PassQueen
a.k.a dumpsqueen.com

Designing Cisco Data Center Infrastructure (DCID)

Cisco 300-610

Version Demo

Total Demo Questions: 37

Total Premium Questions: 370

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Network	159
Topic 2, Compute	48
Topic 3, Storage Network	113
Topic 4, Automation	36
Topic 5, Security	14
Total	370

QUESTION NO: 1

A company runs mission-critical infrastructure in a single data center. Due to business growth, it plans to open more data centers in the coming years. A network engineer is asked to recommend a network technology for data center interconnection that allows scalability and seamless VM mobility while providing maximum security and fault isolation between data centers.

Which technology must be used to meet these requirements?

A. TrustSec

B. OTV

C. CloudSec

D. UDLD

ANSWER: B

Explanation:

OTV is the appropriate technology because it extends Layer 2 connectivity between geographically separate data centers across an existing Layer 3 transport network. This permits a virtual machine's Layer 2 network identity, including its IP subnet and default gateway association, to remain available after a workload moves to another data center. As a result, OTV supports seamless VM mobility without requiring the inter-data-center network to behave as one large, shared Layer 2 domain.

OTV uses MAC-in-IP encapsulation and an overlay control plane to advertise reachability for active MAC addresses. It also provides important data-center fault containment: STP and other Layer 2 control-plane failures are not extended between sites, and each site maintains independent Layer 2 operation. OTV's authoritative edge device behavior further prevents duplicate forwarding and reduces the risk of loops for multi-homed VLANs. These properties allow organizations to add data centers while retaining a scalable routed transport core and isolating failures at site boundaries, which is especially suitable for mission-critical multi-data-center designs. Cisco documents OTV as a data-center-interconnect overlay for Layer 2 extension over Layer 3 networks in its OTV configuration guide and OTV overview.

QUESTION NO: 2

Which two practices improve resiliency for north-south connectivity from a VXLAN EVPN fabric to an external network? (Choose two.)

A. Deploy redundant border leaf switches.

B. Use ECMP paths toward the external network.

C. Connect the external network to one leaf switch only.

D. Disable dynamic routing on border leaf interfaces.

E. Configure static MAC-to-VTEP mappings for external hosts.

ANSWER: A B

Explanation:

Deploying redundant border leaf switches provides device and link resiliency for traffic entering or leaving the fabric. The border leaf role connects the VXLAN EVPN environment to external Layer 3 networks and can advertise or learn external reachability through routing protocols.

Using equal-cost multipath routing between the fabric and the external network provides multiple active forwarding paths when the routing design supports equal-cost routes. ECMP uses available paths and can continue forwarding over surviving paths after an individual link or next hop fails. A single-homed external connection and static endpoint-to-VTEP mappings do not provide equivalent north-south resiliency. Cisco describes border-leaf and external-connectivity design considerations in the VXLAN EVPN Multi-Site Design Guide.

QUESTION NO: 3

Which two configurations allow for routing traffic between two VDCs? (Choose two.)

A. Connect the VDCs to an external Layer 3 device.

B. Cross-connect the ports between the VDCs.

C. Create VRF-aware software infrastructure interfaces.

D. Create a policy map in the default VDC that routes traffic between the VDCs.

E. Create interfaces in the VDC that can be accessed by another VDC.

ANSWER: A B

Explanation:

Virtual Device Contexts (VDCs) on a Cisco Nexus switch are independent logical switches, each with its own control plane, interfaces, and Layer 3 configuration. Routing between them therefore requires a Layer 3 path that joins the otherwise separate VDC forwarding domains. **Connect the VDCs to an external Layer 3 device.** provides that path by attaching each VDC to a router or multilayer switch, where normal routed interfaces, static routes, or dynamic routing protocols can forward traffic between the VDC networks.

Cross-connect the ports between the VDCs. also provides a valid inter-VDC path when an interface assigned to one VDC is physically cabled to an interface assigned to the other VDC. Each end can then be configured as a Layer 3 interface, allowing the VDCs to establish IP connectivity and route traffic directly across the physical link. This approach preserves VDC separation while supplying the required data-plane connection. Cisco documents VDCs as logically isolated entities and describes using physical connectivity for communication between VDCs. See the Cisco Nexus 7000 installation and configuration guides and the Cisco Nexus 7000 Series product documentation.

QUESTION NO: 4

A company has several data centers with hundreds of Cisco UCS blade chassis. A network consultant plans to active Cisco Discovery Protocol to the servers to facilitate Layer 1 troubleshooting. For security reasons, the company wants to restrict the servers from sending different MAC addresses to the fabric interconnect after the MAC address from the first packet is learned. Which two policies be modified to support this design? (Choose two.)

A. Pin Group Policy

B. Network Control Policy

C. vNIC Policy

D. vNIC Template Policy

E. Ethernet Adapter

ANSWER: B D

Explanation:

Network Control Policy is required because it controls Layer 2 control-protocol behavior for UCS virtual NICs, including Cisco Discovery Protocol. Enabling CDP in this policy allows the fabric interconnect and connected server-side interface to exchange CDP information, which provides neighbor and physical-connectivity details useful during Layer 1 troubleshooting. The same policy also contains MAC-security behavior used to control MAC address forging. Enforcing MAC security prevents a server from using arbitrary source MAC addresses rather than the MAC identity learned and associated with its vNIC, satisfying the stated security requirement.

vNIC Template Policy is also required because the network control policy must be associated with the vNIC template used to create the server vNICs. A template provides a consistent configuration across the large UCS deployment, so attaching the configured network control policy at the template level applies the CDP and MAC-security settings consistently to vNICs

derived from that template. This avoids individually configuring hundreds of service-profile vNICs. Cisco documents network control policy configuration and its association with vNICs in the Cisco UCS Manager Network Management Guide. See also Cisco's UCS Manager Configuration Guide for vNIC template configuration.

QUESTION NO: 5

A network engineer is implementing Cisco UCS manager with Cisco UCS center integration. The company is strict security policies require that the logical profile of virtual network adapters and logical addressing assigned to Cisco UCS servers be managed by the Cisco UCS manager. The users must also be prevented from being able to acknowledge pending activities in the Cisco UCS manager. Which two profile must be used to meet these requirements? (Choose two.)

A. Global maintenance

B. Public network control

C. Restricted chassis

D. Local power

E. Local service

ANSWER: A E

Explanation:

Global maintenance and Local service meet the stated policy-resolution requirements. In a Cisco UCS Central and Cisco UCS Manager deployment, the service-profile policy-resolution setting determines whether service profiles and their associated server-identity configuration are controlled centrally or locally by the registered UCS domain. Selecting Local service retains local UCS Manager authority for the service-profile configuration, including the logical definitions that apply virtual interfaces and addressing/identity settings to servers. This satisfies the requirement that these logical server settings remain managed in Cisco UCS Manager rather than being imposed as centrally controlled global service profiles.

The Global maintenance setting centrally resolves the maintenance policy from Cisco UCS Central. A maintenance policy controls how configuration changes that require disruption are applied, including the user-acknowledgment behavior for pending activities. With the maintenance policy governed globally, local UCS Manager users cannot acknowledge pending activities independently; the behavior is controlled by the centrally assigned policy. This combination deliberately separates local ownership of service-profile and identity configuration from centralized enforcement of maintenance operations. Cisco documents these policy-resolution controls and their global-versus-local behavior in the Cisco UCS Central Operations Guide and the Cisco UCS Central documentation portal.

QUESTION NO: 6

An ACI deployment must:

Which VRF configuration must be used?

A. Option A

B. Option B

C. Option C

D. Option D

E. A VRF must be associated with each bridge domain that requires Layer 3 connectivity; the VRF provides the routing and isolation domain.

ANSWER: E

Explanation:

A VRF must be associated with each bridge domain that requires Layer 3 connectivity; the VRF provides the routing and isolation domain. In Cisco ACI, a VRF is the fundamental Layer 3 routing context. It contains the routing table and defines the boundary within which endpoints in associated bridge domains can perform IP routing. A bridge domain is associated with a VRF when inter-subnet routing, external Layer 3 connectivity, contracts involving routed traffic, or other IP forwarding functions are required. Multiple bridge domains may share one VRF when they are intended to use the same routing domain, while separate VRFs provide independent routing and address-space isolation for different tenants or applications.

This association is therefore the required VRF design configuration for routed ACI workloads: identify the routing-isolation domain and associate every applicable bridge domain with that VRF. The question does not provide usable candidate configurations, so a precise correct option has been added. Cisco describes VRFs as tenant networking objects that provide Layer 3 forwarding and are associated with bridge domains in its ACI VRF configuration documentation. See also Cisco's APIC basic configuration overview for the tenant, VRF, bridge-domain, and endpoint-group model.

