



- Expert Verified, Online, **Free**.

Which is a family of solutions for data center designs that span compute, storage, networking, and management, serving as a blueprint for a customer's Software Defined Data Center (SDDC) implementations? (Choose the best answer.)

- A. VMware SDDC Design
- B. VMware Validated Design
- C. VMware POC Design
- D. VMware Cloud Foundation

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Revant** Highly Voted 👍 2 years, 12 months ago

Correct Answer: B
upvoted 5 times

🗨️ 👤 **redtop** Most Recent 🕒 5 months, 1 week ago

Selected Answer: B
<https://docs.vmware.com/en/VMware-Validated-Design/index.html>
upvoted 1 times

🗨️ 👤 **[Removed]** 1 year ago

This is certainly going to be VCF as VVD is a standard for VMware Validated Designs (Architecture Standards) which includes more than what's listed here. While VCF is a product stack that includes the entire SDDC with Automation.
upvoted 1 times

🗨️ 👤 **Alchot** 1 year, 4 months ago

B is correct.
For NSX-T 3.0 the correct answer is VMware Validated Design. In 2022 this changed from VVD to VCF.
upvoted 1 times

🗨️ 👤 **PSE_IT** 1 year, 5 months ago

Selected Answer: B
Correct Answer: B
upvoted 1 times

🗨️ 👤 **jurajJN** 1 year, 6 months ago

Correct Answer: B
upvoted 1 times

🗨️ 👤 **Me_Loi1** 2 years, 9 months ago

Correct Answer: B
upvoted 3 times

🗨️ 👤 **VMwareARCHI** 2 years, 11 months ago

Correct Answer: B
upvoted 3 times

Which three IPv6 features are supported in an NSX-T Data Center design? (Choose three.)

- A. IPv6 OSPF
- B. IPv6 static routing
- C. IPv6 switch security
- D. IPv6 DNS
- E. IPv6 Distributed Firewall
- F. IPv6 VXLAN

Suggested Answer: BCE

Reference:

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

 **tungdt** Highly Voted 2 years, 11 months ago


BCE is correct

upvoted 6 times

 **vRockStar** Most Recent 4 months, 1 week ago

now IPV6 TEP is also supported !! but its Geneve not VxLAN so BCE is still valid


upvoted 1 times

 **Alchot** 1 year, 4 months ago

BCE are correct. Some details in this link with added on version 2.4 onwards

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

upvoted 1 times

 **Me_Loi1** 2 years, 9 months ago

BCE is correct

upvoted 3 times

 **NetRock1** 2 years, 11 months ago

Agree, BCE is correct!

upvoted 3 times

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

This information was gathered during a workshop:

- ⇒ Some workloads should be moved to a Cloud Provider.
- ⇒ Extend network's VLAN or VNI across sites on the same broadcast domain.
- ⇒ Enable VM mobility use cases such as migration and disaster recovery without IP address changes.
- ⇒ Support 1500 byte MTU between sites.

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

- A. Load Balancer
- B. Reflexive NAT
- C. SSL VPN
- D. L2 VPN

Suggested Answer: D

Community vote distribution

D (100%)

- 🗉 **NetRock1** Highly Voted 2 years, 11 months ago
D is correct - "Enable VM mobility use cases such as migration and disaster recovery without IP address changes."
upvoted 5 times
- 🗉 **vines** Most Recent 5 months, 2 weeks ago
Selected Answer: D
D is correct
upvoted 1 times
- 🗉 **Alchot** 1 year, 4 months ago
D is correct. This option allows extending the on-prem network and recovery without changing IP
upvoted 2 times
- 🗉 **abhi707** 1 year, 5 months ago
D is correct
upvoted 1 times
- 🗉 **abo2020** 1 year, 6 months ago
D.. same broadcast domain
upvoted 2 times
- 🗉 **Me_Loi1** 2 years, 9 months ago
D is correct
upvoted 2 times
- 🗉 **tungdt** 2 years, 11 months ago
D is correct
upvoted 4 times

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

This information was gathered during a workshop:

- ⇒ There are six hosts and hardware has already been purchased.
- ⇒ Customer is planning a collapsed Management/Edge/Compute cluster.
- ⇒ Each host has two 10Gb NICs connected to a pair of switches.
- ⇒ There should be no single point of failure in any proposed design.

Which virtual switch design should the architect recommend to the organization? (Choose the best answer.)

- A. Create a vSphere Distributed Switch (vDS) for Management VMkernel traffic and assign one NIC. Also, create an NSX-T Virtual Distributed Switch (N-VDS) for overlay traffic and assign one NIC.
- B. Create an NSX-T Virtual Distributed Switch (N-VDS) for Management VMkernel traffic and assign one NIC. Also, create an NSX-T Virtual Distributed Switch (N-VDS) for overlay traffic and assign one NIC.
- C. Create an NSX-T Virtual Distributed Switch (N-VDS) for Management VMKernel and overlay traffic and assign both NICs.
- D. Create an NSX-T Virtual Distributed Switch (N-VDS) for Management VMkernel and overlay traffic and assign a new virtual NIC.

Suggested Answer: A

Community vote distribution

C (100%)

🗳️ 👤 **RevantT** Highly Voted 👍 2 years, 6 months ago

Correct answer: C

upvoted 13 times

🗳️ 👤 **Mezze** Most Recent 🕒 7 months, 1 week ago

C is correct

upvoted 1 times

🗳️ 👤 **vGAN** 7 months, 2 weeks ago

with NSX-T 3.0 and vSphere 7.0 we can use vDS instead of nVDS , ill go with C with the options I have

upvoted 1 times

🗳️ 👤 **TristynW** 10 months ago

Selected Answer: C

"No single point of failure"

You have no choice but to assign MGMT and overlay traffic on each nic

upvoted 2 times

🗳️ 👤 **Alchot** 10 months, 2 weeks ago

C is correct.

Both 10Gb NICs must be assigned to a virtual switch to avoid a single point of failure.

upvoted 1 times

🗳️ 👤 **abo2020** 1 year ago

Selected Answer: C

C is the correct one for using the 2 nics

upvoted 1 times

🗳️ 👤 **airmouse1234** 1 year, 3 months ago

Correct answer: C

upvoted 1 times

🗳️ 👤 **shahedhasib** 2 years, 2 months ago

Correct answer is unquestionably C.

upvoted 4 times

🗳️ 👤 **Me_Loi1** 2 years, 3 months ago

Correct answer: C

upvoted 4 times

🗨️ 👤 **diegof1** 2 years, 4 months ago

C is correct.

upvoted 2 times

🗨️ 👤 **NetRock1** 2 years, 5 months ago

C is the correct answr!

upvoted 1 times

🗨️ 👤 **tungdt** 2 years, 5 months ago

C is correct

upvoted 1 times

🗨️ 👤 **MJ86** 2 years, 5 months ago

The correct answer is C.

Options A & B is incorrect as it talks about a single NIC for both VDS & N-VDS. As per the requirement, there should not be a single failure in the proposed design.

D is wrong as hosts only have 2 PNICs. Option D talks about adding a new NIC to NVDS

upvoted 4 times

🗨️ 👤 **VMwareARCHI** 2 years, 6 months ago

Correct answer: C

upvoted 1 times

What selection is the key design benefit provided by a dedicated Edge Cluster VM or Bare Metal? (Choose the best answer.)

- A. reduced administrative overhead
- B. predictable network performance
- C. multiple Tier-0 gateways per Edge Node Cluster
- D. support for Edge Node Clusters with more than 10 Edge Nodes

Suggested Answer: B

Community vote distribution

B (100%)

 **MJ86** Highly Voted 2 years, 11 months ago

Answer B is correct.

<https://docs.vmware.com/en/VMware-Validated-Design/services/deployment-of-nsx-t-edge-nodes-on-bare-metal-hardware-for-vmware-cloud-foundation-40/GUID-563E93F0-65C8-4649-B62F-9AFE89B08B50.html>

<https://vxplanet.com/2019/06/13/nsx-t-edges-baremetal-vs-vm-comparison/>
upvoted 10 times

 **Ansari678** Most Recent 3 months, 1 week ago

Selected Answer: B

B. Predictable network performance

A dedicated Edge Cluster ensures that the network performance for edge services, such as routing and firewalling, can be consistently maintained and doesn't compete with compute workloads, reducing the risk of performance variations and ensuring predictable network performance.


upvoted 1 times

 **misterto** 1 year ago

Answer B. Faster Convergence with BGP peering. 3 Seconds for Edges, 500ms for Bare Metal
upvoted 1 times

 **shahedhasib** 2 years, 8 months ago

The answer is undoubtedly A
upvoted 2 times

 **mk0123** 2 years, 3 months ago

I see this question as more tricky than it would seem on the first sight - you see word performance and go with option B immediately, but the wording is different:

PREDICTABLE network performance

So what does make bare metal edge MORE predicatble than VM performance wise? In my opinion the fact that you have more raw compute power and thus more performance doesn't make it predictable.

