. VMware's NSX documentation describes NSX as separating management, control, and data-plane responsibilities, with transport nodes providing the infrastructure that carries and processes workload network traffic. See the [VMware NSX documentation portal](#) and the [NSX architecture documentation](#).

QUESTION NO: 3

A global bank has decided to overhaul its network infrastructure and adopt VMware NSX to enhance security and streamline management. The bank handles sensitive financial data and has a massive customer base, making it a potential target for cyber threats. Therefore, security is of paramount importance in this project.

A Network Solutions Architect is tasked with developing an NSX security design that incorporates security policy methodologies and adheres to NSX security best practices. They must ensure the micro-segmentation of network components, implement distributed firewalling, and create security policies that align with the bank's data protection requirements.

When considering NSX security VMware practices for the bank's deployment, what aspect is essential for enhancing the security posture?

- A. Avoid the use of distributed firewalls as they can complicate the network design.
- B. Implement a Zero Trust model and enforce policies at the Gateway level.
- C. Implement a Zero Trust model and enforce policies at the workload level.
- D. Deploy NSX in a single, large segment for simplicity.

ANSWER: C

Explanation:

Implement a Zero Trust model and enforce policies at the workload level is essential because it applies security controls as close as possible to the protected application, virtual machine, or workload interface. NSX Distributed Firewall enables this approach by enforcing stateful firewall policy in the hypervisor kernel, rather than relying solely on a centralized network chokepoint. Policies can be defined using workload context such as groups, tags, and identities, allowing the bank to express least-privilege communication requirements that align with application and data-protection boundaries.

For a financial institution, workload-level enforcement substantially limits the blast radius of a compromise. Even when an attacker gains access to one workload, explicitly permitted flows are the only paths available to other tiers, services, or sensitive-data systems. This directly supports micro-segmentation and reduces opportunities for unauthorized east-west movement. A Zero Trust design also supports consistent policy as workloads are provisioned, moved, or scaled, which is particularly important in a large environment with changing business services. VMware describes micro-segmentation as a core capability for applying granular security controls between workloads; see the [VMware NSX micro-segmentation guidance](#) and the [NSX security administration documentation](#).

QUESTION NO: 4

A customer requires north-south connectivity to remain available after the loss of one ESXi host or one top-of-rack switch. The customer will deploy an NSX Edge cluster for Tier-0 gateway services.

Which two design choices should the Solutions Architect recommend? (Choose two.)

- A. Deploy the Edge nodes on separate ESXi hosts.
- B. Connect Edge-node uplinks through separate top-of-rack switches.
- C. Deploy all Edge nodes on the same ESXi host to reduce east-west latency.
- D. Connect every Edge-node uplink to a single top-of-rack switch for consistent routing.
- E. Create additional Tier-1 gateways for each Edge node.

ANSWER: A B

Explanation:

NSX Edge nodes that provide a highly available gateway service should be deployed on separate ESXi hosts so that a single host failure does not remove all Edge datapath capacity. The hosts should also use physical uplinks connected through separate top-of-rack switches or equivalent independent physical failure domains. This avoids a single switch failure isolating every Edge node from the physical network.

Deploying all Edge nodes on one host creates a clear single point of failure. Connecting all Edge uplinks to one top-of-rack switch similarly leaves the design vulnerable to a switch or upstream-path failure. Adding more Tier-1 gateways does not provide resiliency for the Tier-0 Edge datapath or its physical-network connectivity.

QUESTION NO: 5

A Solutions Architect is validating the physical underlay for a new NSX deployment. The customer expects high east-west traffic volumes and requires resilient IP connectivity between all transport nodes.

Which two underlay design principles should the architect recommend? (Choose two.)

- A. Use a routed leaf-and-spine topology with ECMP-capable paths.
- B. Configure a consistent end-to-end MTU for Geneve-encapsulated traffic.
- C. Use a Layer 2 spanning-tree topology with intentionally blocked inter-leaf links.
- D. Configure transport-node reachability through a single physical default gateway.
- E. Terminate all compute-node traffic directly on the NSX Manager cluster.

ANSWER: A B

Explanation:

A routed leaf-and-spine underlay provides scalable IP reachability between transport-node tunnel endpoints and supports multiple equal-cost paths through the fabric. ECMP allows traffic to use available paths and provides resiliency when a link or fabric component fails.

The underlay must also use a consistent MTU that accommodates encapsulated Geneve traffic across the complete path. A Layer 2 spanning-tree design with blocked links does not provide the same predictable ECMP forwarding model for a modern NSX transport fabric. Configuring a default gateway only on NSX Edge nodes does not establish the required IP reachability between all compute transport nodes.