Exam Questions 2V0-13.24

VMware Cloud Foundation 5.2 Architect

https://www.2passeasy.com/dumps/2V0-13.24/

NEW QUESTION 1
An architect has been tasked with reviewing a VMware Cloud Foundation design document. Observe the following requirements:
• REQ01: The solution must provide the ability to request new tenant creation with multi- site and different size options.
• REQ02: The solution must provide the capability to monitor the software-defined data center for capacity and performance.
• REQ03: The solution must provide the ability to generate reports with customized metrics to meet business requests.
• REQ04: The solution should report all capacity planning components (such as current capacity usage monthly and annual usage growth).
• REQ05: The solution must provide the ability to provision new virtual machines from predefined templates.
• REQ06: The solution must provide a self-service catalog for end-users to consume services.
Observe the following design decisions:
• DD01: There will be a centralized deployment of Aria Operations Management.
• DD02: There will be customized super-metrics based on existing metrics.
Based on the stated requirements and design decisions, which three requirements does this design decision satisfy? (Choose three.)

A. REQ05
B. REQ01
C. REQ06
D. REQ04
E. REQ03
F. REQ02

Answer: DEF

Explanation:
Reference:VMware Aria Operations 8.10 Administration Guide, Capacity and Custom Metrics; VMware Cloud Foundation 5.2 Architect Study Guide, Monitoring Solutions.

NEW QUESTION 2
A customer defined a requirement for the newly deployed SDDC infrastructure which will host one of the applications responsible for video streaming. Application will run as part of a VI Workload Domain with dedicated NSX instance and virtual machines. Required network throughput was defined as 250 Gb/s. Additionally, the application should provide the lowest possible latency. Which design decision should be recommended by an architect for the NSX Edge deployment?

A. Deploy 2 NSX Edges using NSX console and add to Edge cluster created in SDDC Manager.
B. Deploy 4 extra large edges using vCenter Server console.
C. Deploy NSX bare-metal Edges and create Edge Cluster using NSX console.
D. Deploy 2 large NSX Edges using SDDC Manager.

Answer: C

Explanation:
Reference:NSX-T 3.2 Reference Design Guide, Edge Node Performance; VMware Cloud Foundation 5.2 Networking Guide, NSX Edge Deployment Options.

NEW QUESTION 3
The following are a set of design decisions related to networking: DD01: Set NSX Distributed Firewall (DFW) to block all traffic by default.
DD02: Use VLANs to separate physical network functions.
DD03: Connect the management interface eth0 of each NSX Edge node to VLAN 100. DD04: Deploy 2x 64-port Cisco Nexus 9300 switches for top-of-rack ESXi host
connectivity.
Which design decision would an architect include in the logical design?

A. DD04
B. DD01
C. DD03
D. DD02

Answer: D

Explanation:
In VMware Cloud Foundation (VCF) 5.2, the logical design outlines high-level architectural decisions that define the system??s structure and behavior, distinct from physical or operational details, as per theVCF 5.2 Design Guide. Networking decisions in the logical design focus on connectivity frameworks, security policies, and scalability. Let??s evaluate each:
Option A: DD04 - Deploy 2x 64-port Cisco Nexus 9300 switches for top-of-rack ESXi host connectivityThis specifies physical hardware (switch model, port count), which belongs in the physical design (e.g., BOM, rack layout). TheVCF 5.2 Architectural Guide classifies hardware selections as physical, not logical, unless they dictate architecture, which isn??t the case here.
Option B: DD01 - Set NSX Distributed Firewall (DFW) to block all traffic by default This is a specific security policy within NSX DFW, defining traffic behavior. While critical, it??s an implementation detail (e.g., rule configuration), not a high-level logical design decision. TheVCF 5.2 Networking Guideplaces DFW rules in detailed design, not the logical overview.
Option C: DD03 - Connect the management interface eth0 of each NSX Edge node to VLAN 100This details a specific interface-to-VLAN mapping, an operational or physical configuration. TheVCF 5.2 Networking Guidetreats such specifics as implementation-level decisions, not logical design elements.
Option D: DD02 - Use VLANs to separate physical network functionsUsing VLANs to segment network functions (e.g., management, vMotion, vSAN) is a foundational networking architecture decision in VCF. It defines the logical separation of traffic types, enhancing security and scalability. TheVCF 5.2 Architectural Guideincludes VLAN segmentation as a core logical design component, aligning with standard VCF networking practices.
Conclusion:Option D (DD02) is included in the logical design, as it defines the architectural approach to network segmentation, a key logical networking decision in VCF 5.2.References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Logical Design and Network Segmentation.
VMware Cloud Foundation 5.2 Networking Guide(docs.vmware.com): VLAN Usage in VCF. VMware Cloud Foundation 5.2 Design Guide(docs.vmware.com): Logical vs. Physical Design.

NEW QUESTION 4
As part of a VMware Cloud Foundation (VCF) design, an architect is responsible for planning for the migration of existing workloads using HCX to a new VCF environment. Which two prerequisites would the architect require to complete the objective? (Choose two.)

A. Extended IP spaces for all moving workloads.
B. DRS enabled within the VCF instance.
C. Service accounts for the applicable appliances.
D. NSX Federation implemented between the VCF instances.
E. Active Directory configured as an authentication source.