Reading VVD <https://docs.vmware.com/en/VMware-Validated-Design/services/deployment-of-nsx-t-edge-nodes-on-bare-metal-hardware-for-vmware-cloud-foundation-40/GUID-563E93F0-65C8-4649-B62F-9AFE89B08B50.html> there are more arrows for the difference about manageability rather than performance. So I'm more inclined for answer A rather than B (C and D are plain wrong).

upvoted 1 times

 **SkylarZ86** 2 years, 1 month ago

there are differences in manageability between the Bare Metal and virtual edge nodes true, but nothing is a reduced overhead with Bare Metal Edge nodes, it even adds complexity to the design and deployment, B is the correct answer.

upvoted 1 times

 **tedybear** 2 years, 2 months ago

it doesnt need to contend with other vm for resources.

upvoted 3 times

🗨️ 👤 **mk0123** 2 years, 2 months ago

That is true and a good point, but it is quite speculative - you can for sure design the solution so there won't be any CPU contention. Looking on it from your perspective maybe more predicatable performance is in the non-existent overhead of the hypervisor layer in the case of bare-metal.

Still the VVD is quite clear (at least I see it there).

upvoted 2 times

🗨️ 👤 **Me_Loi1** 2 years, 9 months ago

B is correct

upvoted 2 times

🗨️ 👤 **diegof1** 2 years, 9 months ago

B is correct

upvoted 3 times

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:

- ⇒ There is a performance based SLA for East to West traffic.
- ⇒ The business critical applications require prioritization of their traffic.
- ⇒ One of the services is a file share and has a high demand for bandwidth.

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.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **UdayShanu** Highly Voted 👍 2 years, 11 months ago

B is correct

upvoted 10 times

🗳️ 👤 **helal** 2 years, 10 months ago

really Confuse

upvoted 1 times

🗳️ 👤 **Ansari678** Most Recent 🕒 3 months, 1 week ago

Selected Answer: B

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

This choice involves setting up Quality of Service (QoS) profiles for network segments to prioritize and manage traffic effectively, which aligns with the need for prioritizing traffic for business-critical applications and ensuring a performance-based SLA for East-West traffic. It addresses the specific requirements mentioned in the assessment phase and allows for the necessary network traffic management.

upvoted 1 times

🗳️ 👤 **lockey37** 1 year, 2 months ago

Mentions some high requirements for East/West traffic ; so not sure about A/ as it mentions "North/South" traffic. I'd go for B as well.

upvoted 1 times

🗳️ 👤 **airmouse1234** 1 year, 9 months ago

B is correct

upvoted 1 times

🗳️ 👤 **Me_Loi1** 2 years, 9 months ago

B is correct

upvoted 1 times

🗳️ 👤 **diegof1** 2 years, 10 months ago

B is Correct

QoS provides high-quality and dedicated network performance for preferred traffic that requires high bandwidth. The QoS mechanism does this by prioritizing sufficient bandwidth, controlling latency and jitter, and reducing data loss for preferred packets even when there is a network congestion.

Taken from <https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.0/administration/GUID-62BB7145-EDD7-4611-A50D-17F4A0EAE57C.html>



upvoted 3 times

🗳️ 👤 **Alefin** 2 years, 10 months ago

B is correct

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.1/administration/GUID-62BB7145-EDD7-4611-A50D-17F4A0EAE57C.html>

upvoted 1 times

  **Revant** 2 years, 12 months ago

Correct Answer: A

upvoted 1 times


Which NSX-T feature is used to allocate the network bandwidth to business-critical applications and to resolve situations where several types of traffic compete for common resources? (Choose the best answer.)

- A. Network I/O Control Profiles
- B. LLDP Profile
- C. LAG Uplink Profile
- D. Transport Node Profiles

Suggested Answer: A

Reference:

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.4/installation/GUID-9A8FD62A-F099-4329-8220-6D5853F9A62D.html>

 **Ansari678** 3 months, 1 week ago

A. Network I/O Control Profiles

Network I/O Control (NIOC) allows you to allocate and prioritize network bandwidth based on predefined profiles, ensuring that critical applications receive the necessary network resources and Quality of Service.

upvoted 1 times

 **airmouse1234** 1 year, 9 months ago


A is correct

upvoted 1 times

 **Me_Loi1** 2 years, 9 months ago

A is correct

upvoted 1 times

 **diegof1** 2 years, 10 months ago

A is correct


Use the Network I/O Control (NIOC) profile to allocate the network bandwidth to business-critical applications and to resolve situations where several types of traffic compete for common resources.

NIOC profile introduces a mechanism to reserve bandwidth for the system traffic based on the capacity of the physical adapters on a host. Version 3 of the Network I/O Control feature offers improved network resource reservation and allocation across the entire switch.

Network I/O Control version 3 for NSX-T Data Center supports the resource management of the system traffic related to virtual machines and to infrastructure services, such as vSphere Fault Tolerance. System traffic is strictly associated with an ESXi host.

Taken from <https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.0/installation/GUID-9A8FD62A-F099-4329-8220-6D5853F9A62D.html>

upvoted 4 times

 **NetRock1** 2 years, 11 months ago

A is the correct answer!

upvoted 1 times

 **tungdt** 2 years, 11 months ago

A is correct

upvoted 1 times

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:

- ⇒ Customer currently has a single 10 host vSphere cluster.
- ⇒ Customer wants to improve network security and automation.
- ⇒ Current cluster utilization and business policies prevent changing the existing vSphere deployment.
- ⇒ High-availability is important to the customer.

Which three selections should the architect include in their design? (Choose three.)

- A. Apply vSphere DRS VM-Host anti-affinity rules to the virtual machines of the NSX-T Edge cluster.
- B. Deploy at least two NSX-T Edge virtual machines in the vSphere cluster.
- C. Deploy the NSX Controllers in the management cluster.
- D. Apply vSphere Distributed Resource Scheduler (vSphere DRS) VM-Host anti-affinity rules to NSX Managers.
- E. Remove 2 hosts from the cluster and create a new edge cluster.
- F. Remove vSphere DRS VM-Host affinity rules to the NSX-T Controller VMs.

Suggested Answer: ACE

Community vote distribution

ABD (100%)

🗳️ **Revant** Highly Voted 2 years, 6 months ago

Correct Answer: ABD

upvoted 16 times

🗳️ **shahedhasib** Highly Voted 2 years, 2 months ago

Correct answer is certainly ABD.

Why C is not correct:

It is strictly mentioned that "Current cluster utilization and business policies prevent changing the existing vSphere deployment."

So we are NOT permitted to create a separate "management cluster"

Why E is not correct:

It is strictly mentioned that "Current cluster utilization and business policies prevent changing the existing vSphere deployment."

So we are NOT permitted to create a separate "new edge cluster"

upvoted 8 times

🗳️ **4ourDS** Most Recent 3 months, 4 weeks ago

Selected Answer: ABD

Since NSX-T, controllers are not deployed separately.

upvoted 1 times

🗳️ **vGAN** 7 months, 2 weeks ago

ABD 100%

upvoted 1 times

🗳️ **Alchot** 10 months, 2 weeks ago

ABD is correct.

upvoted 1 times

🗳️ **Mezze** 1 year ago

ABD is correct

upvoted 1 times

🗳️ **abo2020** 1 year ago

ABD , Prefect explanation is below from MJ86

upvoted 1 times

🗳️ **airmouse1234** 1 year, 3 months ago

Correct Answer: ABD

upvoted 1 times

🗨️ 👤 **Me_Loi1** 2 years, 3 months ago

Correct Answer: ABD

upvoted 2 times

🗨️ 👤 **diegof1** 2 years, 4 months ago

The correct answer is ABD

upvoted 3 times

🗨️ 👤 **tungdt** 2 years, 5 months ago

ABD is correct

upvoted 1 times

🗨️ 👤 **MJ86** 2 years, 5 months ago

The correct answer is A B D

F is wrong because NSX-T 2.4 onwards, there is no controller VM's. Controllers were merged with NSX-T Manager. Option C is also wrong for the same reason.

E is wrong because it talks about modifying the vSphere cluster. The 3rd requirement of design says no changes can be made to vSphere deployment.

upvoted 6 times

An architect is helping an organization with the Conceptual Design of an NSX-T Data Center solution. This information was gathered by the architect during the Discover Task of the Engagement Lifecycle:

- ⇒ There are applications which use IPv6 addressing.
- ⇒ Network administrators are not familiar with NSX-T Data Center solutions.
- ⇒ Hosts can only be configured with two physical NICs.
- ⇒ There is an existing management cluster to deploy the NSX-T components.
- ⇒ Dynamic routing should be configured between the physical and virtual network.
- ⇒ There is a storage array available to deploy NSX-T components.

