PassQueen
a.k.a dumpsqueen.com

Cisco Data Center Unified Computing Infrastructure Design (DCICUC)

Cisco 500-901

Version Demo

Total Demo Questions: 8

Total Premium Questions: 89

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

You just installed a group ol new Cisco UCS Servers in an ESX Cluster. Now, you want to deploy Nexus 1000V in this virtual environment. Which two statements regarding the scenario are true? (Choose two.)

A. Nexus 1000v VEM will act as the Control Plane, while Nexus 1000v VSM acts as a remote line card

B. Nexus 1000V VSM will act as the Control Plane, while Nexus 1000v VEM acts as a remote line card

C. Nexus 1000v is not supported on ESX

D. Nexus 1000v VSM is deployed in vCentet. Nexus 1000v VEM is deployed in each ESX Host

E. Nexus 10OOv VEM is deployed in vCentet, Nexus 10OOv VSM is deployed in each ESX Host

ANSWER: B D

Explanation:

Cisco Nexus 1000V uses a distributed architecture that separates switch control functions from packet-forwarding functions. “Nexus 1000V VSM will act as the Control Plane, while Nexus 1000v VEM acts as a remote line card” is correct because the Virtual Supervisor Module (VSM) supplies the centralized control plane and management-plane capabilities. It maintains the virtual switch configuration, integrates with vCenter, and applies Cisco NX-OS-style operational and policy management to the distributed switch environment.

“Nexus 1000v VSM is deployed in vCenter. Nexus 1000v VEM is deployed in each ESX Host” is also correct in architectural terms. The VSM is deployed as a virtual appliance managed through the vSphere/vCenter environment, while the Virtual Ethernet Module (VEM) is installed on every ESX/ESXi host that participates in the Nexus 1000V distributed virtual switch. Each VEM performs the local data-plane switching for that host’s virtual machines and operates under the VSM’s centralized control. This design gives administrators a consistent network configuration and policy model across all hosts in the cluster. Cisco describes the VSM as the supervisor/control component and the VEM as the host-resident virtual line-card component in its Nexus 1000V support documentation and Nexus 1000V architecture overview.

QUESTION NO: 2

What is a UniversallyUnique Identifier in the context of Cisco Unified Computing System?

A. A unique identifierused w identify an object of entity on the internet

B. Aunique identifier used in serverprofilesenabling servermobility

C. A unique identifier used in application profiles enabling server mobility.

D. A unique identifier used m serviceprofiles enacting server mobility

ANSWER: D

Explanation:

A unique identifier used m serviceprofiles enacting server mobility is the correct answer. In Cisco UCS, a Universally Unique Identifier (UUID) is a server identity attribute that can be defined through a UUID pool and assigned to a service profile. The service profile carries this identity independently of a particular physical blade or rack server. When the profile is associated with another compatible server, the assigned UUID follows the service profile, helping preserve the server identity that operating systems, hypervisors, and management software can use to recognize the system. This identity abstraction is a key part of UCS service-profile mobility: the compute configuration and identity can be moved to replacement or alternate hardware without requiring the environment to treat it as an entirely new server. Cisco UCS Manager supports UUID pools so administrators can centrally allocate unique UUID values and consistently apply them to service profiles. Cisco describes service profiles as defining a server’s identity and configuration, while its UCS Manager documentation lists UUIDs among the identity information managed for servers. See Cisco’s UCS Manager overview and UCS Manager configuration guides.

QUESTION NO: 3

Which three options can a virtual switch do? (Choose three.)

A. VLAN tagging is done in UCS Manager

B. Provide a trunk port through which virtual machines impose VLAN tags

C. Impose VLAN tags and manage VLAN tag allocation

D. Set IEEE 802.1p Class-of-Service bits

E. Allow physical switch to impose the VLAN tag

ANSWER: B C D

Explanation:

A virtual switch can provide a trunk connection that permits a virtual machine to send traffic with guest-applied VLAN tags. This is commonly called Virtual Guest Tagging: the guest operating system or application is VLAN-aware and applies the appropriate IEEE 802.1Q tag before transmitting frames through its virtual NIC. The virtual switch must support the trunk and allow the permitted VLANs for that port.

A virtual switch can also impose VLAN tags itself and maintain the VLAN-to-virtual-port association. In switch tagging mode, the virtual-switch port or port group is assigned to a VLAN, and the switch applies the VLAN classification and tag as frames leave toward the physical network. This allows administrators to centrally associate workloads and virtual interfaces with VLANs.

Finally, a policy-capable virtual switch, such as Cisco Nexus 1000V, can apply quality-of-service treatment including Layer 2 Class of Service marking. These markings allow downstream infrastructure to classify and prioritize traffic consistently. In practical terms, this is IEEE 802.1p CoS marking in the Ethernet VLAN header rather than an IP-layer marking; IP-layer differentiation is normally expressed with DSCP. Cisco virtual-switch policies use port profiles to associate VLAN and QoS behavior with virtual interfaces. See the Cisco Nexus 1000V configuration guides and VMware's VLAN tagging overview.

QUESTION NO: 4

An administrator has configured an NPV traffic map lot a new server interface, none of the uplinks the traffic map are operational. What is the resulting action?