Answer: CE

Explanation:
VMware HCX (Hybrid Cloud Extension) is a key workload migration tool in VMware Cloud Foundation (VCF) 5.2, enabling seamless movement of VMs between on- premises environments and VCF instances (or between VCF instances). To plan an HCX- based migration, the architect must ensure prerequisites are met for deployment, connectivity, and operation. Let??s evaluate each option:
Option A: Extended IP spaces for all moving workloadsThis is incorrect. HCX supports migrations with or without extending IP spaces. Features like HCX vMotion and Bulk Migration allow VMs to retain their IP addresses (Layer 2 extension via Network Extension), while HCX Mobility Optimized Networking (MON) can adapt IPs if needed. Extended IP space is a design choice, not a prerequisite, making this option unnecessary for completing the objective.
Option B: DRS enabled within the VCF instanceThis is incorrect. VMware Distributed Resource Scheduler (DRS) optimizes VM placement and load balancing within a cluster but is not required for HCX migrations. HCX operates independently of DRS, handling VM mobility across environments (e.g., from a source vSphere to a VCF destination). While DRS might enhance resource management post-migration, it??s not a prerequisite for HCX functionality.
Option C: Service accounts for the applicable appliancesThis is correct. HCX requires service accounts with appropriate permissions to interact with source anddestination environments (e.g., vCenter Server, NSX). In VCF 5.2, HCX appliances (e.g., HCX Manager, Interconnect, WAN Optimizer) need credentials to authenticate and perform operations like VM discovery, migration, and network extension. The architect must ensure these accounts are configured with sufficient privileges (e.g., read/write access in vCenter), making this a critical prerequisite.
Option D: NSX Federation implemented between the VCF instancesThis is incorrect. NSX Federation is a multi-site networking construct for unified policy management across NSX deployments, but it??s not required for HCX migrations. HCX leverages its own Network Extension service to stretch Layer 2 networks between sites, independent of NSX Federation. While NSX is part of VCF, Federation is an advanced feature unrelated to HCX??s core migration capabilities.
Option E: Active Directory configured as an authentication sourceThis is correct. In VCF 5.2, HCX integrates with the VCF identity management framework, which typically uses Active Directory (AD) via vSphere SSO for authentication. Configuring AD as an authentication source ensures that HCX administrators can log in using centralized
credentials, aligning with VCF??s security model. This is a prerequisite for managing HCX appliances and executing migrations securely.
Conclusion:The two prerequisites required for HCX migration in VCF 5.2 areservice accounts for the applicable appliances(Option C) to enable HCX operations andActive Directory configured as an authentication source(Option E) for secure access management. These align with HCX deployment and integration requirements in the VCF ecosystem.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: HCX Integration)
VMware HCX User Guide (VCF 5.2 compatible): Prerequisites and Configuration VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Identity and Access Management)

NEW QUESTION 5
An architect is planning resources for a new cluster that will be integrated into an existing VI Workload Domain. The cluster??s primary purpose is to support a mission-critical application with five resource-intensive virtual machines. Which design recommendation should the architect provide to prevent resource bottlenecks while meeting the N+1 availability requirement and keeping the overall investment cost minimal?

A. Establish a cluster with four hosts and implement rules to prioritize resources for the application virtual machines.
B. Establish a cluster with three hosts and exclusively run the application virtual machines on this cluster.
C. Establish a cluster with six hosts and implement automated placement rules to keep the application virtual machines together.
D. Establish a cluster with six hosts and implement automated placement rules to distribute the application virtual machines.

Answer: A

Explanation:
Reference:VMware Cloud Foundation 5.2 Design Guide, Cluster Sizing; VMware vSphere 7.0 DRS Documentation.

NEW QUESTION 6
During a design discussion, the VMware Cloud Foundation Architect was presented with a requirement to reduce power utilization across all workload domains including management. The architect has suggested to use vSphere Distributed Power Management (DPM) to satisfy this requirement. Which recommendation should the architect provide?

A. vSphere DPM for Management Workload Domain (excluding when vSAN is a principal storage).
B. vSphere DPM for VI Workload Domains (excluding when vSAN is a principal storage).
C. vSphere DPM for Management Workload Domain (only when hosted within a Hyperscaler Data Center).
D. vSphere DPM for VI Workload Domains (any principal storage).
E. vSphere DPM for Management Workload Domain (any principal storage).

Answer: B

Explanation:
Reference:VMware Cloud Foundation 5.2 Administration Guide, Power Management; VMware vSphere 7.0 Resource Management Guide, DPM Considerations.

NEW QUESTION 7
An architect is working on higher-scale NSX Grouping and security design requirements for Management and VI Workload Domains in VMware Cloud Foundation. Which NSX Manager appliance size will be considered for use?

A. Extra Large
B. Large
C. Medium
D. Small

Answer: B

Explanation:
In VMware Cloud Foundation (VCF) 5.2, NSX Manager appliances manage networking and security (e.g., grouping, policies, firewalls) for Management and VI Workload Domains. The appliance size—Small,Medium, Large, Extra Large—determines its capacity to handle scale, such as the number of hosts, VMs, and security objects. The phrase ??higher scale?? implies a larger-than-minimum deployment. Let??s evaluate:
NSX Manager Appliance Sizes (VCF 5.2 with NSX-T 3.2):
Small: 4 vCPUs, 16 GB RAM, 300 GB disk. Supports up to 16 hosts, basic deployments (e.g., lab environments).
Medium: 6 vCPUs, 24 GB RAM, 300 GB disk. Supports up to 64 hosts, suitable for small to medium production environments.
Large: 12 vCPUs, 48 GB RAM, 300 GB disk. Supports up to 512 hosts, 10,000 VMs, and complex security policies—standard for production VCF.
Extra Large: 24 vCPUs, 64 GB RAM, 300 GB disk. Supports over 512 hosts, massive scale (e.g., service providers, multi-VCF instances).
VCF Context:
Management Domain: Minimum 4 hosts, often 6-7 for HA, with NSX for overlay networking.
VI Workload Domains: Variable host counts, but ??higher scale?? suggests multiple domains or significant workload growth.
Security Design: Grouping and policies (e.g., distributed firewall rules, tags) increase NSX Manager load, especially at scale.
Evaluation:
Small: Insufficient for production VCF, limited to 16 hosts. Unsuitable for a Management Domain (4-7 hosts) plus VI Workload Domains.
Medium: Adequate for small VCF deployments (up to 64 hosts), but ??higher scale?? implies more hosts or complex security, exceeding its capacity.
Large: The default and recommended size for VCF 5.2 production environments. It supports up to 512 hosts, thousands of VMs, and extensive security policies, fitting a Management Domain and multiple VI Workload Domains with ??higher scale?? needs.
Extra Large: Overkill unless managing hundreds of hosts or multiple VCF instances, which isn??t indicated here.
Conclusion:TheLargeNSX Manager appliance size (Option B) is appropriate for a higher- scale NSX design in VCF 5.2. It balances capacity and performance for Management and VI Workload Domains with advanced security requirements, aligning with VMware??s standard recommendation.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: NSX Manager Sizing)
NSX-T 3.2 Installation Guide (integrated in VCF 5.2): Appliance Size Specifications VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Security Design)