Which constraint was documented by the architect? (Choose the best answer.)

- A. Dynamic routing should be configured between the physical and virtual network.
- B. There are applications which use IPv6 addressing.
- C. Hosts can only be configured with two physical NICs.
- D. There are enough CPU and memory resources in the existing management cluster.

Suggested Answer: A

Community vote distribution

C (100%)

🗳️ **VMwareARCHI** Highly Voted 2 years, 12 months ago
C. Hosts can only be configured with two physical NICs.
upvoted 12 times

🗳️ **diegof1** Highly Voted 2 years, 10 months ago
C is the correct.

A constraint is something that will limit the choices you have for a design.
upvoted 6 times

🗳️ **Ansari678** Most Recent 3 months, 1 week ago
Selected Answer: C
C. Hosts can only be configured with two physical NICs.

This constraint indicates that the hosts in the organization can only have two physical network interface cards (NICs) available for use, which may impact the network design and resource allocation in the NSX-T Data Center solution.
upvoted 1 times

🗳️ **Nabil1986** 10 months, 3 weeks ago
C is correct
upvoted 1 times

🗳️ **Alchot** 1 year, 4 months ago
By the way the statement is written C sounds more correct.
On that scenario A sounds more accurate. Dynamic routing with IPv6 is possible only with BGP.
I go with A
upvoted 1 times

🗳️ **airmouse1234** 1 year, 9 months ago
C is correct answer.
upvoted 1 times

🗳️ **Praveen2929** 1 year, 12 months ago
Selected Answer: C
C is correct answer
upvoted 1 times

🗳️ **nick2u** 2 years, 2 months ago

I would say A as support of dynamic routing in NSX-T is limited plus it is for IPv6. C is a constraint as well but not necessarily applicable to this use case basing on the info provided.

upvoted 3 times

🗨️ 👤 **SkylarZ86** 2 years, 1 month ago

C is the correct answer, A is a requirement.

upvoted 3 times

🗨️ 👤 **Me_Loi1** 2 years, 9 months ago

C is correct

upvoted 2 times

🗨️ 👤 **NetRock1** 2 years, 11 months ago

C is a constraint!

upvoted 2 times

🗨️ 👤 **tungdt** 2 years, 11 months ago

C is correct

upvoted 2 times

Which two benefits can be achieved using in-band management of an NSX Bare Metal Edge Node? (Choose two.)

- A. Reduces storage requirements.
- B. Reduces cost.
- C. Preserves packet locality.
- D. Reduces egress data.
- E. Preserves switchports.

Suggested Answer: *BD*

Community vote distribution

BE (100%)

🗳️ 👤 **diegof1** Highly Voted 📈 2 years, 4 months ago
B & E are correct.

When using in-band management interface, the management interface and Edge TEP traffic are on the same NIC.
<https://docs.vmware.com/en/VMware-Validated-Design/services/deployment-of-nsx-t-edge-nodes-on-bare-metal-hardware-for-vmware-cloud-foundation-40/GUID-AAA3EDD0-2F19-49F8-B9B3-F5B1505CBB28.html>
upvoted 10 times

🗳️ 👤 **4ourDS** Most Recent 🕒 3 months, 4 weeks ago
Selected Answer: BE
We can't reduce traffic, so we're separating them.
upvoted 1 times

🗳️ 👤 **Nabil1986** 4 months, 2 weeks ago
B&E are correct
upvoted 1 times

🗳️ 👤 **abo2020** 1 year ago
i would go for BE
Inband : mgmt and TEP traffic would go out from the same nic hence we achieve - preserve switch ports which will reduce costs
upvoted 3 times

🗳️ 👤 **Arund** 1 year, 1 month ago
B & E my choice!
upvoted 1 times

🗳️ 👤 **airmouse1234** 1 year, 3 months ago
B & E are correct.
upvoted 1 times

🗳️ 👤 **Praveen2929** 1 year, 6 months ago
Selected Answer: BE
B and E is indeed correct answers
upvoted 1 times

🗳️ 👤 **fahadnaif** 1 year, 11 months ago
since when the bare metal reduce the cost?.
upvoted 1 times

🗳️ 👤 **abo2020** 1 year ago
i would go for BE
Inband : mgmt and TEP traffic would go out from the same nic hence we achieve - preserve switch ports which will reduce costs
upvoted 1 times

🗳️ 👤 **Me_Loi1** 2 years, 3 months ago
B & E are correct.

upvoted 1 times

🗨️ 👤 **NetRock1** 2 years, 5 months ago

B & E are correct.

When you don't use an extra physical network port, you Reduces cost and Preserves switchports!

upvoted 2 times

🗨️ 👤 **VMwareARCHI** 2 years, 5 months ago

Correct Answer are B & E

upvoted 1 times

🗨️ 👤 **salei** 2 years, 5 months ago

Correct Answer are B & E. You reduce cost because you don't need to use a dedicated port to management and, hence, you preserve ports on the switch

upvoted 3 times

🗨️ 👤 **Revant** 2 years, 6 months ago

Correct Answer: DE

upvoted 1 times

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

This information was gathered during a workshop:


- ⇒ Any proposed solution must provide low latency.
- ⇒ Any proposed solution must provide high throughput.
- ⇒ Customer is running stock trading applications.

Which two selections should the architect recommend to meet high-performance workload requirements? (Choose two.)

- A. Leverage ESXi as the compute host.
- B. Use LACP for all uplink profiles.
- C. Leverage KVM as the compute host.
- D. Enable enhanced data path mode on the N-VDS.
- E. Enable latency sensitivity mode on the N-VDS.

Suggested Answer: AD

Community vote distribution



🗨️ **diegof1** Highly Voted 3 years, 4 months ago

A & D are correct.

N-VDS and VDS support the following modes:

Standard: Provide forwarding capabilities on both KVM and ESX transport nodes and does not require a specialized hardware.

Enhanced Datapath: Available for ESXi transport nodes only and provide Enhanced Network Stack (ENS) targeted for Network Functions Virtualization (NFV) applications that require a faster performance data path.

Take from NSX-T ICM 3.0 Lecture Manual

upvoted 13 times

🗨️ **WQL** Highly Voted 3 years, 5 months ago

AD is correct.

The enhanced data path is available for EXSi only.

upvoted 5 times

🗨️ **VMwareARCHI** 3 years, 5 months ago

you are right! A/D

upvoted 5 times

🗨️ **AT45816** Most Recent 6 months ago

Selected Answer: BD

B & D are correct

upvoted 1 times

🗨️ **82bb1ba** 6 months, 3 weeks ago

Per AI: Considering the nature of stock trading applications, where low latency is crucial, the architect might prioritize options B and D to address the high-throughput and low-latency requirements. The choice between ESXi and KVM (Options A and C) would depend on other organizational factors and preferences. Option E could also be considered if further optimization for latency sensitivity is needed.

upvoted 1 times

🗨️ **Ansari678** 9 months, 2 weeks ago

Selected Answer: AD

A. Leverage ESXi as the compute host.

ESXi is a VMware hypervisor known for its performance capabilities and is often chosen for low-latency, high-throughput applications.

D. Enable enhanced data path mode on the N-VDS.

Enhanced data path mode is designed to improve performance and reduce latency, making it suitable for high-performance workloads.

upvoted 1 times

🗨️ **Tachinsky** 1 year ago

Selected Answer: AD

Enhanced Datapath is correct and requires ESXi -> AD

upvoted 1 times

🗨️ **borracho13** 1 year ago

Selected Answer: BD

Using ESXi isn't a selection whereas LACP and Enhanced Datapath are configuration selection. Since Enhanced Datapath can only be used with ESXi, it is assumed if you select D, you are running ESXi already.

upvoted 2 times

🗨️ **borracho13** 1 year ago

Correction A & D

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/installation/GUID-F459E3E4-F5F2-4032-A723-07D4051EFF8D.html>

upvoted 1 times

🗨️ **4ourDS** 1 year, 3 months ago

Selected Answer: BD

The compute host is the environment in which the VM runs, so it seems pointless to mention it.

upvoted 3 times

🗨️ **vGAN** 1 year, 7 months ago

B seems better compared to A ..

upvoted 1 times

🗨️ **nick2u** 2 years, 8 months ago

BD. ESXi might not provide better throughput comparing to KVM. But LAG/LACP would.

upvoted 1 times

🗨️ **nick2u** 2 years, 8 months ago

I took it back...AD...

upvoted 2 times

🗨️ **shahedhasib** 3 years, 2 months ago

Correct answers are undoubtedly B,D

upvoted 4 times

🗨️ **Me_Loi1** 3 years, 3 months ago

A & D are correct.

upvoted 2 times

🗨️ **tungdt** 3 years, 5 months ago

AD is correct

upvoted 2 times

🗨️ **VMwareARCHI** 3 years, 5 months ago

BD is correct

upvoted 2 times

🗨️ **UdayShanu** 3 years, 5 months ago

BD is correct

upvoted 2 times

Which selection is associated with the Review Task of the Engagement Lifecycle? (Choose the best answer.)