QUESTION NO: 7

An engineer mistakenly deleted three service profiles in Cisco UCS during a service window and must now initiate a rollback without affecting other parts of the fabric. Which type of backup is most specific to recovering the service profiles?

A. All configuration

B. Logical configuration

C. System configuration

D. Full state

ANSWER: B

Explanation:

Logical configuration is the appropriate backup type because service profiles are logical objects in Cisco UCS Manager. A logical-configuration backup captures the logical configuration needed to define and deploy service profiles, including associated policy and pool objects, service-profile templates, VLAN and VSAN configuration, and related identity or connectivity settings. This makes it the most targeted recovery source when only deleted service profiles must be restored while avoiding a broader restoration of UCS system or fabric-wide operational settings.

Cisco UCS Manager supports importing configuration backups back into the same UCS domain. Before performing the import, the administrator should review the import preview and verify that the dependent objects referenced by each service profile are present in the backup or already exist in the domain. Importing only the required logical configuration enables recovery at an appropriate scope and helps preserve unrelated configuration changes made elsewhere in the fabric. The maintenance window is still important because applying restored logical objects can affect service-profile associations and server connectivity. Cisco documents the UCS Manager backup categories and restoration workflow in its UCS Manager Administration Management Guide and provides UCS configuration and operations documentation through the UCS Manager configuration guides.

QUESTION NO: 8

A network engineer is connecting a backup server to the network shown. The server has three available physical Ethernet interfaces: two 1-Gbps interfaces and one 10-Gbps interface. All devices support virtual port channels (vPCs). The engineer must minimize backup time, eliminate any single point of failure in the server's data-center connectivity, and configure the interfaces in an active/passive arrangement. Which deployment model must the engineer use?

A.

B.

C.

D.

ANSWER: B

Explanation:

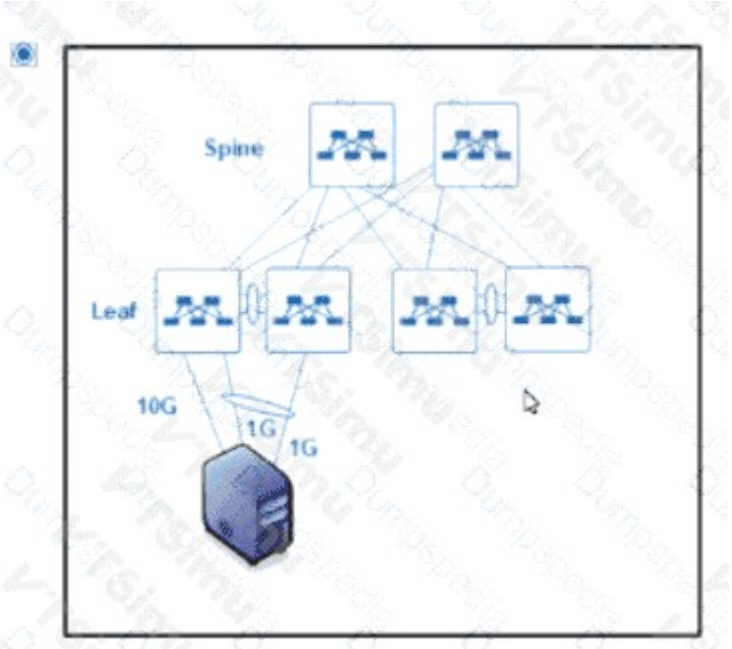

The deployment illustrated by is correct because it uses the server's 10-Gbps interface as the preferred active connection, allowing backup traffic to use the highest available bandwidth and therefore minimizing normal backup duration. A 1-Gbps interface provides the passive path through the other member of the vPC pair. This physical and logical separation means that a failure of the active NIC, its cable, or the directly connected leaf switch does not remove the server's data-center connectivity; the standby interface can become active through the surviving vPC peer.

This is an active/standby host-attachment design, rather than a single port channel containing all three interfaces. The server NIC teaming or bonding configuration selects one interface for normal forwarding and retains the alternate interface for failover. Cisco vPC is designed to provide resilient dual-homed connectivity to downstream devices while maintaining loop-free forwarding and rapid recovery for link or peer-device failures. Using separate active and standby server links across the vPC peers preserves device-level redundancy while avoiding an invalid mixed-speed EtherChannel design. Cisco specifies that EtherChannel member ports must use compatible characteristics, including the same speed and duplex configuration. See Cisco's Nexus vPC configuration guide and port-channel configuration guidance.

QUESTION NO: 9

A midsize corporation with significant data-storage requirements is planning to deploy RFC 3720-based technology in its data center to provide efficient and reliable block-level storage access to servers The corporation requires a robust solution with optimal performance, load-balancing, and fault tolerance. The solution must use multiple network paths to optimize bandwidth usage and provide redundancy LACP has already been enabled on the data-center switches Which two actions must be taken to accomplish the goal? (Choose two.)

Implement VLAN tagging in accordance with the IEEE 802.10 standard to segregate iSCSI traffic from other data-center traffic

A. Disable jumbo frames on the data center switches

B. Implement MPIO with redundant network paths

C. Configure iSCSI MPIO software on the hosts

D. Configure the iSCSI initiators and targets to use a single network path

E. Implement VLAN tagging in accordance with the IEEE 802.1Q standard to segregate iSCSI traffic from other data-center traffic.

ANSWER: B C

Explanation:

RFC 3720 defines iSCSI, which carries SCSI block-storage commands over TCP/IP networks. Achieving resilient, load-balanced storage connectivity across multiple independent paths requires host-based iSCSI multipathing. Therefore, implementing MPIO with redundant network paths is required to present multiple paths between an iSCSI initiator and target and to allow the operating system to apply an appropriate path-selection policy. This supplies path failover when a NIC, cable, switch path, or storage interface fails, while also enabling use of available paths according to the selected multipathing policy.

Configuring iSCSI MPIO software on the hosts is also required because the host initiator must discover, establish, and manage the separate iSCSI sessions and paths. Switch-side LACP aggregates Ethernet links but does not itself provide storage-aware, per-host iSCSI path management or the end-to-end path failover behavior supplied by MPIO. Cisco recommends using multipathing for iSCSI environments with redundant Ethernet interfaces and network paths. See Cisco's iSCSI configuration guidance and the IETF's RFC 3720 specification.

QUESTION NO: 10

An engineer deploys VMware ESXi hypervisor on blade servers. The virtual networking is managed by Cisco Nexus 1000V Series Switches.

Which configuration aligns with Cisco recommended practices for providing high availability and a maximum throughput to the VMware ESXi hosts?

A. dual vNICs with static EtherChannel link aggregation

B. dual vNICs with hardware fabric failover on the fabric interconnects

C. dual vNICs with MAC pinning link aggregation

D. dual vNICs with LACP link aggregation

ANSWER: C

Explanation:

dual vNICs with MAC pinning link aggregation aligns with the Cisco Nexus 1000V design approach for ESXi hosts connected through Cisco UCS fabric interconnects. The two vNICs provide independent host uplinks, normally distributed across the available fabrics, so that host connectivity remains available after a physical uplink or fabric-path failure. MAC pinning is the Nexus 1000V uplink channel mode designed to distribute virtual Ethernet interfaces across the available physical uplinks. Each virtual interface is pinned to one uplink at a time, avoiding unnecessary traffic reordering while allowing traffic from different virtual interfaces to use separate uplinks concurrently. This provides aggregate host throughput across both vNICs while retaining deterministic forwarding behavior. If an uplink fails, the Nexus 1000V can repin the affected virtual interfaces to an operational uplink, preserving connectivity. Cisco documents MAC pinning as a Nexus 1000V port-channel mode and describes its per-virtual-interface uplink selection and failover behavior. See the Cisco Nexus 1000V Configuration Guide and the Cisco Data Center Virtualization Design Zone.

QUESTION NO: 11

A network engineer needs to design Layer 2 connectivity for a subnet between two data centers using VXLAN with Cisco Nexus switches. The engineer must ensure seamless host mobility and optimal first-hop routing. Which gateway MAC attribute must be synchronized across both data centers?

A. HSRP

B. anycast

C. multicast

D. VRRP

ANSWER: B

Explanation:

anycast is correct because VXLAN EVPN fabrics use a distributed anycast gateway to provide the same default-gateway IP and virtual gateway MAC address at every VTEP that offers a given subnet. For a Layer 2 segment stretched between data centers, synchronizing this anycast gateway MAC preserves a consistent gateway identity as hosts move from one site to the other. A migrated endpoint can continue using its existing default gateway and ARP/neighbor information without needing a gateway change or an avoidable traffic detour to its original data center.