NEW QUESTION 8
A customer has stated the following requirements for Aria Automation within their VCF implementation:
• Users must have access to specific resources based on their company organization
• Developers must only be able to provision to the Development environment
• Production workloads can be placed on DMZ or Production clusters
What two design decisions must be implemented to satisfy these requirements? (Choose two.)

A. Separate cloud zones will be configured for Development and Production.
B. Users?? access to resources will be controlled by project membership.
C. Users?? access to resources will be controlled by tenant membership.
D. Separate tenants will be configured for Development and Production.

Answer: AB

Explanation:
Reference:VMware Aria Automation 8.10 Configuration Guide, Cloud Zones and Projects; VMware Cloud Foundation 5.2 Automation Guide.

NEW QUESTION 9
An Architect is designing a VMware Cloud Foundation (VCF)-based private cloud solution for a customer. During the requirements gathering workshop, the customer stated the following:
• All users must only have access to the solution components to fulfill their defined role.
• All administrative users must be authenticated to a separate approved identity source for administrator accounts only.
• All service users must be authenticated to the central approved identity source.
• All service account passwords must be stored centrally in an approved secrets management platform.
When creating the design, how should the Architect classify all the stated requirements?

A. Security
B. Manageability
C. Recoverability
D. Availability

Answer: A

Explanation:
Reference:VMware Cloud Foundation 5.2 Architect Study Guide, Chapter 3: Design Qualities, Section on Security Requirements; VMware Validated Design 6.2 (applicable to 5.2), Security Architecture.

NEW QUESTION 10
During the requirements gathering workshop for a new VMware Cloud Foundation (VCF)- based Private Cloud solution, the customer states that the solution must:
• Provide a single interface for monitoring all components of the solution.
• Minimize the effort required to maintain the solution to N-1 software versions. When creating the design document, under which design quality should the architect
classify these stated requirements?

A. Manageability
B. Recoverability
C. Availability
D. Performance

Answer: A

Explanation:

Reference:VMware Cloud Foundation 5.2 Architect Study Guide, Chapter 3: Design Qualities, Manageability Section.

NEW QUESTION 10
An architect is sizing the workloads that will run in a new VMware Cloud Foundation (VCF) Management Domain. The customer has a requirement to use Aria Operations to provide effective monitoring of the new VCF solution. What is the minimum Aria Operations Analytics node size requirement when AriaSuite Lifecycle is in VCF-aware mode?

A. Small
B. Extra Large
C. Medium
D. Large

Answer: C

Explanation:
VMware Aria Operations (formerly vRealize Operations) integrates with VMware Cloud Foundation 5.2 to monitor the Management Domain, including SDDC Manager, vCenter, NSX, and ESXi hosts. When deployed via VMware Aria Suite Lifecycle in VCF-aware mode, Aria Operations nodes must be sized to handle the monitoring workload effectively. The node size (Small, Medium, Large, Extra Large) determines resource capacity (CPU, memory, disk) and the number of objects (e.g., VMs, hosts) it can monitor. Let??s determine the minimum requirement:
Aria Operations Node Sizing in VCF 5.2:
Small: 4 vCPUs, 16 GB RAM, monitors up to 1,500 objects or 150 hosts. Suitable for small environments.
Medium: 8 vCPUs, 32 GB RAM, monitors up to 6,000 objects or 600 hosts. Suitable for medium to large environments.
Large: 16 vCPUs, 64 GB RAM, monitors up to 15,000 objects or 1,500 hosts. For large- scale deployments.
Extra Large: 24 vCPUs, 128 GB RAM, monitors over 15,000 objects or 1,500 hosts. For very large or dense environments.
VCF Management Domain Context:
The Management Domain in VCF 5.2 typically includes:
4-7 ESXi hosts (minimum 4 for HA, often 6-7 for resilience).
Management VMs (e.g., SDDC Manager, vCenter, NSX Managers, Aria Suite components).
Typically, fewer than 50-100 objects (VMs, hosts, networks) in a standard deployment. Aria Suite Lifecycle in VCF-aware mode deploys Aria Operations to monitor this domain, integrating with SDDC Manager for automated discovery and configuration.
Evaluation:
Small: Can monitor up to 150 hosts or 1,500 objects. For a Management Domain with ~7
hosts and <100 objects, this is sufficient capacity-wise but not the recommended minimum in VCF-aware mode due to integration overhead and future growth.
Medium: Supports up to 600 hosts or 6,000 objects. This size is recommended as the minimum for VCF deployments because it accommodates the Management Domain??s complexity (e.g., NSX, vSAN metrics) and allows headroom for additional monitoring (e.g., future Workload Domains).
Large/Extra Large: Overkill for a single Management Domain, designed for multi-domain or large-scale environments.
VMware Guidance:
The VMware Aria Operations documentation and VCF integration guides specify that in VCF-aware mode (via Aria Suite Lifecycle), theMediumnode size is the minimum recommended for effective monitoring of a Management Domain. This ensures performance for real-time analytics, dashboards, and integration with SDDC Manager, even if the initial object count is low. The Small size, while technically feasible for tiny setups, is not advised due to potential limitations in handling VCF-specific metrics and scalability.
Conclusion:The minimum Aria Operations Analytics node size requirement when Aria Suite Lifecycle is in VCF-aware mode isMedium(Option C). This balances resource needs with effective monitoring for the VCF 5.2 Management Domain.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Aria Operations Integration)
VMware Aria Operations 8.10 Sizing Guidelines (integrated in VCF 5.2): Node Size Recommendations
VMware Aria Suite Lifecycle 8.10 Documentation (VCF-aware mode requirements)