- A. Gather and document requirements, assumptions, and constraints.
- B. Build, deploy, implement, and test the design.
- C. Measure performance against customer's objective.
- D. Create and document the logical and physical design.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Rawatam** Highly Voted 📌 2 years, 3 months ago
C is the correct. Just checked the design guide
upvoted 5 times

🗳️ 👤 **Ansari678** Most Recent 🕒 3 months, 1 week ago
Selected Answer: C
C. Measure performance against the customer's objective.

During the Review Task, you assess and evaluate the design's performance against the customer's requirements and objectives, ensuring that it meets the intended goals.

upvoted 1 times

🗳️ 👤 **Nabil1986** 10 months, 3 weeks ago
C is correct
upvoted 1 times

🗳️ 👤 **shahedhasib** 2 years, 8 months ago
In my opinion the correct answer is D as review tasks in related to "Create and document the logical and physical design".
Answer C is not correct as "Measure performance against customer's objective." is a task that is done during the system validation phase.
upvoted 1 times

🗳️ 👤 **BrianOC** 2 years, 9 months ago
C is a correct answer.
<https://docs.vmware.com/en/VMware-Validated-Design/6.1/sddc-architecture-and-design-for-the-management-domain/GUID-1117D50C-096D-40B8-84C0-B9D636E322C6.html>
upvoted 4 times

🗳️ 👤 **Me_Loi1** 2 years, 9 months ago
I think C is correct because you need the gathered measurement criterias to post review the entire Engagement Lifecycle process.
upvoted 3 times

🗳️ 👤 **diegof1** 2 years, 9 months ago
For me A is the correct answer.

The engagement lifecycle is associated with gather and document requirements, assumptions, and constraints. In fact, a lot of question associated with risk, requirements, assumptions, and constraints says something like: "During the Discover Task of the Engagement Lifecycle"...

upvoted 3 times

🗳️ 👤 **abo2020** 1 year, 6 months ago
it's C he is asking about the review task
upvoted 1 times

🗳️ 👤 **Pheakdey** 2 years, 10 months ago
Correct Answer: C
upvoted 2 times

🗳️ 👤 **tungdt** 2 years, 11 months ago

C is correct

upvoted 2 times

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:

- ⇒ NSX-T will span across two sites for disaster recovery.
- ⇒ Public Load Balancer VIP should be accessible from a secondary site.
- ⇒ Distributed Firewall Policies should be available at a secondary site.
- ⇒ Routing capabilities should be maintained after failure.
- ⇒ NAT capabilities are required.

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

- A. Use of the same ISPs across sites.
- B. Use two separate ISPs across sites.
- C. Use MTU to 1550 between sites.
- D. Set MTU to 1550 between sites.
- E. Use IP sets or groups to configure DFW rules.

Suggested Answer: AE

Reference:

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/nsx/vmware-multi-site-solutions-cross-vcenter-nsx-design-guide.pdf>

Community vote distribution

AE (100%)

 **MBrown** Highly Voted 2 years, 5 months ago

A and E are the correct answers, the MTU can't be less than 1600.

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.0/administration/GUID-5D7E3D43-6497-4273-99C1-77613C36AD75.html>

upvoted 10 times

 **NetRock1** 2 years, 4 months ago

Agree with you A and E are the correct answers.

upvoted 3 times

 **4ourDS** Most Recent 3 months, 4 weeks ago

Selected Answer: AE


The question seems to be about making sure the architect is aware of how to make use of the multi-region IPs provided by the ISP.

upvoted 1 times

 **joncursio** 4 months, 2 weeks ago


Why the F would you use the same ISP? You can advertise the same IP space to multiple suppliers. There's better fault tolerance using different at each site..?

upvoted 1 times

 **Nabil1986** 4 months, 2 weeks ago

A & E is corr

upvoted 1 times

 **Me_Loi1** 2 years, 3 months ago

I am not sure, as long as you span the NSX layer across both site you are also aware of the NSX/vSphere logics, therefore you don't need to use IP sets, you can easily use all object types for the DFW. Answer A is definitely correct as it is a mandatory requirement stated on the NSX

documentation: <https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.0/administration/GUID-5D7E3D43-6497-4273-99C1-77613C36AD75.html>

But I don't think that E is correct, instead D would be more suitable since the GENEVE header needs exactly 50 Bytes (1550Bytes in total) extra like VXLAN back then. It only depends how much extra metadata is implmented on the TLV from VMware. But technically it could work with 1550Bytes, but I never tested it...

upvoted 1 times

🗨️ 👤 **tedybear** 1 year, 8 months ago

If this is talking about multisite, then the mtu requirement is for 1600, so CD are wrong.

upvoted 1 times

🗨️ 👤 **diegof1** 2 years, 3 months ago

A & E are correct.

upvoted 3 times

🗨️ 👤 **VMwareARCHI** 2 years, 5 months ago

Correct Answer: AD

upvoted 1 times

🗨️ 👤 **RevanT** 2 years, 6 months ago

Correct Answer: AD

upvoted 2 times

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

Which risk is documented by an architect? (Choose the best answer.)

- A. The security team has a firewall communication matrix documented.
- B. The team is not trained for NSX-T but have a very strong experience with vSphere.
- C. Open communication between different application tiers is not allowed.
- D. Aggregate N-S throughput at any given time should be at least 10G.

Suggested Answer: B

Community vote distribution

B (100%)

  **tungdt** Highly Voted 2 years, 5 months ago

B is correct
upvoted 7 times


  **Nabil1986** Most Recent 4 months, 2 weeks ago

Selected Answer: B



B is correct
upvoted 1 times

  **Rawatam** 1 year, 9 months ago

B is good
upvoted 2 times

  **Me_Loi1** 2 years, 3 months ago

B is correct.
upvoted 4 times

  **diegof1** 2 years, 4 months ago

B is correct.
upvoted 4 times

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:

- ⇒ Data between two networks connected over a public network needs to be encrypted.
- ⇒ Certificate authentication is required.
- ⇒ Dynamic route learning is preferred.

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.

Suggested Answer: C

🗨️ **MJ86** Highly Voted 11 months, 3 weeks ago

C is the correct answer.

Route-based IPSec VPN provides tunneling on traffic based on the static routes or routes learned dynamically.

Tier-0 gateway in the active-standby state supports the following services:

NAT
Load balancing
Stateful firewall
VPN

Reference Docs:

<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/administration/GUID-C0E5AF10-576D-493A-A079-C4C95D8F5373.html>

<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/administration/GUID-7B0CD287-C5EB-493C-A57F-EEA8782A741A.html#GUID-7B0CD287-C5EB-493C-A57F-EEA8782A741A>

upvoted 8 times

🗨️ **Rawatam** Most Recent 3 months, 2 weeks ago

C is correct due to dynamic route learning.

upvoted 3 times

🗨️ **Me_Loi1** 9 months, 2 weeks ago

C is correct

upvoted 4 times

🗨️ **diegof1** 10 months ago

C is correct.

<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/administration/GUID-DF689847-252E-451E-84B5-DB507CC010AC.html>

upvoted 4 times

🗨️ **tungdt** 11 months, 2 weeks ago

C is correct

upvoted 2 times

🗨️ **VMwareARCHI** 11 months, 2 weeks ago

correct answer; C

upvoted 2 times

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

This information was gathered during a workshop:

- ⇒ Migrating existing data center to KVM hosts.
- ⇒ Redundancy and high availability are required.
- ⇒ No component can be a single point of failure.

Which selection should the architect recommend? (Choose the best answer.)

- A. Linux Bridge redundancy with Active/Active Mode and multiple pNICs with necessary binding
- B. Linux Bridge redundancy with Active/Active Mode and single pNIC with static binding
- C. vSS/vDS in Active/Standby Mode with necessary binding
- D. vSS/vDS in Active/Active Mode with necessary pNICs and required binding modes

Suggested Answer: C

Community vote distribution

A (100%)

🗨️ **MBrown** Highly Voted 2 years, 5 months ago

As the new DC where the NSX-T will be deployed is based on KVM-Hosts [including both computes and mgmt clusters], so this question is referring basically to the mgmt. cluster which will use Linux bridging in active/active mode to multiple pNICs as uplinks. so answer is A.
upvoted 10 times

🗨️ **diegof1** Highly Voted 2 years, 4 months ago

A is correct.

