3V0-41.19Q&As

Advanced Design NSX-T Data Center 2.4

Pass VMware 3V0-41.19 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/3v0-41-19.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by VMware Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

Refer to Exhibit.

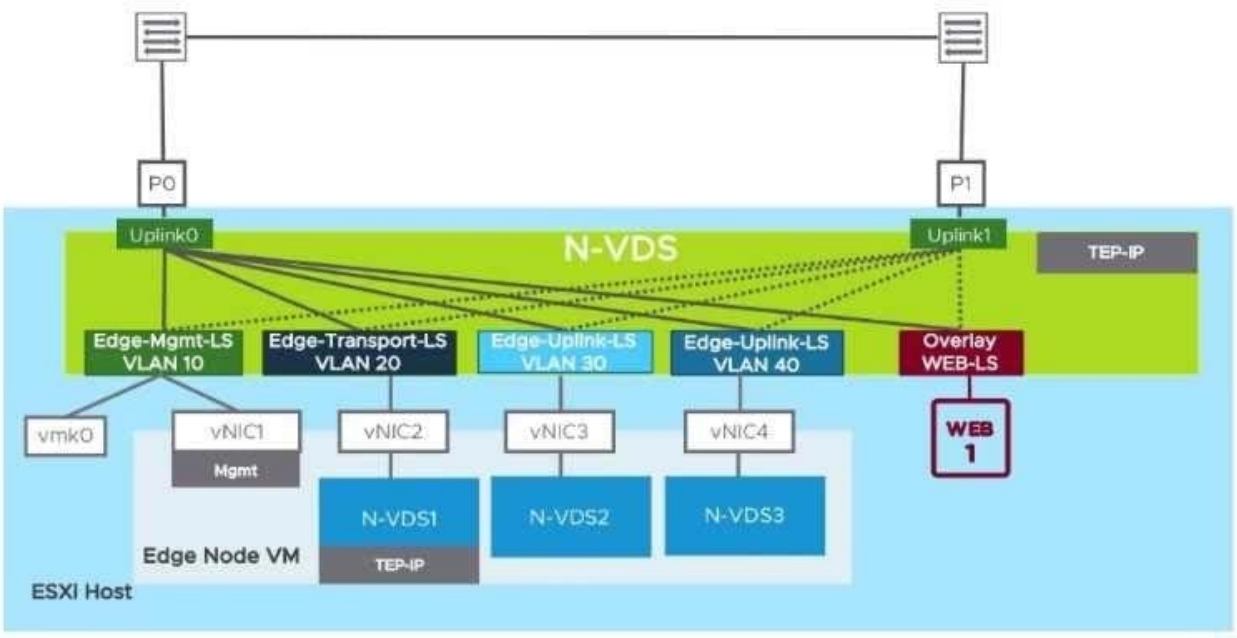

To meet the technical requirements for NSX Edge VM, which two design choices are required to satisfy this architectural design. (Choose two.)

A. NSX Edge TEP and ESXi TEP need to be in different VLANs.

B. ESXi host should be prepared as a Transport Node and use VLAN backend segments to connect Edge Node Interfaces.

C. ESXi host must have more than 2 pNICs available to create another N-VDS. D NSX Edge should run as a physical device.

D. vmk ports need to be on VDS instead of N-VDS, with onepNIC for each virtual switch providing greater functionality.

Correct Answer: AB

I believe this was supposed to have 5 answers as (C) looks like it has two answers on the same line unless it is saying "ESXi host must have more than 2pNICs available to create another N-VDS or NSX Edge should run as a physical device". Either one of those statements is still incorrect based on that diagram though.

(D)

is wrong as 1 pNIC per vSwitch is a bad design.

(C)

is wrong because you can do a 2 pNIC design with NSX-T and an Edge VM running on a N-VDS

https://vxplanet.com/2019/07/08/deploying-and-configuring-nsx-t-edges-on-n-vds-networking/

QUESTION 2

Refer to exhibit: An NSX architect is creating a Greenfield NSX-T Data Center solution using IPv6 addressing. This solution will form the starting point for a migration away from IPv4 addressing in the data center.