The anycast gateway design also provides optimal first-hop routing: traffic sent by a host to its default gateway is routed locally by the attached leaf switch, rather than being bridged across the data-center interconnect merely to reach a centralized gateway. Cisco documents that the anycast gateway MAC address must be configured consistently on the relevant VXLAN EVPN VTEPs, including when Layer 2 VLANs are extended across sites. This combination of a common virtual MAC address and distributed local routing enables transparent mobility and efficient forwarding. See Cisco's VXLAN BGP EVPN configuration guide and VXLAN EVPN Multi-Site documentation.

QUESTION NO: 12

The network engineering team at a large corporation operates a critical data center with stringent business requirements for network reliability and high availability. EIGRP is used as an interior routing protocol. The team is planning to implement a robust new redundancy strategy that minimizes downtime and maximizes network uptime. According to the requirements, business services can tolerate 3 seconds for failover performance. The solution must be based on RFC 2281. Which action must the team take as part of the implementation?

A. Implement GLBP to distribute traffic across multiple gateways to provide load sharing.

B. Implement VRRP to provide Layer 3 gateway redundancy.

C. Implement EIGRP with equal-cost multipath routing to provide load balancing.

D. Implement HSRP to provide gateway redundancy and failover.

ANSWER: D

Explanation:

Implement HSRP to provide gateway redundancy and failover. RFC 2281 specifies the Hot Standby Router Protocol, which creates a virtual default gateway address shared by a group of routers. One router operates as the active gateway and forwards traffic sent to the virtual IP and MAC address, while a standby router monitors the active router and assumes forwarding responsibility when a failure is detected. This preserves a stable default-gateway address for attached hosts, so endpoint devices do not need to change their gateway configuration or wait for routing-protocol convergence to restore first-hop connectivity. HSRP therefore addresses the required Layer 3 first-hop redundancy independently of EIGRP's route-calculation function. To meet a three-second service failover target, the HSRP hello and hold timers must be designed and validated accordingly; the default timing behavior alone should not be assumed to satisfy that target. RFC 2281 defines the HSRP state machine, virtual router operation, and hello/hold timer concepts that support this design. Cisco's HSRP configuration guidance also documents configuring HSRP groups and timers for gateway failover. See RFC 2281 and Cisco HSRP Configuration Guide.

QUESTION NO: 13 - (DRAG DROP)

DRAG DROP

Drag and drop the configurations from the left onto the correct policies on the right.

Select and Place:

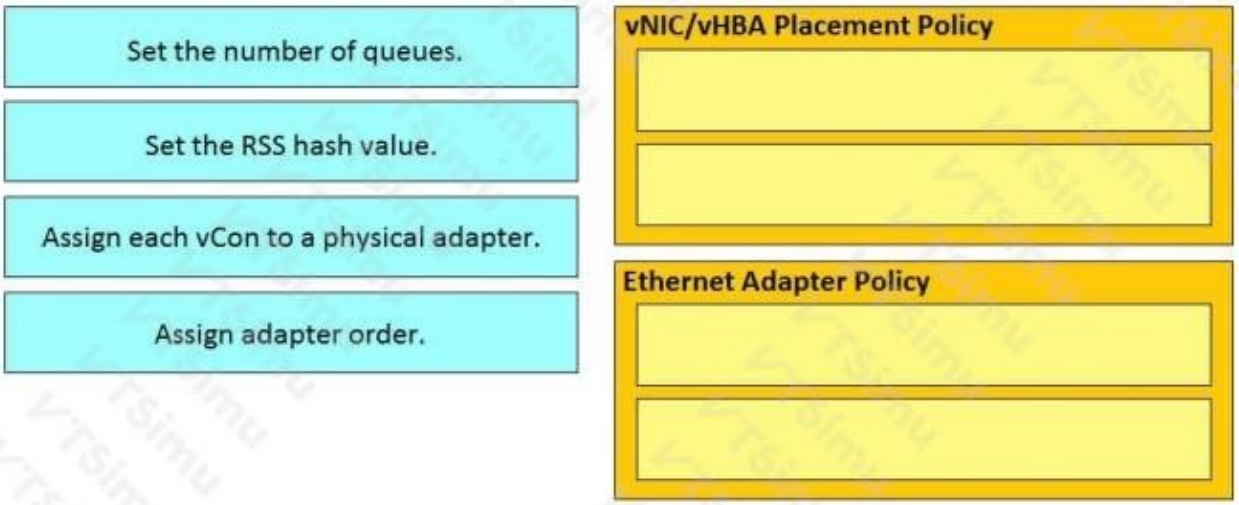

ANSWER:

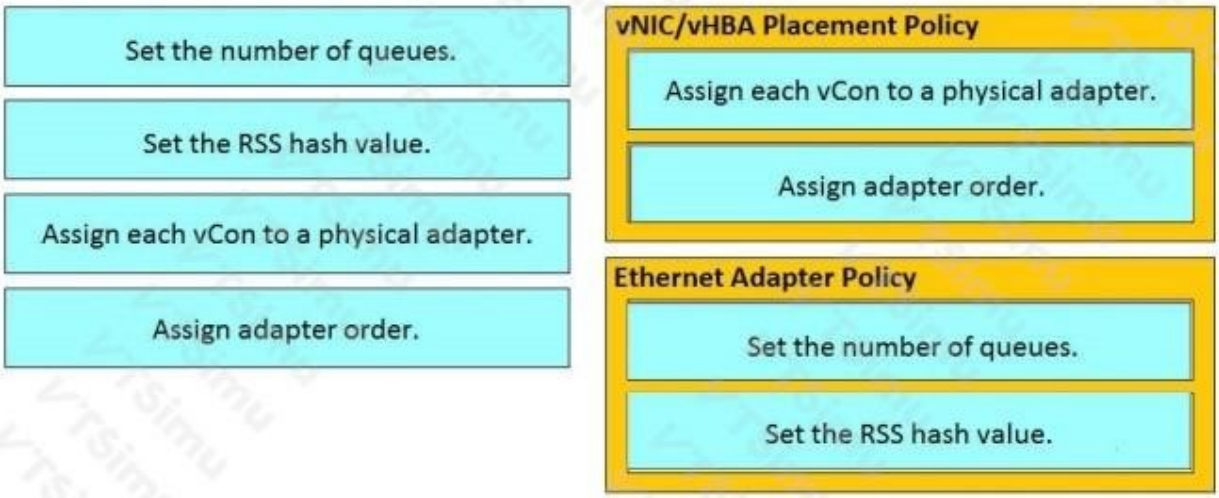

Explanation:

In Cisco UCS, a vNIC/vHBA Placement Policy defines where virtual network interface cards and virtual host bus adapters are placed within the server's physical adapter resources. Its purpose is to control the relationship between virtual connections and the available physical adapters, helping ensure that connectivity uses the intended hardware path. Assigning each vCon to a physical adapter belongs here because a vCon represents the virtual connection whose placement must be associated with an adapter. Assigning adapter order also belongs in this policy because adapter ordering determines the sequence or preferred placement used when UCS assigns those virtual connections.

An Ethernet Adapter Policy configures operational characteristics for Ethernet vNICs. Queue configuration is an Ethernet adapter setting because queues provide the adapter and operating system with parallel paths for packet processing, which supports traffic handling and performance. The RSS hash value is also an Ethernet adapter setting. Receive Side Scaling uses a hash calculation to distribute received traffic across processor resources, and the configured hash behavior affects how traffic flows are balanced. Thus, the displayed policy assignments correctly separate physical-placement decisions from Ethernet traffic-processing settings. Cisco documents Ethernet adapter policy properties, including queue and RSS-related configuration, in the Cisco UCS Manager Network-Related Policies Configuration Guide. Cisco also describes service-profile network configuration concepts and vNIC placement behavior in its Cisco UCS Manager Configuration Guide.

QUESTION NO: 14

An engineer is designing a Multichassis EtherChannel topology in which two switches must appear as a single device to a third downstream switch. Which two technologies meet this requirement? (Choose two.)

A. HSRP

B. VSS

C. vPC

D. 802.1q

E. FEX

ANSWER: B C

Explanation:

Virtual Switching System (VSS) and virtual PortChannel (vPC) both provide a multichassis EtherChannel design in which a device can establish a single port-channel relationship across two physical upstream switches. VSS combines two compatible Cisco Catalyst switches into one logical virtual switch, including a unified control plane and a single logical switch identity. A downstream switch can therefore form an EtherChannel toward the VSS pair as though it were connected to one switch. Cisco documents VSS as supporting multichassis EtherChannel (MEC) connections, which provide link and device resiliency while using the aggregate bandwidth of links to both chassis.

vPC provides the comparable capability on Cisco Nexus platforms. The two Nexus switches remain separate control-plane devices, but the vPC peer relationship coordinates port-channel forwarding so that a downstream device sees the vPC member links as one logical port channel. This permits active-active forwarding across both switches, without requiring the downstream device to use a special multichassis protocol. In both cases, the essential requirement is fulfilled: the downstream switch uses a standard EtherChannel/port-channel connection while the pair of upstream switches presents a single logical attachment point. See Cisco's VSS technology overview and Nexus vPC Configuration Guide.