Since migrating to KVM exist then the answers with vSS/vDS do not apply. The requirement is that no component has a single point of failure so B does not apply either.

upvoted 6 times

🗨️ **4ourDS** Most Recent 3 months, 4 weeks ago

Selected Answer: A

no ESXi in the Transport Nodes, HA required.

upvoted 2 times

🗨️ **Nabil1986** 4 months, 2 weeks ago

A is correct

upvoted 1 times

🗨️ **vGAN** 7 months, 2 weeks ago

A , on KVM we don't have vSS/vDS

upvoted 1 times

🗨️ **airmouse1234** 1 year, 3 months ago

A is correct.

upvoted 1 times

🗨️ **Rawatam** 1 year, 9 months ago

A looks fine

upvoted 1 times

🗨️ **Pal68** 1 year, 10 months ago

I agree with C.. <https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/administration/GUID-A3674EA2-F368-4A53-9078-0694F76624F4.html>

upvoted 1 times

🗨️ **DCL202** 1 year, 9 months ago

I believe this is A, as the test is not based on NSX-T 3.x but rather NSX-T 2.x, which does not support vSS/vDS.

Even if it is NSX-T 3, why would the answer not be D then (Active-Active)?

upvoted 1 times

  **NetRock1** 2 years, 4 months ago

A is the correct answer!

upvoted 2 times

An architect is helping an organization with the Conceptual Design of an NSX-T Data Center solution. This information was gathered by the architect during the Discover Task of the Engagement Lifecycle:

- ⇒ There are applications which use IPv6 addressing.
- ⇒ Network administrators are not familiar with NSX-T Data Center solutions.
- ⇒ Hosts can only be configured with two physical NICs.
- ⇒ There is an existing management cluster to deploy the NSX-T components.
- ⇒ Dynamic routing should be configured between the physical and virtual network.
- ⇒ There is a storage array available to deploy NSX-T components.

Which risk was documented by the architect? (Choose the best answer.)

- A. Network administrators are not familiar with NSX-T Data Center solutions.
- B. Dynamic routing should be configured between the physical and virtual network.
- C. There are applications which use IPv6 addressing.
- D. There is a storage array available to deploy NSX-T components.

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ **VMwareARCHI** Highly Voted 2 years, 5 months ago

A is correct
upvoted 15 times

🗳️ **NetRock1** Highly Voted 2 years, 4 months ago

A is correct.
upvoted 5 times

🗳️ **4ourDS** Most Recent 3 months, 4 weeks ago

Selected Answer: A
Human resources are more important.
upvoted 2 times

🗳️ **Athyl** 6 months, 3 weeks ago

Selected Answer: A
A is correct
upvoted 1 times

🗳️ **vGAN** 7 months, 2 weeks ago

Staff is not trained is a "BIG RISK"
upvoted 1 times

🗳️ **airmouse1234** 1 year, 3 months ago

A is Right.
upvoted 1 times

🗳️ **KD01** 1 year, 5 months ago

A is correct
upvoted 1 times

🗳️ **ShyamC** 1 year, 7 months ago

Selected Answer: A
How is B a risk by using dynamic routing. No knowledge of the employees with NSX is the biggest risk
upvoted 1 times

🗳️ **Rawatam** 1 year, 9 months ago

A is vorrect
upvoted 2 times

🗨️ 👤 **Me_Loi1** 2 years, 3 months ago

A is correct

upvoted 3 times

🗨️ 👤 **diegof1** 2 years, 3 months ago

A is the correct answer.

upvoted 4 times

🗨️ 👤 **tungdt** 2 years, 5 months ago

D is correct

upvoted 3 times

🗨️ 👤 **abo2020** 1 year ago

this is never a risk

upvoted 1 times

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

Suggested Answer: DE

Community vote distribution

AC (100%)

🗨️ **RevanT** Highly Voted 2 years, 6 months ago

Correct Answer: AC
upvoted 22 times

🗨️ **SkylarZ86** 1 year, 7 months ago

AC is correct, Uplink profiles within NSX-T shows that you can only use Load balance Source and Load balance Source MAC, failover order doesnt provide any load balancing or high throughput

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.1/installation/GUID-50FDFDFB-F660-4269-9503-39AE2BBA95B4.html>

upvoted 1 times

🗨️ **MJ86** Highly Voted 2 years, 5 months ago

Correct answer should be C & E. We have to keep in mind that following criteria should be met: low latency, high throughput, redundancy, and performance

Check out this blog post which explains about type of teaming policies

<https://www.vstellar.com/2020/08/05/nsx-t-3-0-seriespart-2-uplink-profiles/#:~:text=Load%20Balance%20Source%20MAC%3A%20This,the%20source%20VM's%20MAC%20address.>

upvoted 6 times

🗨️ **rajavalas** 2 years, 3 months ago

E is Failover Order. So One uplink will be in Standby always. With one link not actively participating in I/O, how do you get high throughput?
upvoted 6 times

🗨️ **vGAN** Most Recent 7 months, 2 weeks ago

A and C

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/installation/GUID-50FDFDFB-F660-4269-9503-39AE2BBA95B4.html>

upvoted 1 times

🗨️ **Alchot** 10 months, 2 weeks ago

AC

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.0/installation/GUID-50FDFDFB-F660-4269-9503-39AE2BBA95B4.html>

upvoted 1 times

🗨️ **rgraghu** 1 year, 7 months ago

Selected Answer: AC

Failover Order: Select an active uplink is specified along with an optional list of standby uplinks. If the active uplink fails, the next uplink in the standby list replaces the active uplink. No actual load balancing is performed with this option.

Load Balance Source: Select a list of active uplinks. When you configure a transport node, you can pin each interface of the transport node to one active uplink. This configuration allows use of several active uplinks at the same time.

Load Balance Source MAC Address: Select an uplink based on a hash of the source Ethernet.

upvoted 4 times

🗨️ 👤 **Rawatam** 1 year, 9 months ago

AC is the correct answer.

upvoted 3 times

🗨️ 👤 **Pal68** 1 year, 10 months ago

AC for Teaming policy with 3.0

- <https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/installation/GUID-50FDFDFB-F660-4269-9503-39AE2BBA95B4.html?hWord=N4IghgNiBclDIHswBMAEAhSYB2BjApqgMolCuATgagLICCawIAL5A>

upvoted 2 times

🗨️ 👤 **BrianOC** 1 year, 12 months ago

NSX-T Virtual Distributed Switch(N-VDS) supports Load balance source and Failover teaming policies. When you use the Load balance source policy, both physical NICs can be active and carry traffic.

<https://docs.vmware.com/en/VMware-Validated-Design/5.1/sddc-architecture-and-design-for-vmware-nsxt-workload-domains/GUID-6D7C584B-8E2F-4FEB-908D-7EDD1FD3C5ED.html>

The correct answers are Load Balance Source and Failover (C & E)

upvoted 1 times

🗨️ 👤 **Me_Loi1** 2 years, 3 months ago

Correct Answer: CE

upvoted 2 times

🗨️ 👤 **diegof1** 2 years, 3 months ago

C & E are the correct answer.

Uplink Policy for ESXi Host Transport Nodes

Uplink profiles define policies for the links from ESXi hosts to NSX-T segments or from NSX-T Edge appliances to top of rack switches. By using uplink profiles, you can apply consistent configuration of capabilities for network adapters across multiple ESXi hosts or NSX-T Edge nodes.

Uplink profiles can use either load balance source or failover order teaming. If using load balance source, multiple uplinks can be active. If using failover order, only a single uplink can be active.

Design Decisions on the Uplink Profile for ESXi Transport Nodes....

Taken from <https://docs.vmware.com/en/VMware-Validated-Design/6.2/sddc-architecture-and-design-for-a-virtual-infrastructure-workload-domain/GUID-B1A08A06-F475-436B-8BA4-31552450D63C.html>

and

<https://docs.vmware.com/en/VMware-Validated-Design/6.2/sddc-architecture-and-design-for-the-management-domain/GUID-BEDFA883-AE2C-4665-972A-179E88BB6B6A.html>

upvoted 5 times

🗨️ 👤 **MBrown** 2 years, 5 months ago

A and D are the correct answer as per the requirements.

upvoted 2 times

🗨️ 👤 **EW321** 2 years, 1 month ago

Agree with u. E is Failover Order => active/standby => does not meet the requirement of high throughput.

Load balanced source Port and Load balance Source Mac is also mentioned in the NSX-T reference design guide under Teaming Policy section

upvoted 3 times

🗨️ 👤 **tungdt** 2 years, 5 months ago

CE is correct

upvoted 3 times

 **VMwareARCHI** 2 years, 5 months ago

Correct Answer: CE

upvoted 3 times

Which two resources can be used by an NSX architect during the Assessment Phase? (Choose two.)