NEW QUESTION 12
A customer is designing a new VMware Cloud Foundation stretched cluster using L2 non- uniform connectivity, where due to a past incident an attacker was able to inject some false routes into their dynamic global routing table. What design decision can be taken to prevent this when configuring the Tier-0 gateway?

A. OSPF MD5 authentication
B. Gateway Firewall with ECMP
C. Implicit deny for any traffic
D. BGP peer password

Answer: D

Explanation:
The scenario involves designing a VMware Cloud Foundation (VCF) stretched cluster with L2 non-uniform connectivity, leveraging NSX (a core component of VCF) for networking. The customer??s past incident, where an attacker injected false routes into their dynamic global routing table, indicates a security vulnerability in the routing protocol. The Tier-0 gateway in NSX handles external connectivity and routing, typically using dynamic routing protocols like BGP (Border Gateway Protocol) or OSPF (Open Shortest Path First) to exchange routes with external routers. The design decision must prevent unauthorized route injection, ensuring the integrity of the routing table.
Context Analysis:
Stretched Cluster with L2 Non-Uniform Connectivity:In VCF 5.2, a stretched cluster spans multiple availability zones (AZs) with L2 connectivity for workload VMs, but the Tier-0 gateway uplinks may use L3 routing to external networks. ??Non-uniform?? suggests varying latency or bandwidth between sites, but this does not directly impact the routing security concern.
False Routes Injection:This implies the attacker exploited a lack of authentication or filtering in the routing protocol, allowing unauthorized route advertisements to be accepted into the Tier-0 gateway??s routing table.
Tier-0 Gateway:In NSX, the Tier-0 gateway is the edge component that peers with external routers (e.g., top-of-rack switches or upstream routers) and supports dynamic routing protocols like BGP and OSPF.
Routing Security in NSX:
NSX Tier-0 gateways commonly use BGP for external connectivity due to its scalability and flexibility in multi-site deployments like stretched clusters. OSPF is also supported but is less common for external peering in VCF designs.
Route injection attacks occur when an unauthorized device advertises routes without validation, often due to missing authentication mechanisms.
Option Analysis:
* A. OSPF MD5 authentication:OSPF supports MD5 authentication to secure routing updates between neighbors. Each OSPF message is hashed with a shared secret key, ensuring only trusted peers can exchange routes. This would prevent false route injection if OSPF were the protocol in use. However, in VCF stretched cluster designs, BGP is the

default and recommended protocol for Tier-0 gateway uplinks to external networks, as per the VMware Cloud Foundation Design Guide. OSPF is typically used for internal NSX routing (e.g., between Tier-0 and Tier-1 gateways) rather than external peering. Without evidence that OSPF is used here, and given BGP??s prevalence in such scenarios, this option is less applicable.
* B. Gateway Firewall with ECMP:The Gateway Firewall on the Tier-0 gateway filters traffic, not routes. Equal-Cost Multi-Path (ECMP) enhances bandwidth by load-balancing across multiple uplinks but does not inherently secure the routing table. While a firewall could block traffic from malicious sources, it cannot prevent the Tier-0 gateway from accepting false route advertisements in the control plane (routing protocol). Route injection occurs at the routing protocol level, not the data plane, so this option does not address theroot issue. The NSX Administration Guide confirms that firewall rules apply to packet forwarding, not route validation, making this incorrect.
* C. Implicit deny for any traffic:An implicit deny rule in the Gateway Firewall blocks all traffic not explicitly allowed, enhancing security for data plane traffic. However, this does not protect the control plane—specifically, the dynamic routing protocol—from accepting false routes. Route injection happens before traffic filtering, as the routing table determines where packets are sent. The VMware Cloud Foundation 5.2 documentation emphasizes that routing security requires protocol-specific measures, not just firewall rules. This option fails to prevent the described attack and is incorrect.
* D. BGP peer password:BGP supports authentication via a peer password (MD5-based in NSX), where each BGP session between the Tier-0 gateway and its external peers (e.g., physical routers) uses a shared secret. This ensures that only authenticated peers can advertise routes, preventing unauthorized devices from injecting false routes into the dynamic routing table. In VCF 5.2 stretched cluster deployments, BGP is the standard protocol for Tier-0 uplinks, as it supports multi-site connectivity and ECMP for redundancy. The NSX-T Data Center Design Guide and VCF documentation recommend BGP authentication to secure routing in such environments, directly addressing the customer??s past incident. This is the most relevant and effective design decision.
Conclusion:The architect should chooseBGP peer password (D)as the design decision for the Tier-0 gateway. This secures the BGP routing protocol—widely used in VCF stretched clusters—against false route injection by requiring authentication, aligning with the scenario??s security requirements and NSX best practices.
References:
VMware Cloud Foundation 5.2 Design Guide (Section: NSX Design for Stretched Clusters) VMware NSX-T Data Center 3.2 Administration Guide (Section: Tier-0 Gateway Routing) VMware Cloud Foundation 5.2 Planning and Preparation Workbook (Section: Networking Security)
VMware Validated Design for Stretched Clusters (Section: Routing Security)

NEW QUESTION 15
An architect is designing a VMware Cloud Foundation (VCF)-based private cloud solution for a customer. The customer has stated the following requirement:
• All management tooling must be resilient against a single ESXi host failure
When considering the design decisions for VMware Aria Suite components, what should the Architect document to support the stated requirement?