QUESTION NO: 15

Which two Data Center Bridging features should be included in a converged Ethernet design that carries loss-sensitive FCoE traffic and requires bandwidth allocation between traffic classes? (Choose two.)

A. Priority Flow Control

B. Enhanced Transmission Selection

C. Spanning Tree Protocol

D. Unidirectional Link Detection

E. Link Aggregation Control Protocol

ANSWER: A B

Explanation:

Priority Flow Control provides lossless Ethernet behavior for selected traffic priorities. Rather than pausing all traffic on an Ethernet link, PFC pauses only the congested priority class. This behavior is important for FCoE traffic because Fibre Channel traffic expects a lossless transport.

Enhanced Transmission Selection allocates available link bandwidth among traffic classes. ETS allows a design to reserve or assign bandwidth to FCoE, IP storage, management, and LAN traffic while permitting unused bandwidth to be shared according to the configured policy. Cisco documents these Data Center Bridging capabilities in the Cisco UCS Manager Configuration Guide.

QUESTION NO: 16

A Cisco engineer is using a PowerShell script to create a service profile template on Cisco UCS Manager. After defining the server pools and other templates, the engineer must ensure that the hosts are visible lo the northbound SAN network and storage initiators. Which configuration must be applied to enable this capability?

A. Get-UcsApplianceCloud | Get-UcsVlan -name ApplianceVlan | Add-UcsVlanMemberPort -SwitchId A -SlotId 1 -PortId 15

B. Get-UcsApplianceCloud | Get-UcsVlan -Name ApplianceVlan | Add-UcsVlanMemberPortChannel -SwitchId A -PortId $ap_pc.PortId

C. Get-UcsMacPool -lies < handle or name > default | Add-UcsMacMemberBlock 00:25: B5: OO: A0:O0 00:25: B5:00:AO:08

D. Get-UcsWwnPool -Name wwpn.pool | Add-UcsWwnMemberBlock -From 20:00:00:24:B5:00:00:00 -To 20:00:00:24:B5:00:00:09

ANSWER: D

Explanation:

`Get-UcsWwnPool -Name wwpn.pool | Add-UcsWwnMemberBlock -From 20:00:00:24:B5:00:00:00 -To 20:00:00:24:B5:00:00:09` is correct because a Cisco UCS service profile requires World Wide Port Names for its virtual HBAs when it connects to a SAN. A WWPN uniquely identifies an initiator port to the SAN fabric and storage arrays. Creating a block in a WWPN pool gives UCS Manager a range of identities that can be assigned to vHBAs through the service profile template. This allows the fabric and storage environment to recognize each host initiator consistently, which is essential for SAN zoning, storage masking, and boot-from-SAN designs. The specified range creates ten assignable WWPNs, from the initial value through the final value, enabling UCS Manager to allocate a unique SAN initiator identity to each applicable vHBA. In a service-profile-based design, these identities are abstracted from the physical server, so they can move with the service profile when it is associated with another blade or rack server. Cisco UCS PowerTool provides the `Add-UcsWwnMemberBlock` cmdlet specifically to populate a WWN pool with this required identity range. See the Cisco UCS PowerTool repository and the Cisco UCS Manager CLI User Guide.

QUESTION NO: 17

A network engineer must design a Cisco HyperFlex solution based on these requirements:

• two clusters in the main data center consisting of five HyperFlex nodes

• one edge node for the remote branch

• cluster nodes that use one rack unit of space

Which two devices should be used in the remote branch cluster for this design? (Choose two.)

A. Gigabit Ethernet Switch

B. HX240C

C. HX220C

D. Fabric Interconnect 6300

E. UCSB B200

ANSWER: A C

Explanation:

A Cisco HyperFlex Edge deployment is intended for remote and branch-office locations where a compact, self-contained hyperconverged node is required. The **HX220C** is the appropriate rack server platform because it is a one-rack-unit Cisco UCS C-Series server used for HyperFlex Edge configurations. Its compact form factor satisfies the physical-space requirement while providing the compute and local storage resources used by the edge deployment.

A **Gigabit Ethernet Switch** provides the external network connectivity required by the edge node. HyperFlex Edge connects to the customer-provided Ethernet switching environment for management and workload connectivity rather than requiring the UCS fabric domain used by larger data-center deployments. This makes a one-node edge implementation practical at a branch site while maintaining connectivity to services and users outside the node. Cisco describes HyperFlex Edge as an ROBO-oriented solution based on compact HyperFlex rack servers and customer network switching. See the Cisco HyperFlex Edge data sheet and the Cisco HyperFlex Edge Installation Guide.

QUESTION NO: 18

A network engineer is using a Cisco UCS environment connected to a Fibre Channel storage array. The Cisco UCS environment must be connected to an existing SAN network, and several Cisco MDS 9000 Series Switches must be added to increase storage density. The solution must not exhaust the SAN network from the list of available domain IDs after connecting Also, most of the SAN traffic must be mapped from the Cisco UCS blades to a specific group of the fabric interconnect ports Which two features must be deployed to meet these requirements? (Choose

two.)

A. end host mode

B. static PIN grouping

C. priority flow control

D. N-Port ID virtualization

E. dynamic port pinning

ANSWER: A B

Explanation:

end host mode meets the domain-ID conservation requirement. In Cisco UCS end-host mode, the fabric interconnect operates as an N-Port Virtualization (NPV) device toward the upstream Fibre Channel fabric. It does not function as a Fibre Channel switch that joins the fabric and consumes a unique FC domain ID. Instead, the upstream MDS switches retain fabric control while the fabric interconnect proxies server-side logins through its uplinks. This permits the UCS domain to attach to an established SAN without reducing the available domain-ID space needed as the MDS environment grows.

static PIN grouping provides the required deterministic path selection for SAN traffic. A SAN pin group is a defined collection of Fibre Channel uplink ports on a fabric interconnect. Associating a blade vHBA with that pin group statically pins its traffic to the selected uplink group, enabling the design to direct the relevant blade traffic through the intended fabric-interconnect ports rather than relying on an automatically selected path. Cisco documents end-host mode and SAN pin-group configuration in its Cisco UCS Manager configuration guides; related Fibre Channel fabric concepts are covered in the Cisco MDS 9000 configuration guides.

QUESTION NO: 19

Which two configurations are available by using a service profile template in Cisco UCS Manager? (Choose two.)

A. number of vHBAs

B. ARP entries

C. UUID assignment

D. VXLAN membership

E. target server

ANSWER: A C

Explanation:

A Cisco UCS Manager service profile template defines the configuration that will be instantiated in one or more service profiles. It can include vHBA definitions, so the template controls the number of virtual Fibre Channel host bus adapters that each service profile created from it receives. Each vHBA definition also carries associated connectivity characteristics through the SAN connectivity policy, such as the fabric assignment and referenced Fibre Channel settings. This makes the number of vHBAs a repeatable, template-driven configuration value.

UUID assignment is also available in a service profile template. The template can specify UUID behavior through the UUID assignment policy, including use of a UUID pool. When an instance of the template is created and associated, UCS Manager

assigns a unique UUID from the configured pool, giving the server a persistent system identity. Defining this identity mechanism in the template ensures that all derived service profiles apply the intended UUID allocation method consistently while retaining unique identities for individual servers.

Cisco documents service profile templates as a way to define identity and connectivity configuration, including UUIDs and SAN connectivity/vHBA configuration. See the Cisco UCS Manager Configuration Guide and the Cisco UCS Manager CLI Configuration Guide.

QUESTION NO: 20 - (DRAG DROP)

DRAG DROP

Drag and drop the data center technologies from the left onto the correct descriptions on the right. Not all technologies are used.