A. The traffic map will be disregarded and an uplink in the interface VSAN will be selected

B. The default VSAN uplink will be used if operational.

C. The server will remain in a nonoperational state.

D. The server interface will signal an FDISC message to the fabric switch.

ANSWER: C

Explanation:

The server will remain in a nonoperational state. NPV traffic maps provide deterministic uplink selection by associating a server-facing interface with specified NP uplinks. This mapping is intended to constrain the server interface's available forwarding paths to the operational uplinks included in its configured traffic map. If no mapped uplink is operational, NPV cannot establish a valid upstream path for the server interface, so the server-facing interface is held nonoperational rather than being allowed to forward over an arbitrary path. When an eligible mapped uplink becomes operational again, the server interface can be brought up and can proceed with normal fabric login and traffic forwarding through that mapped path. This behavior preserves the path-control and isolation objectives for which traffic maps are configured, including designs that require hosts or storage connections to use particular uplinks. Cisco documents NPV traffic mapping and its role in selecting NP uplinks for server interfaces in the Cisco MDS NX-OS Fabric Configuration Guide. Additional NPV design and operational context is available in Cisco's NX-OS NPV configuration documentation.

QUESTION NO: 5

Which three scalability options are applicable lo Cisco HyperFlex? (Choose three.)

A. Add additional hyperconverged nodes to expand both storage capacity and compute resources

B. Add compute-only nodes to expand compute resources without adding additional storage capacity

C. Add memory-only nodes to expand memory resources without adding storage capacity

D. Add network cards to increase to without adding node capacity

E. Add additional storage to existing nodes to efficiently expand storage capacity without adding compote resources

F. Add clusters to improve overall system performance and manage w*h UCS Central

ANSWER: A B F

Explanation:

Cisco HyperFlex supports scale-out growth at the cluster level by adding hyperconverged nodes. Each such node contributes both CPU and memory resources and local disks that are incorporated into the HyperFlex Data Platform storage cluster, increasing compute and storage together. HyperFlex also supports compute-only nodes, allowing additional virtual-machine compute capacity to be introduced without contributing storage disks to the distributed storage pool. This enables a design to accommodate workloads whose processor and memory requirements grow faster than their storage requirements.

At a broader architectural level, organizations can scale across multiple HyperFlex clusters. Separate clusters can be deployed to provide additional aggregate capacity, workload isolation, or location-specific infrastructure, while centralized UCS management can provide consistent policy and operational visibility across the UCS domain. This is distinct from expanding the resources of one individual HyperFlex cluster, but it is a valid scalability approach for the overall HyperFlex environment. Cisco's HyperFlex installation documentation describes cluster expansion and compute-only node support, and Cisco UCS Central provides centralized management capabilities for UCS infrastructure: Cisco HyperFlex Systems Installation Guide; Cisco UCS Central.

QUESTION NO: 6

Wherecan World Wide Node Name be defined?

A. In neither vHBA template nor SAN connectivitypolicy

B. In vHBA template only

C. In SAN connectivity policy only

D. In both. vHBA template or SAN connectivity policy

ANSWER: C

Explanation:

In Cisco UCS Manager, the World Wide Node Name is a SAN identity attribute that applies to the server's complete set of virtual Fibre Channel HBAs rather than to one individual vHBA. A SAN connectivity policy is therefore the appropriate location to define the WWNN, either by selecting a WWNN pool or by assigning a static value. When the policy is associated with a service profile or service profile template, UCS Manager uses that WWNN as the common node identity for the vHBAs configured by the policy. This design lets storage systems recognize multiple vHBA ports as belonging to the same host node while each vHBA receives its own distinct World Wide Port Name. Cisco's UCS Manager documentation describes SAN connectivity as the policy area in which WWNN assignment is configured and identifies vHBA templates as reusable definitions for the individual vHBA characteristics. This separation supports consistent, portable SAN identity configuration across service profiles. See Cisco's UCS Manager installation and configuration guides and the Cisco UCS Manager documentation for SAN connectivity-policy configuration details.

QUESTION NO: 7

Which three practices are recommended when designing Fibre Channel zoning? (Choose three.)

A. Use a single initiator in each zone.

B. Zone each initiator only to required target ports.

C. Activate the zoneset that contains the required zones.

D. Place all initiators and targets in one common zone.

E. Disable the name server for each VSAN.

ANSWER: A B C

Explanation:

A Fibre Channel zoning design should use **single-initiator zones** to limit the scope of host-to-storage access and simplify troubleshooting. Each initiator should be zoned only to the **required target ports**, which enforces least-privilege storage connectivity. The configured zones must also be included in an **active zoneset** for the zoning configuration to become effective in the VSAN. These practices improve isolation, reduce unnecessary registered state change notifications, and make SAN access control easier to manage. Cisco provides zoning guidance in the Cisco MDS 9000 Series Fabric Configuration Guide.

QUESTION NO: 8

An administrator is setting up a new VSAN and needs to use IVR lo allow access lo backup tape storage targets located in a different VSAN. There is an existing Nexus 5K switch in the SAN fabric that the admin would like to use; however, the switch is not IVR-enabled. How can the Cisco Nexus switch be used?

A. It cannot be used, since only MOS Switches implement IVR

B. As a Transit switch.

C. As a Border switch or IVR NAT gateway

D. As an Edge switch

ANSWER: B

Explanation:

As a Transit switch. is correct because an IVR deployment distinguishes between switches that perform IVR control-plane and routing functions and switches that simply carry the Fibre Channel traffic between those IVR-capable devices. The switch connected to devices requiring cross-VSAN access must participate in the IVR topology in the appropriate edge or border role, while an IVR NAT gateway performs address translation when required by the design. A Nexus 5000 that is not IVR-enabled cannot assume those IVR-specific roles, but it can remain in the SAN path as a transit switch. In this role, it forwards the required traffic over the fabric without needing to originate IVR services or maintain IVR reachability information. This enables the administrator to retain the existing switch infrastructure while placing IVR-capable MDS switches at the points where inter-VSAN routing functionality is required. Cisco's SAN configuration guidance describes IVR topology roles and the use of transit switches in an IVR fabric. See the Cisco MDS NX-OS Inter-VSAN Routing configuration guide and Cisco's MDS configuration-guide library.