A. The solution will deploy the VCF Workload domain in a stretched topology across two sites.
B. The solution will deploy three Aria Automation appliances in a clustered topology.
C. The solution will deploy Aria Suite Lifecycle in a clustered topology.
D. The solution will deploy an external load balancer for Aria Operations Cloud Proxies.

Answer: B

Explanation:
Reference:VMware Aria Automation 8.10 Installation Guide, Section on High Availability Configuration; VMware Cloud Foundation 5.2 Architecture and Deployment Guide, Management Domain HA.

NEW QUESTION 20
The following requirements were identified in an architecture workshop for a VMware Cloud Foundation (VCF) design project utilizing vSAN for its primary storage solution:
REQ001: Application must maintain a minimum of 1,000 transactions per second (TPS) during business hours excluding disaster recovery (DR) scenarios.
REQ002: Automatic DRS and HA must be utilized.
REQ003: Planned maintenance must be executed outside of business hours.
Which of the following test scenarios should be added and performed to validate these requirements?

A. Trigger a Virtual Machine vMotion operation.
B. Trigger a vCenter Server update.
C. Trigger a vSAN disk group evacuation.
D. Trigger a failure of an ESXi host.

Answer: D

Explanation:
To validate the stated requirements, the test scenario must address all three: application performance (1,000 TPS), automatic DRS and HA functionality, and maintenance timing (implying minimal disruption during business hours). In a VCF environment with vSAN, test scenarios should simulate real-world conditions that challenge these requirements. Let??s evaluate each option:
Option A: Trigger a Virtual Machine vMotion operationvMotion tests DRS??s ability to migrate VMs for load balancing, which aligns with REQ002??s ??automatic DRS?? mandate. It can be scheduled outside business hours (REQ003) to minimize impact. However, it doesn??t fully test HA (automatic failover) or ensure 1,000 TPS (REQ001) under failure conditions, as vMotion is a planned operation, not a failure scenario. This is a partial match but not comprehensive.
Option B: Trigger a vCenter Server updateUpdating vCenter tests management plane resilience but doesn??t directly validate application performance (REQ001), DRS/HA automation (REQ002), or vSAN-specific behavior. While it could relate to maintenance
(REQ003), it??s unrelated to workload or storage functionality in the VCF design, making it irrelevant here.
Option C: Trigger a vSAN disk group evacuationEvacuating a vSAN disk group simulates maintenance (REQ003) by moving data to other nodes, testing vSAN??s resilience. It may involve DRS for VM migration (REQ002), but it doesn??t trigger HA failover. While it could indirectly affect TPS (REQ001), the requirement excludes DR scenarios, and this test doesn??t guarantee performance validation during business hours under normal operations or host failure.
Option D: Trigger a failure of an ESXi hostSimulating an ESXi host failure directly tests REQ002: HA automatically restarts VMs on other hosts, and DRS balances the load post- failure. In a vSAN environment, it also validates data availability (vSAN rebuilds objects), ensuring 1,000 TPS (REQ001) is maintained during business hours under failure conditions (excluding DR, as this is a single-host failure within a site). While not a maintenance task (REQ003), it implicitly ensures maintenance-like disruptions (e.g., host failure) don??t violate performance, aligning with VCF??s HA/DRS automation goals. TheVCF 5.2 Administration Guiderecommends host failure testing to validate HA and vSAN resilience.
Conclusion:Option D comprehensively validates REQ001 (TPS under failure), REQ002 (automatic DRS and HA), and indirectly supports REQ003 by ensuring business-hour performance during unplanned events, making it the best test scenario.References: VMware Cloud Foundation 5.2 Administration Guide(docs.vmware.com): vSAN and HA/DRS Testing Scenarios.
vSphere Availability Guide(docs.vmware.com): HA Failover Testing.
vSAN Administration Guide(docs.vmware.com): Disk Group Evacuation and Failure Scenarios.

NEW QUESTION 25

An architect is working with an organization on the creation of a new Private Cloud Platform. The organization has provided the following business objectives they wish to achieve with the new platform:
• Reduce the operating costs associated with running separate areas of hosting capacity and separate/duplicate systems.
• Reduce the risks, time, and effort associated with managing platforms that are out of vendor support.
• Reduce the operating costs associated with Public Cloud usage.
• Reduce the risks associated with having incomplete documentation for application inventory and dependency mappings.
They have grouped these business objectives into a set of use cases:
• Migration - Provide a platform that supports the migration of virtualized workloads from existing platforms.
• Containerization - Provide a platform that supports the deployment of containerized workloads.
• Centralization and Consolidation - Provide a central private cloud platform accessible to all relevant areas of the business.
When considering these objectives and use cases, what should the architect include in the design documentation as a part of the Conceptual Model?

A. An assumption that the new platform will co-exist with the existing platforms for a period of time to allow workloads to be migrated in a phased approach
B. A risk that the existing platforms are running Linux Operating Systems that are out of vendor support
C. An assumption that a complete mapping of application dependencies is not available
D. A requirement that the solution will provide the capability to migrate Kubernetes-based workloads from the Public Cloud

Answer: A

Explanation:
Reference:VMware Cloud Foundation 5.2 Architect Study Guide, Chapter 1: Conceptual Design; VMware Migration Planning Guide for VCF.