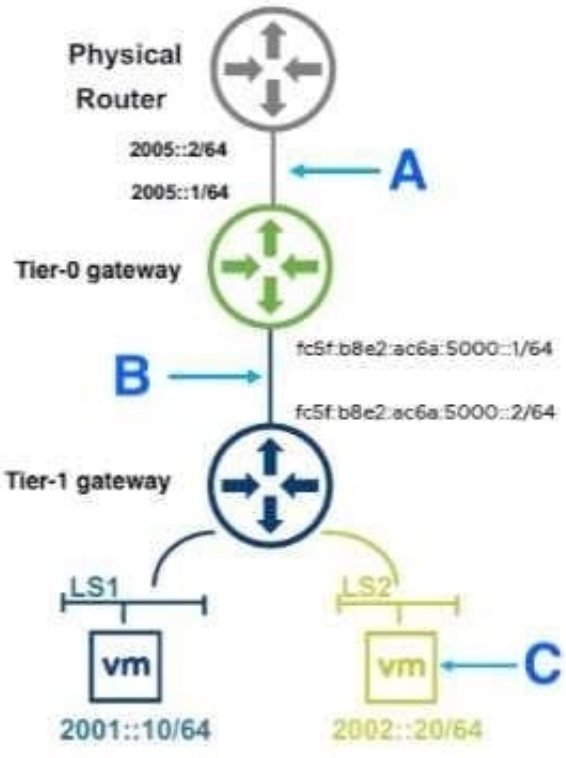

What are three correct labels for locations A, B, and C in the exhibit? (Choose three.)

A. Static IPv4 Addresses

B. Auto assigned from fc5f:b8e2:ac6a::/48 Unique Local

C. Static IPv6 addresses

D. Auto Assigned from fe:::::/48 Unique Local

E. DHCP relay for IPv4

F. Static IPv6 addresses and DHCP Relay

Correct Answer: BCF

https://blogs.vmware.com/networkvirtualization/2019/02/ipv6-support-in-nsx-t-2-4.html/

QUESTION 3

An architect is helping an organization with the Logical Design of a Layer 2 bridging solution. This information was gathered during the Assessment Phase:

1.

Workloads are running on ESXI hosts.

2.

Workloads are running on KVM hosts.

3.

Workloads on both type of hypervisors should use bridging services.

4.

VLAN 50 is used for Tier-0 uplink connectivity.

Which should the architect include in their design?

A. Create an NSX Edge Bridge Cluster and configure the bridging profile with VLAN 60.

B. Create an ESXi Bridge Cluster and configure the bridging profile with VLAN 60.

C. Create an NSX Edge Bridge Cluster and configure the bridging profile with VLAN 50.

D. Create an ESXi Bridge Cluster and configure the bridging profile with VLAN 50.

Correct Answer: C

https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.3/com.vmware.nsxt.admin.doc/GUID-E57A4794-93BF-4E1C-B5D2-23C575C00EEC.html VLAN 50 is used in the example -Given that along with required support for ESXi and KVM, and given that KVM is not supported on ESXi Bridge Cluster, C would be the correct answer https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.3/ com.vmware.nsxt.admin.doc/GUID- 7B21DF3D-C9DB-4C10-A32F-B16642266538.html--vetted You can configure layer 2 bridging using either ESXi host transport nodes or NSX Edge transport nodes. Edge bridging is preferred over ESXi bridging.

QUESTION 4

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:

1.

Maximum performance and availability is required between the physical and virtual network.

2.

Load Balancing service is required for back-end web servers.

3.

NAT is required.

Which should the architect include in their design?

A. Deploy a Tler-1 gateway and connect It to an Active/Active Tier-0 gateway with ECMP configured.

B. Deploy an Active/Active Tier-0 gateway and configure ECMP.

C. Create two separate VLANs to connect the Tier-0 gateway upstream traffic and configure ECMP.

D. Deploy an Active/Passive Tier-0 gateway and configure ECMP.

Correct Answer: A

Option A is required (even though BandC are technically correct for parts of the requirement).