- A. vRealize Network Insight
- B. VMware Validated Design
- C. VMware customer references
- D. key stakeholder interviews
- E. application licensing

Suggested Answer: AE

Community vote distribution

AD (67%)

AB (33%)

🗳️ **VMwareARCHI** Highly Voted 3 years, 6 months ago

Correct Answer: A / D
upvoted 9 times

🗳️ **AT45816** Most Recent 6 months ago

Selected Answer: AD

Correct Answer: A & D
upvoted 1 times

🗳️ **82bb1ba** 6 months, 3 weeks ago

Per AI: the two most directly relevant resources are:

A. vRealize Network Insight: This tool is used to assess the existing network environment, providing visibility into network traffic, security, and performance.

D. Key stakeholder interviews: Interviews with key stakeholders are crucial for understanding the organization's needs, challenges, and requirements. This information is essential for the assessment phase.
upvoted 1 times

🗳️ **sierrabobs** 1 year ago

Selected Answer: AB

In the assessment phase the following is done:

1. Use the vmware tools to collect required data
2. Finalize an assessment report
3. Present a transition workshop

The question is asking for "resources". Since we are limited to "resources", I believe A and B to be the correct answers.

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/docs/vmware-nsx-migration-standard-for-vmware-cloud-director-datasheet.pdf>

upvoted 1 times

🗳️ **4ourDS** 1 year, 3 months ago

Selected Answer: AD

VMware provides an assessment pre-sale scenario with a 60-day evaluation period.
upvoted 1 times

🗳️ **Alchot** 1 year, 10 months ago

A - VRNI can do network assessment for free.
D- Interviews to gather requirements and constraints
upvoted 1 times

🗳️ **Johannes2109** 1 year, 10 months ago

I prefer B,D.
Why not A - How can I expect to have a vRNI installed in the assessment phase?

B - I can use VMware Validated Design as a resource at any time to design the preferred target design for the stakeholder

Why not C - customer reference is not a resource to help draft the target design.

D - Interviews with key stakeholders are needed as a resource to get the key information to shape the target design.



E - I'm a little unsure. What is meant by application licensing? Customer Application Licensing or VMware Component Licensing? If it's the last one, then it might as well be E. Because the customer has to decide how much he is willing to pay for it. This affects the design decisions.

upvoted 1 times

  **airmouse1234** 2 years, 3 months ago



Correct Answer: A / D

upvoted 1 times

  **Rawatam** 2 years, 9 months ago



Why A is correct if it's assessment phase. D is good and E should possible because this may be constraint, one of things to find in asses6

upvoted 1 times

  **nick2u** 2 years, 8 months ago


Question is asking "resources". VNI is a resource. How can app licensing be a resource?

upvoted 1 times

  **diegof1** 3 years, 3 months ago



A and D are correct for the assessment phase.

upvoted 4 times

  **NetRock1** 3 years, 4 months ago



A and D are correct!

upvoted 2 times

  **RevanT** 3 years, 5 months ago

Agreed, forgot the A in the first post.

upvoted 2 times

  **RevanT** 3 years, 6 months ago

Correct Answer: D

upvoted 2 times

A customer wants to place their NSX Managers in different subnets.

Which would an architect recommend to support the request? (Choose the best answer.)

- A. Use a load balancer.
- B. Use round-robin DNS.
- C. Use NAT.
- D. Use a cluster Virtual IP.

Suggested Answer: D

Reference:

<http://www.cloudxtreme.info/nsx-t-manager-clustering/>

Community vote distribution

A (100%)

🗃️ 👤 **Revant** Highly Voted 👍 2 years, 6 months ago

Correct Answer: A

upvoted 16 times

🗃️ 👤 **diegof1** Highly Voted 👍 2 years, 3 months ago

The Correct answer is A.

NSX Management Cluster with Load Balancer

A load balancer provides high availability to the NSX Management Cluster:

- All nodes are active
- GUI and API are available on all manager.
- Traffic to the virtual IP address is load balanced to multiple manager node
- The Manager nodes can be in different subnets.

Taken from NSX-T ICM 3.0 Lecture Manual.

upvoted 5 times

🗃️ 👤 **4ourDS** Most Recent 🕒 3 months, 4 weeks ago

Selected Answer: A

LB provides wider operational compatibility than VIP.

upvoted 1 times

🗃️ 👤 **vGAN** 7 months, 2 weeks ago

A , for cluster VIP they need to be same Segment

upvoted 1 times

🗃️ 👤 **vivekbristol** 7 months, 2 weeks ago

Selected Answer: A

A is correct

upvoted 1 times

🗃️ 👤 **Alchot** 10 months, 2 weeks ago

A is correct.

An external load balancer can provide the following benefits:

Load balance across the NSX Managers.

The NSX Managers can be in different subnets.

Fast recovery time in the event of a Manager node failure.

upvoted 1 times

🗨️ 👤 **airmouse1234** 1 year, 3 months ago

Correct Answer: A

upvoted 1 times

🗨️ 👤 **KD01** 1 year, 5 months ago

A is correct

upvoted 1 times

🗨️ 👤 **Rawatam** 1 year, 9 months ago

A is right

upvoted 2 times

🗨️ 👤 **Me_Loi1** 2 years, 3 months ago

Answer A, because the "integrated" cluster VIP need to have all its NSX Managers in the same subnet due the GARP process during a NSX Manager node failure therefore only an external LB can be used.

upvoted 2 times

🗨️ 👤 **vKramer** 2 years, 5 months ago

Correct Answer is A, cluster VIP is only supported if all members of the management cluster are in the same subnet.

upvoted 2 times

🗨️ 👤 **MJ86** 2 years, 5 months ago

A is indeed the right answer. External load balancer is needed when NSX-T managers are not in the same L2 domain. This is explained here: <http://www.cloudxtreme.info/nsx-t-manager-clustering/>

upvoted 4 times

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

This information was gathered during a workshop:

- ⇒ Current hypervisor of choice is KVM.
- ⇒ Cost reduction is important.

Which two selections should the architect recommend to the organization? (Choose two.)

- A. Deploy Edge VM Nodes using ISO.
- B. Deploy NSX Manager using OVF.
- C. Deploy NSX Manager using QCOW2.
- D. Deploy bare metal Edge Nodes.
- E. Deploy Edge VM Nodes on KVM.

Suggested Answer: CD

Community vote distribution

CD (100%)

 **altah3r** Highly Voted 2 years, 4 months ago

C and D correct

A and C is wrong you cannot setup edge nodes on KVM

<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.1/installation/GUID-11417AA2-5EBC-49C7-8A86-EB94604261A6.html>

B wrong because NSX manager can be installed only using qcow

<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.1/installation/GUID-5229A83D-1B97-4203-BA30-F52716F68F7F.html>

upvoted 9 times

 **Alefin** Highly Voted 2 years, 4 months ago

C-D are correct

upvoted 7 times

 **4ourDS** Most Recent 3 months, 4 weeks ago

Selected Answer: CD

Situations where there should be no vSphere license purchases.

upvoted 1 times

 **private0330** 11 months, 3 weeks ago

Selected Answer: CD

C is correct because you must use QCOW2 to deploy on KVM Hypervisor

D is correct because NSX Edge is supported only on ESXi or on bare metal.


<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.1/installation/GUID-11417AA2-5EBC-49C7-8A86-EB94604261A6.html>

upvoted 1 times

 **KD01** 1 year, 5 months ago


C & D are correct

upvoted 1 times

 **Me_Loi1** 2 years, 3 months ago

Correct is CD but if cost reduction is an important strategy then you should not invest in NSX-T to be honest :)

upvoted 4 times

 **diegof1** 2 years, 3 months ago

C & D are correct.



Due hypervisors are KVM NSX Manager needs to be deployed from QCOW2 image.

<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/installation/GUID-A65FE3DD-C4F1-47EC-B952-DEDF1A3DD0CF.html>

Also, for the KVM environment NSX Edge is not supported on KVM. So, we have to deploy it on bare-metal form factor. NSX Edge Node just support VM on an ESXi host.


<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.0/installation/GUID-5EF2998C-4867-4DA6-B1C6-8A6F8EBCC411.html>

upvoted 6 times

  **tungdt** 2 years, 5 months ago

AC is correct

upvoted 1 times

  **BlackTai** 2 years, 5 months ago

Answer: A,C

upvoted 1 times

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

This information was gathered during a workshop:

- ⇒ The company will use a Leaf and Spine physical network architecture with Layer 3 gateways for top of rack switches.
- ⇒ The company is planning to deploy 120 ESX hosts across 10 racks.
- ⇒ There will be a total of 12 clusters where each cluster has one host per rack.

What should the architect recommend to allow applications to run on any host in the cluster? (Choose the best answer.)

- A. Deploy all application networks on NSX segments.
- B. Deploy an L2 VPN to allow the networks to extend to each host.
- C. Deploy a Tier-0 gateway per Rack and configure BGP between racks.
- D. Deploy a Tier-1 gateway per Rack and configure BGP between racks.