NEW QUESTION 29
The following are a list of design decisions made relating to networking: NSX Distributed Firewall (DFW) rule to block all traffic by default. Implement overlay network technology to scale across data centers.
Configure Cisco Discovery Protocol (CDP) - Listen mode on all Distributed Virtual Switches (DVS).
Use of 2x 64-port Cisco Nexus 9300 for top-of-rack ESXi host switches. Which design decision would an architect document within the logical design?

A. Use of 2x 64-port Cisco Nexus 9300 for top-of-rack ESXi host switches.
B. NSX Distributed Firewall (DFW) rule to block all traffic by default.
C. Implement overlay network technology to scale across data centers.
D. Configure Cisco Discovery Protocol (CDP) - Listen mode on all Distributed Virtual Switches (DVS).

Answer: C

Explanation:
In VCF 5.2, the logical design focuses on high-level architectural decisions that define the system??s structure and behavior, as opposed to physical or operational details. Networking decisions in the logical design emphasize scalability, security policies, and connectivity frameworks, per theVCF 5.2 Architectural Guide. Let??s evaluate each: Option A: Use of 2x 64-port Cisco Nexus 9300 for top-of-rack ESXi host switches This specifies physical hardware, a detail typically documented in the physical design (e.g., BOM, rack layout). TheVCF 5.2 Design Guidedistinguishes hardware choices as physical, not logical, unless they dictate architecture (e.g., spine-leaf), which isn??t implied here. Option B: NSX Distributed Firewall (DFW) rule to block all traffic by defaultThis is a security policy configuration within NSX, defining how traffic is controlled. While critical, it??s an operational or detailed design decision (e.g., rule set), not a high-level logical design element. TheVCF 5.2 Networking Guideplaces DFW rules in implementation details, not the logical overview.
Option C: Implement overlay network technology to scale across data centers Overlay networking (e.g., NSX VXLAN or Geneve) is a foundational architectural decision in VCF, enabling scalability, multi-site connectivity, and logical separation of networks. The VCF 5.2 Architectural Guidehighlights overlays as a core logical design component, directly impacting how the solution scales across data centers, making it a prime candidate for the logical design.
Option D: Configure Cisco Discovery Protocol (CDP) - Listen mode on all Distributed Virtual Switches (DVS)CDP in Listen mode aids network discovery and troubleshooting on DVS. This is a configuration setting, not a logical design decision. TheVCF 5.2 Networking Guidetreats such protocol settings as operational details, not architectural choices.
Conclusion:Option C belongs in the logical design, as it defines a scalable networking architecture critical to VCF 5.2??s multi-data center capabilities.References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Logical Design and Overlay Networking.
VMware Cloud Foundation 5.2 Networking Guide(docs.vmware.com): NSX and DVS Configuration.
VMware Cloud Foundation 5.2 Design Guide(docs.vmware.com): Logical vs. Physical Design.

NEW QUESTION 33
An architect is working on a leaf-spine design requirement for NSX Federation in VMware Cloud Foundation. Which recommendation should the architect document?

A. Use a physical network that is configured for EIGRP routing adjacency.
B. Layer 3 device that supports OSPF.
C. Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 1500 ms.
D. Jumbo frames on the components of the physical network between the VMware Cloud Foundation instances.

Answer: D

Explanation:
NSX Federation in VMware Cloud Foundation (VCF) 5.2 extends networking and security across multiple VCF instances (e.g., across data centers) using a leaf-spine underlay network. The architect must recommend a physical network design that supports this. Let??s evaluate:
Option A: Use a physical network that is configured for EIGRP routing adjacency
Enhanced Interior Gateway Routing Protocol (EIGRP) is a Cisco-proprietary routing protocol. NSX Federation requires a Layer 3 underlay with dynamic routing (e.g., BGP, OSPF), but EIGRP isn??t a VMware-recommended standard for NSX leaf-spine designs. BGP is preferred for its scalability and interoperability in NSX-T 3.2 (used in VCF 5.2). This option is not optimal.
Option B: Layer 3 device that supports OSPF
Open Shortest Path First (OSPF) is a supported routing protocol for NSX underlays, alongside BGP. A Layer 3 device with OSPF could work in a leaf-spine topology, but VMware documentation emphasizes BGP as the primary choice for NSX Federation due to its robustness in multi-site scenarios. OSPF is valid but not the strongest recommendation for Federation-specific designs.
Option C: Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 1500 ms
NSX Federation requires low latency between sites for control plane consistency (Global Manager to Local Managers). The maximum supported latency is 150 ms (not 1500 ms), per VMware specs. 1500 ms (1.5 seconds) is far too high and would disrupt Federation operations, making this incorrect.
Option D: Jumbo frames on the components of the physical network between the VMware Cloud Foundation instances
This is correct. NSX Federation relies on NSX-T overlay traffic (Geneve encapsulation) across sites, which benefits from jumbo frames (MTU 9000) to reduce

fragmentation and improve performance. In a leaf-spine design, enabling jumbo frames on all physical network components (switches, routers) between VCF instances ensures efficient transport of tunneled traffic (e.g., for stretched networks). VMware strongly recommends this for NSX underlays, making it the best recommendation.
Conclusion:The architect should documentD: Jumbo frames on the components of the physical network between the VMware Cloud Foundation instances. This aligns with VCF 5.2 and NSX Federation??s leaf-spine design requirements for optimal performance and scalability.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: NSX Federation Networking)
NSX-T 3.2 Reference Design (integrated in VCF 5.2): Leaf-Spine Underlay Requirements VMware NSX-T 3.2 Installation Guide: Jumbo Frame Recommendations

NEW QUESTION 38
An architect is designing a new VMware Cloud Foundation (VCF) solution. During the discovery workshops, the customer explained that the solution will initially be used to host a single business application and some internal management tooling. The customer provided the following background information:
The business application consists of two virtual machines.
The business application is sensitive to changes in its storage I/O.
The business application must be available during the company??s business hours of 9 AM - 5 PM on weekdays.
The architect has made the following design decisions in response to the customer??s requirements and the additional information provided during discovery:
The solution will use the VCF consolidated architecture model. A single cluster will be created, consisting of six ESXi hosts.
Which design decision should the architect include in the design to mitigate the risk of impacting the business application?