Stateful services (LB) can\\'t be on the same gateway as ECMPgateway.

https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.4/administration/GUID-DAEF8457-83634F33-84DA-68AA36A2DE3C.html

https://vnuggets.com/2019/09/13/nsx-t-inline-and-onearm-load-balancing-part1/ https://

nsx.techzone.vmware.com/resource/vmware-nsx-t-design-guide-designing-environments-nsx-t

QUESTION 5

Refer to Exhibit:

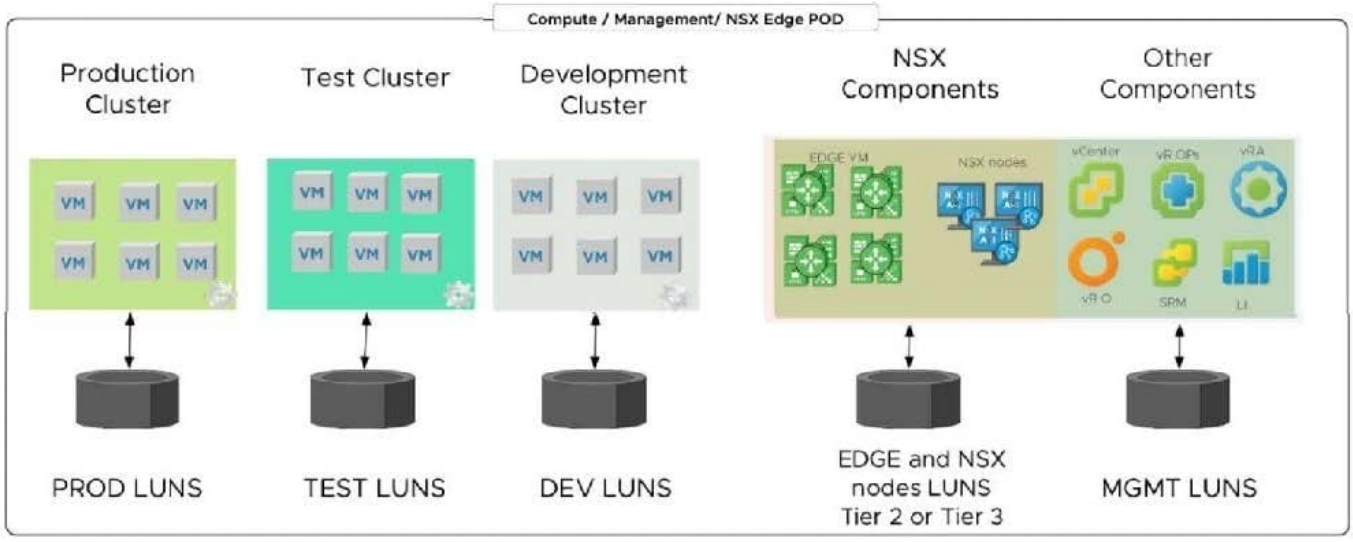

An NSX-T architect has been asked to review and recommend improvements for an NSX-T Data Center Logical Design, as shown in the drawing. The design must allow workload bursts for tenants to and from the public cloud and accommodate 30% yearly growth.

What two VMware recommended changes will Improve the Logical design? (Choose two.)

A. A separate POD is required for the NSX Edge nodes since the amount of traffic will be heavy.

B. An additional POD will be required to pivot workloads to Public Cloud.

C. Automation tools will be required to reduce time for workloads to be vMotioned.

D. Load balancers should be added to the design to support bursts from the Public Cloud.

E. NSX-T Datacenter components needs to be placed on the Public Cloud for cost reduction.

Correct Answer: CD

You aren\\'t placing NSX-T components in the cloud so (E) is wrong. It talks about bursting "to and from" the cloud, which lends itself to possibly being a VMware HCX (automation tools) play for (C) (A) With a whole separate "POD" (covering everything in the graphic) based on this logical design would be overkilled for the NSX Edges (B) no additional pods are required for pivoting/moving workloads to the public cloud