Suggested Answer: D

Community vote distribution

A (67%) D (33%)

🗳️ **vkramer** Highly Voted 2 years, 5 months ago

Correct Answer is A. NSX Segments can span across multiple racks and subnets. D is not correct since it is not possible to configure BGP between Tier-1 gateways.

upvoted 14 times

🗳️ **tester2478** Most Recent 3 months, 2 weeks ago

A looks like to be the best answer since question says What should the architect recommend to allow "applications" to run on any host in the cluster?

upvoted 1 times

🗳️ **misterto** 6 months, 2 weeks ago

I would say C. Here is my justification.

A: Deploy all application networks on NSX segments: Doesn't precise if VLAN backed or Overlay Segment. Not precise enough.

D: Impossible to use BGP on T1

B: Possible only if C is done previously

C: good answer as we will have 1 cluster per rack, so VM won't move with HA b/w racks.

upvoted 2 times

🗳️ **lqh61727** 9 months, 2 weeks ago

Selected Answer: A

BGP is configured on Tier-0 and not Tier-1.

upvoted 2 times

🗳️ **airmouse1234** 1 year, 3 months ago

Correct Answer is A

upvoted 1 times

🗳️ **KD01** 1 year, 5 months ago

A is correct

upvoted 1 times

🗳️ **ShyamC** 1 year, 6 months ago

Selected Answer: A

No BGP routing possible on T1

upvoted 1 times

🗳️ **rgraghu** 1 year, 7 months ago

Selected Answer: A

Even though in multiple racks we use bgp for Layer-3 vmotion, we never deploy a T1 on each rack, so i will have to correct my answer to A

upvoted 1 times

🗳️ **rgraghu** 1 year, 7 months ago

Selected Answer: D

In a spine-leaf architecture use of BGP in layer3 to move the workloads across the racks seamlessly, and the question is about moving vm's within the cluster which is spread across multiple racks, so its E-W traffic utilizes T1 within the cluster and when it has to go across the racks the traffic utilizes BGP on Data ToR's\
upvoted 2 times

🗨️ **BrianOC** 1 year, 11 months ago

"D" is the wrong answer. BGP is configured on Tier-0 and not Tier-1. "A" is the correct answer.
upvoted 1 times

🗨️ **Me_Loi1** 2 years, 3 months ago

Correct is A
upvoted 2 times

🗨️ **diegof1** 2 years, 3 months ago

A is the correct answer.
upvoted 2 times

🗨️ **VMwareARCHI** 2 years, 5 months ago

A. Deploy all application networks on NSX segments.
upvoted 4 times

An architect is designing a solution for containerization. The solution will include high availability and security using NSX-T Data Center. The architect plans to provide a basic required components list in the Logical Design. Which solution should the architect recommend? (Choose the best answer.)

- A. 3 NSX Managers, 3 virtual NSX Edges, two Tier-0 gateways in Active/Standby, BGP configuration
- B. 2 NSX Managers, 2 virtual NSX Edges, one Tier-0 gateway, BGP configuration and a static route
- C. 3 NSX Managers, 3 virtual NSX Edges, one Tier-0 gateway and a static route and OSPF
- D. 1 NSX Manager, 2 virtual NSX Edges, two Tier-0 gateways in Active/Active, BGP configuration

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **MJ86** Highly Voted 2 years, 5 months ago

B & D can't be the right answer as high availability is intended. to achieve HA in the NSX-T environment, 3 is the minimum number of nodes needed.

C is incorrect as OSPF is only supported with NSX-T 3.1.1 which was released just last week. This exam is based on NSX-T 3.0 and this version don't support ospf at all.

That leaves us with only one option..which is A.
upvoted 19 times

🗳️ **diegof1** Highly Voted 2 years, 3 months ago

A is Correct. I am absolutely agree with MJ86.
upvoted 6 times

🗳️ **outlawww** Most Recent 4 months, 3 weeks ago

Selected Answer: A

NSX manager cluster is min 3 nodes, so B&D are out of the options. OSPF is never recommended in SDDC. Only A left...
upvoted 1 times

🗳️ **Saida001** 8 months, 2 weeks ago

Correct answer is A
upvoted 1 times

🗳️ **vKramer** 2 years, 5 months ago

Correct Answer is A
upvoted 3 times

🗳️ **VMwareARCHI** 2 years, 5 months ago

Answer : A
upvoted 4 times

🗳️ **BlackTai** 2 years, 5 months ago

Answer : A
upvoted 3 times

🗳️ **UdayShanu** 2 years, 5 months ago

B is the write answer
upvoted 1 times

A Solutions Architect is assisting a service provider with designing an NSX-T Data Center solution for these environments:

- ⇒ Virtual Data Center to Virtual Data Center connectivity
- ⇒ Tenant workload on-boarding to Virtual Data Centers.

These requirements must be met:

- ⇒ scalability across 5 data centers
- ⇒ all sites have a latency of 180ms
- ⇒ MTU between sites is 1800
- ⇒ bandwidth is 100Mbps between sites
- ⇒ multi-tenancy

Which two selections should the Solutions Architect propose to the service provider? (Choose two.)

- A. Configure Remote TEPs for stretching network services between Virtual Data Centers.
- B. Utilize SSL VPN for workloads on-boarding from on-premises to Virtual Data Centers.
- C. Configure IPSec VPN for Tenant T0 gateways for Virtual Data Centers connectivity
- D. Configure IPSec VPN for Tenant T1 gateways for Virtual Data Centers connectivity.
- E. Utilize L2 VPN for workloads on-boarding from on-premises to Virtual Data Centers.

Suggested Answer: DE

Community vote distribution



Aymanovitchy Highly Voted 3 years, 4 months ago

C,E

since IPSEC tunnel is done on T0

and L2 VPN to expand Layer 2 to onboard VMs

upvoted 8 times

udo2020 1 year, 10 months ago

IPSec tunnel can also be done on T1

upvoted 1 times

Pheakdey Highly Voted 3 years, 4 months ago

AE is correct

upvoted 8 times

tedybear 2 years, 8 months ago

Rtep requires latency of 150ms between LM,RLM&GM. the latency here exceeds that

upvoted 4 times

AT45816 Most Recent 6 months ago

Selected Answer: CE

CE is correct

upvoted 1 times

Gayan84 8 months ago

Tier-0 gateways are more suitable for connecting multiple data centers with multi-tenancy than Tier-1 gateways. This is because Tier-0 gateways are designed to handle north-south traffic, which is the traffic that flows between the tenants and the external network. Tier-1 gateways are designed to handle east-west traffic, which is the traffic that flows between the tenants themselves.

upvoted 1 times

Gayan84 8 months ago

Selected Answer: CE

Configuring IPSec VPN for Tenant T0 gateways ensures secure connectivity between Virtual Data Centers. T0 gateways are suitable for north-south routing between data centers, aligning with the requirement for connectivity between Virtual Data Centers

upvoted 1 times

Gayan84 8 months ago

Configuring IPsec VPN for Tenant T0 gateways ensures secure connectivity between Virtual Data Centers. T0 gateways are suitable for north-south routing between data centers, aligning with the requirement for connectivity between Virtual Data Centers

upvoted 1 times

🗨️ 👤 **Gayan84** 8 months ago

Selected Answer: CE

Configuring IPsec VPN for Tenant T0 gateways ensures secure connectivity between Virtual Data Centers. T0 gateways are suitable for north-south routing between data centers, aligning with the requirement for connectivity between Virtual Data Centers

upvoted 1 times

🗨️ 👤 **4ourDS** 1 year, 3 months ago

Selected Answer: DE

A high MTU indicates that the packet size due to VPN communication will be large.

upvoted 1 times

🗨️ 👤 **outlawww** 1 year, 4 months ago

Selected Answer: DE

T1 because of multitenancy.

upvoted 2 times

🗨️ 👤 **Alchot** 1 year, 10 months ago

DE is correct

T1 is needed for multitenancy on service provider

Latency is higher than supported so each site will have its own NSX-T

upvoted 1 times

🗨️ 👤 **nick2u** 2 years, 8 months ago

DE, T1 is required for multi-tenancy

upvoted 5 times

🗨️ 👤 **tedybear** 2 years, 8 months ago

CD IPsec VPN is supported for T0 and T1 gateways

NSX-T Data Center supports IPsec Virtual Private Network (IPsec VPN) and Layer 2 VPN (L2 VPN) on an NSX Edge node. IPsec VPN offers site-to-site connectivity between an NSX Edge node and remote sites. With L2 VPN, you can extend your data center by enabling virtual machines to keep their network connectivity across geographical boundaries while using the same IP address.

Onboarding from onprem to remote site requires long distance vMotion with mtu of 150

upvoted 1 times

🗨️ 👤 **DCL202** 2 years, 9 months ago

A and E are not supported, due to 180ms latency: RTT latency is less than or equal to 150 ms, which is required for vMotion to work across two sites. This rules out stretch networking and L2 for migrating workloads.