A. Use resource pools to apply CPU and memory reservations on the business application virtual machines.
B. Implement FTT=6 for the business application virtual machines.
C. Perform ESXi host maintenance activities outside of the stated business hours.
D. Replace the vSAN shared storage exclusively with an All-Flash Fibre Channel shared storage solution.
E. Use Anti-Affinity Distributed Resource Scheduler (DRS) rules on the business application virtual machines.

Answer: C

Explanation:
The VCF 5.2 design must ensure the business application (two VMs) remains available during business hours (9 AM - 5 PM weekdays) and is protected from storage I/O disruptions in a consolidated architecture with a single six-host cluster using vSAN. The goal is to mitigate risks to the application??s performance and availability. Let??s evaluate each option:
Option A: Use resource pools to apply CPU and memory reservations on the business application virtual machinesResource pools with reservations ensure CPU and memory availability, which could help performance. However, the application??s sensitivity is tostorage I/O, not CPU/memory, and the availability requirement (business hours) isn??t directly addressed by reservations. While useful, this doesn??t fully mitigate the primary risks identified, making it less optimal.
Option B: Implement FTT=6 for the business application virtual machinesThis is incorrect and infeasible. In vSAN, Failures to Tolerate (FTT) defines the number of host or disk failures a storage object can withstand, with a maximum FTT dependent on cluster size. FTT=6 requires at least 13 hosts (2n+1 where n=6), but the cluster has only six hosts, supporting a maximum FTT=2 (RAID-5/6). Even if feasible, FTT addresses data redundancy, not runtime availability or I/O sensitivity during business hours, making this irrelevant to the stated risks.
Option C: Perform ESXi host maintenance activities outside of the stated business hoursThis is the correct answer. In a vSAN-based VCF cluster, ESXi host maintenance (e.g., patching, reboots) triggers data resyncs and VM migrations (via vMotion), which can impact storage I/O performance and potentially cause brief disruptions. The application??s sensitivity to storage I/O and its availability requirement (9 AM - 5 PM weekdays) mean maintenance during business hours poses a risk. Scheduling maintenance outside these hours (e.g., nights or weekends) mitigates this by ensuring uninterrupted I/O performance and availability during critical times, directly addressing the customer??s needs.
Option D: Replace the vSAN shared storage exclusively with an All-Flash Fibre Channel shared storage solutionThis is incorrect. While an All-Flash Fibre Channel array might offer better I/O performance, VCF??s consolidated architecture relies on vSAN as the primary storage for management and workload domains. Replacing vSAN entirely contradicts the chosen architecture and introduces unnecessarycomplexity and cost. The sensitivity to storage I/O changes doesn??t justify abandoning vSAN, especially since All- Flash vSAN could meet performance needs if properly tuned.
Option E: Use Anti-Affinity Distributed Resource Scheduler (DRS) rules on the business application virtual machinesAnti-Affinity DRS rules ensure the two VMs run on separate hosts, improving availability by avoiding a single host failure impacting both. While this mitigates some risk, it doesn??t address storage I/O sensitivity (a vSAN-wide concern) or guarantee availability during business hours if maintenance occurs. It??s a partial solution but less effective than scheduling maintenance outside business hours.
Conclusion:The best design decision is toperform ESXi host maintenance activities outside of the stated business hours(Option C). This directly mitigates the risk of storage I/O disruptions and ensures availability during 9 AM - 5 PM weekdays, aligning with the customer??s requirements in the VCF 5.2 consolidated architecture.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Consolidated Architecture Design)
VMware vSAN 7.0U3 Planning and Deployment Guide (integrated in VCF 5.2): Maintenance Mode Considerations
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Availability and Performance Design)

NEW QUESTION 40
When sizing a VMware Cloud Foundation VI Workload Domain, which three factors should be considered when calculating usable compute capacity? (Choose three.)

A. NSX
B. vSphere HA
C. vSAN
D. NIOC
E. Storage DRS
F. Core Dumps

Answer: BCD

Explanation:
When sizing a VMware Cloud Foundation (VCF) VI Workload Domain, calculating usable compute capacity involves determining the resources available for workloads after accounting for overheads and system-level requirements. In VCF 5.2, a VI Workload Domain integrates vSphere, vSAN, and NSX, and certain factors directly impact the compute capacity available to virtual machines. Based on the official VMware Cloud Foundation 5.2 documentation, the three key factors to consider are vSphere HA, vSAN, and NIOC.

NEW QUESTION 44
As a VMware Cloud Foundation architect, you are provided with the following requirements:

All administrative access to the cloud management components must be trusted. All cloud management components?? communications must be encrypted.
Enhancement of lifecycle management should always be considered.
Which design decision fulfills the requirements?

A. Integrate the SDDC Manager with a supported 3rd-party certificate authority (CA).
B. Integrate the SDDC Manager with the vCenter Server in VMCA mode.
C. Write a PowerCLI script to run on all virtual appliances and force a redirection on port 443.
D. Write an Aria Orchestrator Workflow to change the ESXi hosts?? certificates in bulk.