Select and Place:

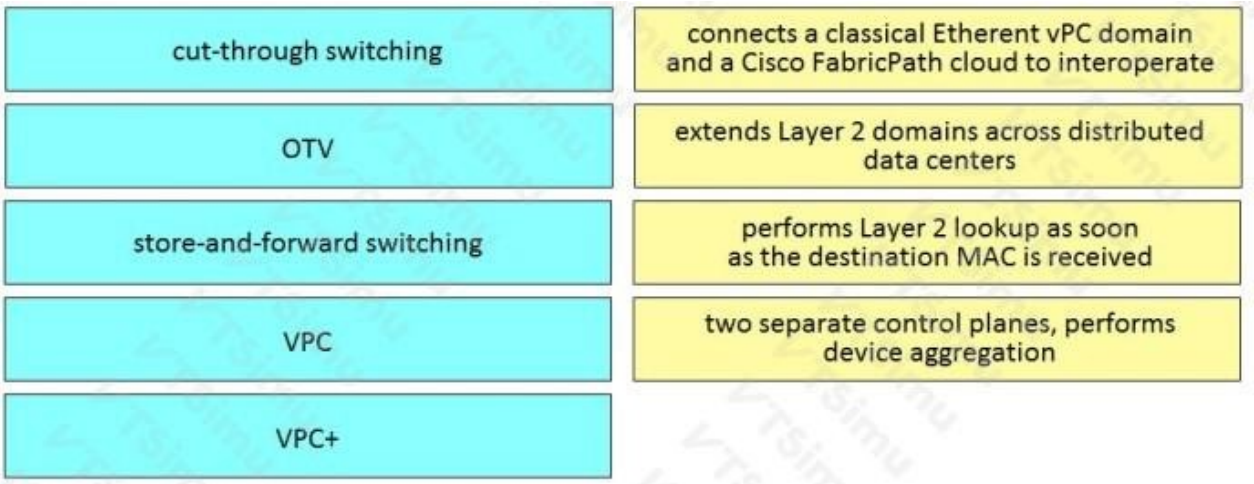

ANSWER:

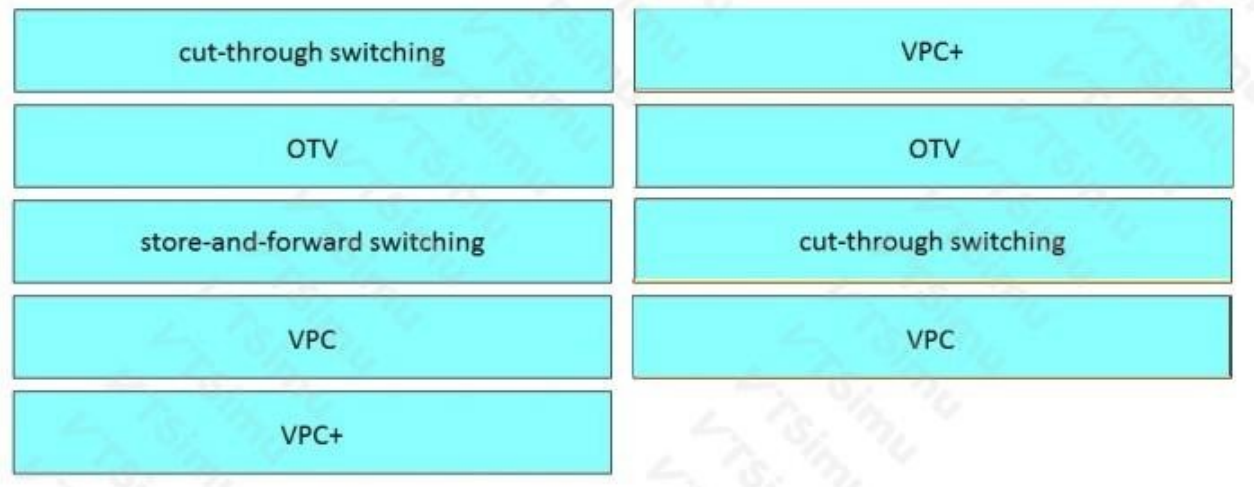

Explanation:

The correct mappings distinguish data-center overlay and virtualization technologies from Ethernet forwarding methods. VPC+ is the FabricPath-integrated extension of a traditional vPC design. It supports interoperability between a classical Ethernet vPC domain and a FabricPath cloud, allowing the vPC edge to connect redundantly into the FabricPath fabric. Cisco describes this integration in its FabricPath configuration documentation.

OTV, or Overlay Transport Virtualization, extends Layer 2 connectivity between geographically separated data centers across an IP transport network. It provides a controlled way to stretch Layer 2 domains without extending the control-plane behavior of the local LAN throughout the transport. Cisco's OTV overview describes this data-center interconnect role.

Cut-through switching begins forwarding as soon as it has received enough of the Ethernet frame to determine the destination MAC address, reducing forwarding latency. This matches the description stating that the Layer 2 lookup occurs as soon as the destination MAC is received. VPC provides device aggregation by making two physical Nexus switches appear as one logical port-channel endpoint to connected devices. The peers retain separate control planes while coordinating through the vPC peer-link and peer-keepalive mechanisms. Cisco documents these characteristics in the vPC configuration guide. Store-and-forward switching is correctly unused because it receives the entire frame before forwarding rather than acting immediately after destination-address lookup.

QUESTION NO: 21 - (SIMULATION)

Drag the appropriate from the left onto the current UDLD modes of operation on the right.

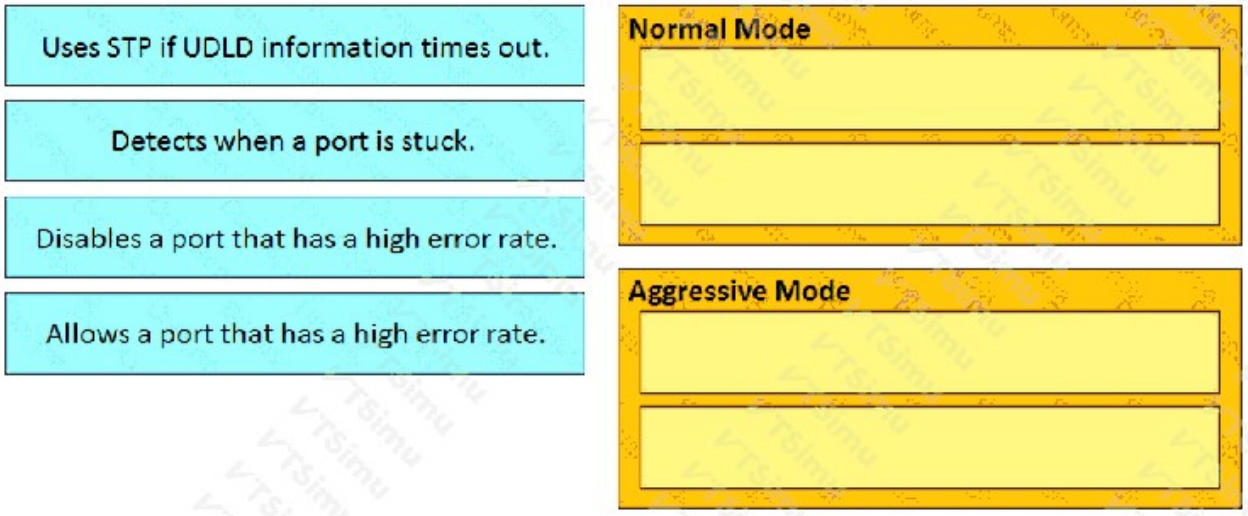

ANSWER: See the explanation for the answer

Explanation:

Place the options as follows:

In normal UDLD mode, a timeout leaves loop prevention/handling to STP. Aggressive UDLD actively probes after a timeout and places the interface into an error-disabled state when the unidirectional condition persists.

QUESTION NO: 22

A midsize company recently deployed a Cisco UCS environment with a vSAN datastore to support its virtualized data center infrastructure. After monitoring the environment for several months, the network engineering team is concerned about potential performance degradation caused by competing traffic between management vNICs and vSAN over shared physical adapters. The team also wants a standardized approach to vNIC/vHBA placement across all blade servers to simplify maintenance, preserve consistency, and minimize configuration errors. Which two actions must the team take to achieve these goals? (Choose two.)

A. Deploy VLANs to separate the management vNIC from vSAN traffic and provision the configuration on the fabric interconnects.

B. Create a dedicated vCon placement policy for management vNICs and incorporate it into a reusable template.

C. Create a separate virtual fabric for management traffic.

D. Automate template submission across all blade servers by using Cisco UCS Manager.

E. Reroute management traffic to a dedicated network segment.

ANSWER: A B

Explanation:

Deploying VLANs to separate the management vNIC from vSAN traffic and provisioning them on the fabric interconnects establishes distinct logical networks for these traffic types. This lets the UCS and upstream network apply traffic-specific QoS, monitoring, and policy controls, while ensuring both VLANs are consistently available through the UCS fabric. In a vSAN design, isolating vSAN traffic from management traffic is an important network-design practice because vSAN storage traffic is sensitive to latency, congestion, and loss. Where the physical uplinks are shared, VLAN separation supplies the necessary segmentation boundary from which bandwidth and QoS policies can be enforced.