Answer is B, D.

upvoted 2 times

🗨️ 👤 **Pal68** 2 years, 10 months ago

From the coursework:

Topologies: •

IPsec VPN services are only available on Tier-0 gateways.

• Segments can be connected to either Tier-0 or Tier-1 gateways to use VPN services. Considerations: • Overlapping networks or multitenancy requires multiple T0-GWs. • VPN services are only available on active-standby T0-GWs. • NSX-T Data Center supports site-to-site IPsec VPNs in tunnel mode. • DPDK-accelerated performance.

upvoted 1 times

🗨️ 👤 **Pal68** 2 years, 10 months ago

I.e. Only CE correct

upvoted 1 times

🗨️ 👤 **Arden101** 3 years, 1 month ago

Correct answers - DE

As mentioned, using Federation for five sites is not possible yet. Therefore, we have to setup L2VPN.

IPsec (needed by L2VPN) can be established from T0 as well as T1 (the same rule applies to L2VPN). However, L2VPN is limited (server or client) to one service per gateway, therefore it's not possible to utilize five L2VPN tunnels from the same T0 and we have to deploy five T1s.

Btw. 3V0-42.20 is based on NSX-T 3.0 GA (i.e. 3.0.0) and the Federation limit is three sites (four sites was increased in 3.0.1)

upvoted 6 times

🗨️ 👤 **vkais** 3 years, 2 months ago

Can't be A, that would be NSX-T Federation, which only supports 4 sites. The questions mentions 5 data centers.

upvoted 1 times

🗨️ 👤 **Fornax** 3 years, 2 months ago

It also can't be A, because federation requires 150ms RTT. <https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/installation/GUID-AD369B9D-4ADC-4CE9-B8DC-BB2B47C7BFBF.html>

upvoted 1 times

🗨️ 👤 **VMwareARCHI** 3 years, 5 months ago

D,E correct

upvoted 5 times

🗨️ 👤 **Aymanovitchy** 3 years, 4 months ago

how come deply ipsec on T1

upvoted 1 times

🗨️ 👤 **diegof1** 3 years, 3 months ago

IPsec and Layer 2 VPN are supported on both Tier-1 and Tier-0.

<https://docs.vmware.com/en/VMware-NSX-TData-Center/3.0/administration/GUID-DF689847-252E-451E-84B5-DB507CC010AC.html>

upvoted 1 times

An architect is helping an organization with the Logical Design of a Layer 2 bridging solution.

This information was gathered during the Assessment Phase:

- ⇒ Workloads are running on ESXi hosts.
- ⇒ Workloads are running on KVM hosts.
- ⇒ Workloads on hypervisors should use bridging services.
- ⇒ VLAN 50 is used for Tier-0 uplink connectivity.

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

- A. Create an NSX Edge Bridge Cluster and configure the bridging profile with VLAN 60.
- B. Create an NSX Edge Bridge Cluster and configure the bridging profile with VLAN 50.
- C. Create an ESXi Bridge Cluster and configure the bridging profile with VLAN 50.
- D. Create an ESXi Bridge Cluster and configure the bridging profile with VLAN 60.

Suggested Answer: B

Community vote distribution

A (50%) B (50%)

🗨️ 👤 **vKramer** Highly Voted 3 years, 5 months ago

Answer A is Correct:

When a bridge instance is mixed with L3 routing on T0, the VLAN ID used by the T0 uplinks for the route peer cannot be used by the bridge instances.

upvoted 16 times

🗨️ 👤 **VMwareARCHI** 3 years, 5 months ago

A must be correct

upvoted 5 times

🗨️ 👤 **BrianOC** Highly Voted 2 years, 12 months ago

NSX-T edge bridge does not support multiple VLAN uplinks (81777); https://kb.vmware.com/s/article/81777?lang=en_US

That means that we will be required to use the same VLAN identified for Tier-0 uplink and that's VLAN ID 50. Therefore, B is the correct answer.

upvoted 7 times

🗨️ 👤 **AT45816** Most Recent 6 months ago

Selected Answer: A

The correct answer is A

upvoted 1 times

🗨️ 👤 **Gayan84** 8 months ago

I am Select "A"

While technically possible to configure NSX-T with the same VLAN for Tier-0 uplink connectivity and bridging, it's recommended to use different VLANs to avoid potential conflicts and maintain proper network segmentation. This approach aligns with standard networking practices and helps ensure a clear and robust network configuration. Always refer to the NSX-T documentation and collaborate with your network and NSX administrators for specific organizational guidelines.

upvoted 1 times

🗨️ 👤 **Adil_Sensei** 1 year, 6 months ago

A is Correct

VLAN conflict on the Edge

An edge uplink cannot have two VLAN segments with the same VLAN ID. This limitation is not directly related to the edge bridge, but its consequences will be apparent in some edge bridge related configuration.

When a user configures the attachment of segment S1 to VLAN X on a specific edge uplink, the edge automatically creates a local segment VLAN X on the N-VDS owning the uplink. This automatic creation will fail if there was already a VLAN segment configured with the same ID on this N-VDS. To be more precise, the configuration of the VLAN segment will succeed, but the realization of the segment will fail.

It is possible to hit this limitation when:

⚠️ Trying to configure attachments of several different segments to the same VLAN ID on the same uplink (only the first attachment will be operational.)

Trying to configure the attachment of a segment to a VLAN ID that is already used by other services on this uplink. For example, if there is a Tier 0 router using this VLAN ID to send traffic on the uplink.

upvoted 1 times

  **private0330** 1 year, 11 months ago

Selected Answer: B

It must be the same VLAN in bridging. B is the correct answer.

upvoted 1 times

  **airmouse1234** 2 years, 3 months ago

Answer C is Correct

NSX Edge Bridge Cluster: Do not exist, so AB are wrong.

ESXi Bridge Cluster: Yes, it exists.

upvoted 1 times

  **SkylarZ86** 2 years, 7 months ago

Answer A is correct


This is a caveat from NSX Bridging document found on VMware communities

"VLAN conflicts on the Edge: two services cannot use the same VLAN ID on an edge uplink.

That means that, for example, it is not possible to bridge traffic to a VLAN ID that is already used by a Tier 0 interface on an edge. "


<https://communities.vmware.com/wbsdv95928/attachments/wbsdv95928/4002/70/2/NSX-T%20Bridge%203.1.pdf>

upvoted 3 times

  **BrianOC** 3 years, 2 months ago



B is the correct answer and you can locate it here: <https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.3/com.vmware.nsxt.admin.doc/GUID-7B21DF3D-C9DB-4C10-A32F-B16642266538.html>

upvoted 3 times

  **diegof1** 3 years, 3 months ago

A should be the correct answer.

upvoted 3 times

  **Clicky** 3 years, 4 months ago

My answer is B:

If the edge cluster is used for bridging, all traffic associated with that T1 does not go to a centralized location. However, if an edge cluster has existing services, it can be reused under the bridge profile.

upvoted 4 times

  **tungdt** 3 years, 5 months ago

B is correct

upvoted 3 times

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.

Suggested Answer: D

Community vote distribution

C (67%)

A (33%)

 **BrianOC** Highly Voted 3 years, 2 months ago

D is the correct answer

upvoted 5 times

 **AT45816** Most Recent 6 months ago

Selected Answer: C

Correct Answer: C

upvoted 1 times

 **outlawww** 1 year, 4 months ago

Selected Answer: C

C is totally feasible and satisfy the requirements. You can manage the steering rules with federation.

upvoted 2 times

 **ertin74** 1 year, 9 months ago

Selected Answer: C

Only C applies the requirements

upvoted 1 times

 **Alchot** 1 year, 10 months ago

A is correct.

D is not correct. How NSX Intelligence will help protect external workloads when only works for overlay networks plus is not a security feature.

upvoted 1 times

 **ShyamC** 2 years, 5 months ago

Selected Answer: A

I feel this should be A, with C you have no control over DFW policies on external workloads, however A addresses network introspection for External and DFW for internal workloads which NSX can do

upvoted 2 times

 **sergejszajbaver** 2 years, 6 months ago

NSX Intelligence is not a protection tool. It just monitors flows and let's you click on DFW rules that it suggests based on that.

upvoted 2 times

 **nick2u** 2 years, 8 months ago

I will go with D if I ran into this question during exam. NSX Intelligence is designed to plan for firewall rules and for deployment.


upvoted 1 times

 **tedybear** 2 years, 8 months ago

C, Can't be A as they want to use NGFW both externally and internally,

can't be B as that only covers NS, and especially can't be D, as does it not mention the NGFW. Federation can be a use case for integration with NGFW

upvoted 2 times

  **nick2u** 2 years, 8 months ago

How would Federation be used in this case?

upvoted 1 times