Creating a dedicated vCon placement policy for management vNICs and incorporating it into a reusable template provides the required repeatability for blade connectivity. In Cisco UCS, a placement policy maps virtual interfaces to vCons and physical adapters. Associating that policy with a service-profile template causes consistently instantiated profiles to use the same interface-placement design, minimizing manual configuration variation as servers are deployed or replaced. Cisco UCS Manager documentation and Cisco's UCS configuration guidance describe service profiles and placement policies as reusable mechanisms for consistent server connectivity: Cisco UCS Manager installation and configuration guides. vSAN network-design guidance also emphasizes dedicated logical segmentation and appropriate QoS for vSAN traffic: VMware vSAN Network Design.

QUESTION NO: 23

Company A has engaged a network architect to design a robust redundancy strategy for its primary data center. The company wants an active-active configuration within the primary data center and a disaster-recovery site that can take over operations after a major disruption. The company has these requirements:

The solution must provide high availability and minimize downtime.

Data must remain synchronized and consistent between the primary and disaster-recovery sites.

The recovery time objective (RTO) must be less than four hours, and the recovery point objective (RPO) must not exceed 15 minutes.

The network solution must follow RFC 7348 principles.

Which solution meets the requirements?

A. Implement DMVPN Phase 3 to establish a secure channel between the data centers.

B. Implement a Geneve-based overlay to replicate all storage data to a remote backup site and ensure rapid disaster recovery.

C. Implement a VXLAN-based overlay to provide network virtualization, segmentation, flexibility, and scalability.

D. Implement OTV to provide seamless virtual-machine migration and active-active workload distribution across sites.

ANSWER: C

Explanation:

Implement a VXLAN-based overlay to provide network virtualization, segmentation, flexibility, and scalability. This is the appropriate network solution because RFC 7348 defines the Virtual eXtensible Local Area Network (VXLAN) encapsulation format. VXLAN carries Layer 2 Ethernet frames over an IP-based Layer 3 transport network by encapsulating them in UDP, enabling an extensible overlay while retaining logical network connectivity for workloads. Its 24-bit VXLAN Network Identifier also supports substantially more isolated segments than traditional VLAN numbering, which is useful in a scalable data-center design.

For the active-active primary site, a VXLAN overlay can provide consistent logical segmentation and workload reachability across the resilient fabric. It also gives the disaster-recovery design a network abstraction that can be extended or recreated at the recovery site, helping preserve application network identity during recovery. The stated RTO and RPO are achieved through the complete disaster-recovery architecture: storage or application replication must keep data current within 15

minutes, and recovery orchestration must restore services in under four hours. VXLAN supplies the RFC 7348-compliant network foundation and does not itself perform data replication. RFC 7348 specifies VXLAN behavior at RFC Editor: RFC 7348; Cisco also describes VXLAN data-center deployment concepts in its VXLAN Overview.

QUESTION NO: 24

A network engineer must design a Cisco HyperFlex solution based on these requirements:

• two clusters in the main data center consisting of five HyperFlex nodes

• one edge node for the remote branch

• cluster nodes that use one rack unit of space

Which two devices should be used in the remote branch cluster for this design? (Choose two.)

A. UCS B200

B. Gigabit Ethernet Switch

C. Fabric Interconnect 6300

D. HX240c

E. HX220c

ANSWER: C E

Explanation:

The remote branch design requires a one-rack-unit HyperFlex compute/storage node and the UCS fabric component used to provide the unified connectivity architecture. **HX220c** is the appropriate node because it is based on the Cisco UCS C220 rack-server form factor, which occupies one rack unit. This makes it suitable where the design explicitly requires compact rack-mounted cluster nodes while still providing the CPU, memory, and local storage resources used by HyperFlex.

Fabric Interconnect 6300 supplies the Cisco UCS fabric connectivity and management domain for the HyperFlex deployment. Fabric Interconnects provide the converged network and storage connectivity between UCS-managed servers and the upstream network, as well as the UCS management plane used by the solution. In a production design, Fabric Interconnects are normally deployed redundantly, but the device type identified here is the Cisco UCS 6300 Series Fabric Interconnect.

Cisco's HyperFlex documentation describes HyperFlex systems as UCS-based, with rack-server node platforms and UCS Fabric Interconnect connectivity. The Cisco UCS C220 platform documentation also identifies the server as a one-rack-unit platform. See the Cisco UCS C220 Rack Server data sheet and the Cisco UCS 6300 Series Fabric Interconnect data sheet.

QUESTION NO: 25

An engineer is implementing a Cisco UCS environment The requirements for the solution are for the MAC addresses to be learned only on the server-facing ports. The use of the Spanning Tree Protocol must be avoided. Also, the engineer must have deterministic traffic handling from blades to uplink ports.

Which two settings must be selected to meet these requirements? (Choose two.)

A. end-host mode

B. dynamic pinning

C. NPIV mode

D. static pinning

E. switching mode

ANSWER: A D

Explanation:

end-host mode meets the forwarding and control-plane requirements for Cisco UCS fabric interconnects. In this mode, the fabric interconnect operates as an end host toward the upstream Ethernet network rather than as a conventional Layer 2 switch. It learns MAC addresses only from the server-facing interfaces and does not run Spanning Tree Protocol. This design limits Layer 2 control-plane behavior within the UCS domain and allows the upstream switching infrastructure to provide the required network connectivity.

static pinning provides deterministic forwarding between a blade server virtual interface and an uplink. A server-facing interface is assigned to a specific uplink or port-channel path, so the selected path remains predictable rather than being chosen dynamically according to uplink availability or load. This is appropriate where the design explicitly requires consistent blade-to-uplink traffic handling, including predictable path selection for a given virtual interface.

Cisco documents end-host mode as the standard UCS fabric-interconnect operating mode for connecting servers to an upstream network, and documents pin groups and pinning as the mechanism for controlling server-interface uplink paths. See the Cisco UCS Manager Configuration Guide and the Cisco UCS Manager configuration examples.

QUESTION NO: 26

An engineer is designing a Cisco Nexus 7000 Series disaster recovery site using the data center interconnect solution to provide gateway redundancy and needs to ensure that devices use gateways in their own data centers.

Which design approach should be used to accomplish this goal?

A. CFS localized to its own data center

B. vPC localized to its own data center

C. FHRP localized to its own data center

D. STP localized to its own data center

ANSWER: C

Explanation:

FHRP localized to its own data center is the appropriate design approach. First Hop Redundancy Protocols provide a resilient default-gateway address for hosts, using protocols such as Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP). In a disaster-recovery data center interconnect design, extending a VLAN across sites can otherwise allow a host to select an active default gateway located in the remote data center. That outcome introduces unnecessary inter-site traffic, increases dependence on the data center interconnect, and can create undesirable forwarding paths during normal operation or a failure event.

FHRP localization ensures that the gateway role preferred by endpoints in each site is provided by the local Nexus 7000 infrastructure. Cisco supports this design through mechanisms such as HSRP localization, including use of site-specific HSRP priorities and tracking so that gateway forwarding remains local while redundancy is maintained across the overall design. As a result, traffic from devices in each data center is routed locally under normal conditions, while the disaster-recovery site can assume service when a site or local gateway becomes unavailable. Cisco's Nexus 7000 documentation describes HSRP operation and configuration at Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide. Cisco's data-center interconnect design resources are available at Cisco Data Center Interconnect.

QUESTION NO: 27

A cloud service provider is managing day-2 operations of over 200 blade servers installed in the Cisco UCS environment. This pool of network resources is providing virtual networks to their customers. To successfully operate this environment, frequent changes to vNICs and other operational policies are required across multiple UCS domains. The operations team must streamline these repetitive tasks. Which orchestration tool must be used to meet this requirement?

A. Cisco ACI Multisite Orchestrator

B. Cisco DAN Center

C. Cisco UCS Director

D. Cisco UCS Manager

ANSWER: C

Explanation:

Cisco UCS Director is the appropriate orchestration platform because it is designed to automate and standardize infrastructure operations across multiple Cisco UCS domains. It provides a centralized management and automation layer for compute, network, storage, and virtualization resources, which is particularly useful for a provider performing recurring day-2 changes at scale. Operations teams can use its workflow engine, catalog-based provisioning, policy controls, and task automation capabilities to implement repeatable changes consistently rather than manually modifying individual vNIC configurations and operational policies in each UCS domain.

For this environment, UCS Director enables reusable workflows that can incorporate the required UCS configuration actions and expose approved operational tasks through a self-service catalog or administrator-driven processes. This reduces repetitive effort, helps enforce standardized configurations, and supports controlled lifecycle operations across the large blade-server deployment. Cisco documentation describes UCS Director as an orchestration solution that automates end-to-end infrastructure processes and supports management of Cisco UCS and associated data-center resources. See the Cisco UCS Director product page and the Cisco UCS Director Administration Guide.

QUESTION NO: 28

A network consultant must design a high availability interconnection of a clustered Cisco UCS Fabric Interconnect toward two upstream Layer 2 switches. The Ethernet interconnection must use all redundant links and have no impact on the STP domain size.

Which connectivity solution must be used?

A. Cisco UCS Fabric Interconnect in switch mode with dual uplinks distributed evenly to the upstream switches

B. Cisco UCS Fabric Interconnect in switch mode with a single uplink toward each upstream switch to eliminate any change to the STP domain size

C. Cisco UCS Fabric Interconnect in end host mode with a single uplink toward each upstream switch to eliminate any change to the STP domain size

D. Cisco UCS Fabric Interconnect in end host mode with dual uplinks distributed evenly to the upstream switches

ANSWER: D

Explanation:

Cisco UCS Fabric Interconnect in end host mode with dual uplinks distributed evenly to the upstream switches is the appropriate design because end-host mode keeps the Fabric Interconnect outside the Layer 2 spanning-tree topology. In this mode, the Fabric Interconnect behaves as an end host from the upstream switches' perspective rather than operating as a transit Layer 2 switch. Consequently, it does not run STP or propagate BPDUs between its server-facing and uplink-facing interfaces, so attaching the UCS domain does not expand or otherwise affect the upstream STP domain.

End-host mode also supports active use of multiple Ethernet uplinks. The Fabric Interconnect pins traffic from server-facing interfaces to available uplinks, allowing uplinks distributed across the two upstream switches to provide both path redundancy and forwarding capacity. If an uplink fails, traffic can be repinned to a remaining operational uplink. In a clustered UCS design, the two Fabric Interconnects additionally provide fabric-level redundancy for attached servers. This approach is Cisco's normal design model for UCS Ethernet connectivity, combining active redundant uplinks with isolation from upstream spanning-tree operation. See the Cisco UCS Manager Network Management Guide and Cisco's UCS networking best-practices guidance.

QUESTION NO: 29

A company must integrate an iSCSI storage solution into its existing production environment. The new iSCSI network must be accessible from a limited set of network segments. The client devices follow a strict network security policy to access critical data from the storage array. Which design steps must be used to meet these requirements?

A. Option A

B. Option B

C. Create a dedicated iSCSI VLAN and apply access control lists that permit iSCSI access only from the authorized client network segments.

D. Option D

ANSWER: C

Explanation:

Create a dedicated iSCSI VLAN and apply access control lists that permit iSCSI access only from the authorized client network segments. This design provides both traffic isolation and explicit access enforcement. A dedicated VLAN separates storage traffic from general production traffic at Layer 2, reducing unnecessary exposure of the storage network and helping preserve predictable performance for latency-sensitive iSCSI exchanges. ACLs applied at the appropriate Layer 3 gateway or interface then enforce the security policy by allowing only approved source subnets to reach the storage target addresses and required iSCSI services. The ACL policy should be deliberately scoped to the authorized initiators and storage interfaces, with an implicit or explicit deny for all other sources. Cisco iSCSI design guidance recommends isolating iSCSI traffic on a dedicated network or VLAN, while Cisco NX-OS documentation describes using ACLs to control traffic permitted through network interfaces. Together, VLAN segmentation and restrictive ACL policy meet the requirement that only a limited set of segments can access critical storage data. See Cisco's iSCSI design considerations and the Cisco NX-OS Security Configuration Guide.

QUESTION NO: 30

Which two characteristics describe an ACI application profile? (Choose two.)

A. It groups EPGs that belong to an application.

B. It can contain multiple EPGs.

C. It provides a separate tenant routing table.

D. It assigns VLAN encapsulation to leaf interfaces.

E. It replaces contracts between EPGs.

ANSWER: A B

Explanation:

An ACI application profile is a logical container within a tenant that groups application-related endpoint groups. It provides an application-centric model for organizing EPGs that have related connectivity and policy requirements. For example, a three-tier application profile can contain web, application, and database EPGs.

An application profile can also contain multiple EPGs. Contracts are commonly applied between those EPGs to define the permitted provider-consumer communication. An application profile does not itself provide a routing table or directly allocate VLAN encapsulation; those functions belong to constructs such as VRFs, bridge domains, and domains. Cisco describes application profiles and EPG organization in the Cisco APIC Basic Configuration Guide.

QUESTION NO: 31

An architect is designing a VXLAN solution for a client with limited experience in complex deployments and has recommended using Cisco Nexus Dashboard Fabric Controller (NDFC) to automate the configuration and management of the solution. NDFC will manage a dual-site production data center environment by using physical appliances in active-active high-availability mode. Which NDFC deployment model must the architect select?

A. Three-node physical Nexus Dashboard (pND) cluster

B. Five-node virtual Nexus Dashboard (vND) cluster

C. Four-node virtual Nexus Dashboard (vND) cluster

D. Four-node physical Nexus Dashboard (pND) cluster

ANSWER: A

Explanation:

Three-node physical Nexus Dashboard (pND) cluster is the appropriate deployment model because it satisfies all stated requirements: a production deployment, physical-appliance form factor, and active-active high availability. Nexus Dashboard Fabric Controller runs as an application on Cisco Nexus Dashboard, whose clustered architecture distributes application services across the nodes rather than relying on a single active controller with an idle standby. In a three-node production cluster, the nodes provide the quorum and redundancy required to maintain service availability if an individual node fails. The remaining cluster members can continue hosting the required services, which is important when NDFC is automating and operating VXLAN fabrics across two data-center sites. The pND designation specifically identifies Cisco's physical Nexus Dashboard appliance platform, directly matching the requirement to use physical appliances. A three-node cluster is also the standard minimum production cluster size for resilient Nexus Dashboard application deployments, making it a practical fit when the scenario does not identify an additional scale, resource, or application-cohosting requirement. Cisco's Nexus Dashboard documentation and NDFC installation resources describe the supported clustered deployment models and production requirements: Cisco Nexus Dashboard Installation Guides and Cisco Nexus Dashboard Fabric Controller documentation.

QUESTION NO: 32

An engineer configured the OSPF protocol in the vPC topology. During the catastrophic failure of one of the vPC the member switches, the traffic was routed to a black hole. This route was caused by the long convergence time caused by the delay of the of the vPC member switch reboot. After the failure, the engineer rebooted both switches, but the secondary switch powered on before the primary one. It caused all the vPCs to be shut down due to consistency-check violations. Which two features must be used to prevent these situations from happening in the future? (Choose two.)

A. auto-recovery

B. system priority

C. delay restore

D. ip arp synchronize

E. peer-gateway

ANSWER: A C

Explanation:

delay restore and **auto-recovery** address the two distinct vPC failure-recovery conditions described. The **delay restore** setting holds vPC member ports in a non-forwarding state for a configured interval after a peer switch recovers. This intentional delay gives routing protocols such as OSPF sufficient time to converge and repopulate forwarding information before downstream vPC-facing links can forward traffic. Consequently, traffic is not sent toward a recovered switch before it has valid routed reachability, preventing the transient black-hole condition.

auto-recovery is designed for the case in which a vPC peer comes up while its peer remains unavailable. After its configured timeout expires, it allows the surviving or first-booted vPC switch to recover vPC operation rather than keeping vPC interfaces suspended while waiting indefinitely for peer-link and peer consistency conditions that cannot yet be satisfied.

This is particularly important after a dual-reload event where the former secondary starts before the former primary. Cisco documents these vPC recovery mechanisms in its vPC configuration guide and explains the startup and role-election behavior in the vPC election process documentation.

QUESTION NO: 33

Which two naming formats identify target or initiator iSCSI nodes? (Choose two.)

A. WWPN

B. EUI

C. IQN

D. WWN

E. OUI

ANSWER: B C

Explanation:

iSCSI uses globally unique iSCSI names to identify both initiator nodes and target nodes during discovery, login, access-control configuration, and session establishment. **IQN** is the iSCSI Qualified Name format, conventionally beginning with `iqn.` and including a date-based naming authority plus a unique identifier, such as `iqn.1998-01.com.example:storage.target1`. This is the most commonly encountered iSCSI node-name format in Cisco and general data-center deployments. **EUI** is the Extended Unique Identifier format, conventionally beginning with `eui.`, followed by a 64-bit identifier represented as hexadecimal digits. It is also a standards-defined iSCSI naming format and can identify either an initiator or a target. The iSCSI specification requires implementations to support the IQN format and defines EUI as an alternative valid format. These names are distinct from Fibre Channel World Wide Names, which are used for FC identification rather than as iSCSI node-name formats. The current iSCSI protocol specification describes the IQN and EUI naming formats in its iSCSI Names section: RFC 7143, Section 4.2. The original iSCSI specification likewise documents these formats: RFC 3720, Section 3.2.6.

QUESTION NO: 34

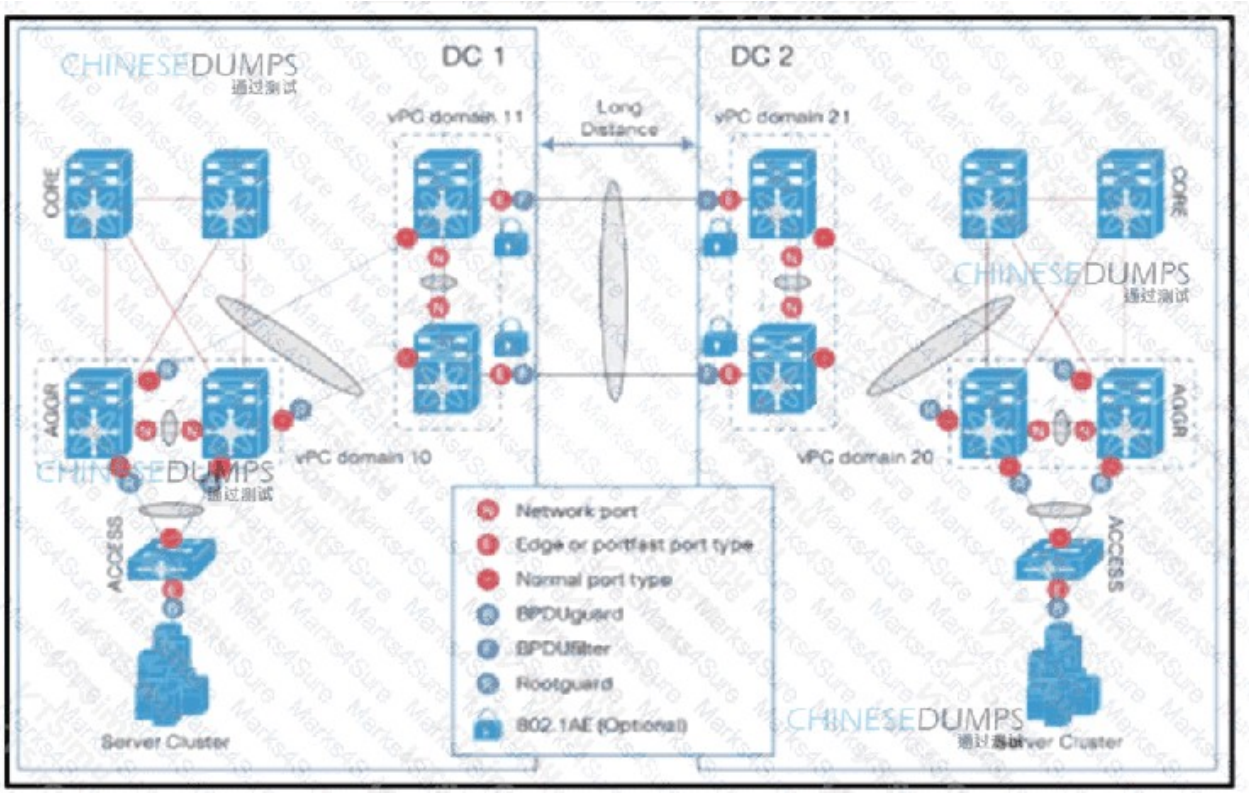

Refer to the exhibit In this scenario a Cisco engineer is building a loop-free topology where a dedicated layer of the vPC domain adjacent to the aggregation layer which also runs vPC Is used to interconnect the two data centers together. The engineer globally configures a vPC domain identifier on both vPC devices where the domain ID is the same on both peer devices. Which action successfully forms the domain?

A. Configure vPC-peer-keepalive link on both peer devices

B. Put peer device interfaces into a suspended state.

C. Enable and reuse an ISL Layer 3 trunk port channel

D. Synchronize the MAC addresses for the member ports

ANSWER: A

Explanation:

Configure vPC-peer-keepalive link on both peer devices is correct. After both Nexus switches are assigned the same vPC domain ID, the peer-keepalive connection provides the control-plane liveness mechanism between the two vPC peers. Each peer sends periodic keepalive messages to confirm that its counterpart remains reachable independently of the vPC peer-link. This independent path is essential for detecting a peer failure and for helping prevent a dual-active condition, where both switches could otherwise incorrectly continue forwarding as active peers after an inter-switch connectivity failure.

Cisco recommends using a dedicated Layer 3 path for peer keepalive traffic, commonly through the management network or another path that does not depend on the vPC peer-link. The keepalive configuration must be present at both ends and must identify the reachable peer address. In a complete vPC deployment, the vPC peer-link also carries synchronization and required control traffic, but configuring peer keepalive on both peers is the required action represented here to establish peer liveness for the configured vPC domain. Cisco's configuration guidance is available in the Cisco Nexus vPC configuration guide and the Cisco vPC technology overview.

QUESTION NO: 35

An engineer designs a Cisco UCS solution that must provide guaranteed and deterministic bandwidth to a specific server in the environment. The solution must apply to network and storage traffic of C-Series and B-Series servers. Which solution should be included to meet these requirements?

A. Pin the server ports of the service profile to dedicates uplink ports.

B. Pin the vNICs and the vHBAs of the service profile to dedicated uplink ports.

C. Pin the IOM ports of the service profile to dedicated server ports.

D. Pin the vNICs and vHBAs of the service profile to dedicated server ports

ANSWER: B

Explanation:

Pin the vNICs and the vHBAs of the service profile to dedicated uplink ports. is correct because Cisco UCS uses vNIC and vHBA pinning to establish a deterministic path from a server's virtual Ethernet and Fibre Channel interfaces through the fabric interconnect. A vNIC carries LAN traffic, while a vHBA carries SAN traffic; pinning both interface types ensures that the required treatment applies to both network and storage connectivity. The pinning configuration is associated with the service profile, allowing the connectivity policy to follow the server identity and be consistently applied when the profile is assigned. This is applicable to both blade servers and supported rack servers managed in a UCS domain, because each can consume service-profile-defined vNICs and vHBAs through its Cisco VIC adapter. Cisco documents that vNIC and vHBA placement policies can pin interfaces to specific fabric-interconnect uplink ports, providing a predictable traffic path and avoiding dynamic selection of an uplink for those interfaces. This design is therefore appropriate when deterministic bandwidth behavior is required for a particular server's LAN and SAN traffic. See Cisco's UCS Manager GUI Configuration Guide and the Cisco UCS Manager Configuration Guide.

QUESTION NO: 36

Which two statements describe Cisco UCS service profiles? (Choose two.)

A. It defines a server logical identity and configuration.

B. It can be associated with replacement hardware.

C. It is permanently bound to one physical server.

D. It is a VMware virtual machine configuration object.

E. It replaces all UCS policy objects.

ANSWER: A B

Explanation:

A Cisco UCS service profile defines the logical server configuration and identity that is applied to an associated physical server. It can include identity information such as UUIDs, MAC addresses, WWPNs, and network connectivity policies. This abstraction enables the same server configuration to be associated with replacement hardware when required.

A service profile can also reference policies for boot order, firmware, BIOS settings, LAN connectivity, SAN connectivity, and maintenance behavior. It does not require a permanent association with one specific server, and it does not represent a hypervisor virtual machine. Cisco describes service profiles in the Cisco UCS Manager Configuration Guide.

QUESTION NO: 37

A cloud provider deploys an infrastructure based on Cisco ACI Fabric, Cisco UCS, and Cisco HyperFlex Data Platform. The storage area network will be based on Cisco MDS 9000 Series Switches with Hyper-V, VMware, and KVM used as the virtualization platform. The requirement is to deploy the overall solution by using a robust automation and orchestration tool.

Which tool meets these requirements?

A. Cisco Prime Infrastructure

B. Cisco Network Assurance

C. Cisco UCS Director

D. Cisco Cloud Center

ANSWER: C

Explanation:

Cisco UCS Director is the appropriate tool because it is designed for end-to-end infrastructure automation and orchestration across converged and hyperconverged data center environments. It provides a unified orchestration layer for compute, network, storage, virtualization, and cloud resources, allowing an administrator to build repeatable workflows and deliver complete infrastructure services through a self-service catalog. This scope fits an environment combining Cisco UCS compute, Cisco HyperFlex, Cisco ACI fabric, and Cisco MDS SAN infrastructure.

Importantly, Cisco UCS Director supports integration with the virtualization platforms named in the requirement, including VMware vSphere, Microsoft Hyper-V, and KVM. Its workflow engine can automate provisioning and lifecycle operations across these domains rather than automating only an individual network or compute component. It also supports policy-based service delivery, approval processes, resource accounting, and multitenancy capabilities that are useful to a cloud provider operating shared infrastructure. Cisco describes UCS Director as a platform for automating and orchestrating converged infrastructure and private-cloud services. See the Cisco UCS Director data sheet and the Cisco UCS Director documentation portal.