Answer: A

Explanation:
The requirements focus on trust, encryption, and lifecycle management for a VMware Cloud Foundation (VCF) 5.2 solution. VCF leverages SDDC Manager, vCenter Server, NSX, and ESXi hosts as core management components, and their security and manageability are critical. Let??s evaluate each option against the requirements:
Option A: Integrate the SDDC Manager with a supported 3rd-party certificate authority (CA)This is the correct answer. In VCF 5.2, integrating SDDC Manager with a 3rd-party CA (e.g., Microsoft CA, OpenSSL) allows it to manage and deploy trusted certificates across all management components (e.g., vCenter, NSX Manager, ESXi hosts). This ensures:
Trusted administrative access: Certificates from a trusted CA secure administrative interfaces (e.g., HTTPS access to SDDC Manager and vCenter), ensuring authenticated and verified connections.
Encrypted communications: All management component interactions (e.g., API calls, UI access) use TLS with CA-signed certificates, encrypting data in transit.
Lifecycle management enhancement: SDDC Manager automates certificate lifecycle operations (e.g., issuance, renewal, replacement), reducing manual effort and improving operational efficiency.The VMware Cloud Foundation documentation explicitly supports this integration as a best practice for security and scalability, fulfilling all three requirements comprehensively.
Option B: Integrate the SDDC Manager with the vCenter Server in VMCA modeThis is
incorrect. The vCenter Server??s VMware Certificate Authority (VMCA) can issue certificates for vSphere components (e.g., ESXi hosts, vCenter itself), but it operates within the vSphere domain, not across the broader VCF stack. SDDC Manager requires a higher- level CA integration to managecertificates for all components (including NSX and itself). VMCA mode doesn??t extend trust to SDDC Manager or NSX Manager natively, nor does it enhance lifecycle management across the entire VCF solution—it??s limited to vSphere. This option fails to fully address the requirements.
Option C: Write a PowerCLI script to run on all virtual appliances and force a redirection on port 443This is incorrect. Forcing redirection to port 443 (HTTPS) via a PowerCLI script might enable encrypted communication for some components, but it??s a manual, ad-hoc solution that:
Doesn??t ensuretrustedaccess (no mention of certificate trust). Doesn??t integrate with a CA for certificate management.
Contradicts lifecycle enhancement, as it requires ongoing manual intervention rather than automation.This approach is not scalable or supported in VCF 5.2 for meeting security requirements.
Option D: Write an Aria Orchestrator Workflow to change the ESXi hosts?? certificates in bulkThis is incorrect. While VMware Aria Orchestrator (formerly vRealize Orchestrator) can automate certificate updates for ESXi hosts, it??s a partial solution that:
Only addresses ESXi hosts, not all management components (e.g., SDDC Manager, NSX). Doesn??t inherently ensure trust unless tied to a trusted CA (not specified here).
Improves lifecycle management only for ESXi certificates, not the broader VCF stack.This option lacks the holistic scope required by the question and isn??t a native VCF design decision.
Conclusion:Integrating SDDC Manager with a 3rd-party CA (Option A) is the only design decision that fully satisfies all requirements. It leverages VCF 5.2??s built-in certificate management capabilities to ensure trust, encryption, and lifecycle efficiency across the entire solution.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Certificate Management)
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Security Design Considerations)
vSphere 7.0U3 Security Configuration Guide (integrated in VCF 5.2): Certificate Authority Integration

NEW QUESTION 47
During a requirement gathering workshop, various Business and Technical requirements were collected from the customer. Which requirement would be categorized as a Business Requirement?

A. The application should be compatible with Windows, macOS, and Linux operating systems.
B. Decrease processing time for service requests by 30%.
C. The system should support 10,000 concurrent users.
D. Data should be encrypted using AES-256 encryption.

Answer: B

Explanation:
Business requirements in VCF articulate organizational objectives that the solution must enable, often focusing on efficiency, cost, or service improvements rather than specific technical implementations. Option B, "Decrease processing time for service requests by 30%," is a business requirement as it targets an operational efficiency goal that benefits the customer??s service delivery, measurable from a business perspective rather than dictating how the system achieves it. Options A, C, and D—specifying OS compatibility, user capacity, and encryption standards—are technical requirements, as they detail system capabilities or security mechanisms that architects must implement within VCF components like vSphere or NSX. The distinction hinges on intent: B focuses on outcome (speed), while others define system properties.
Reference: VMware Cloud Foundation 5.2 Architect Study Guide, Chapter 2: Requirements
Classification, Section on Business vs. Technical Requirements.

NEW QUESTION 51
An architect is responsible for updating the design of a VMware Cloud Foundation solution for a pharmaceuticals customer to include the creation of a new cluster that will be used for a new research project. The applications that will be deployed as part of the new project will include a number of applications that are latency-sensitive. The customer has recently completed a right-sizing exercise using VMware Aria Operations that has resulted in a number of ESXi hosts becoming available for use. There is no additional budget for purchasing hardware. Each ESXi host is configured with:
2 CPU sockets (each with 10 cores)
512 GB RAM divided evenly between sockets
The architect has made the following design decisions with regard to the logical workload design:
The maximum supported number of vCPUs per virtual machine size will be 10. The maximum supported amount of RAM (GB) per virtual machine will be 256.
What should the architect record as the justification for these decisions in the design document?

A. The maximum resource configuration will ensure efficient use of RAM by sharing memory pages between virtual machines.
B. The maximum resource configuration will ensure the virtual machines will cross NUMA node boundaries.

C. The maximum resource configuration will ensure the virtual machines will adhere to a single NUMA node boundary.
D. The maximum resource configuration will ensure each virtual machine will exclusively consume a whole CPU socket.

Answer: C

Explanation:
The architect??s design decisions for the VMware Cloud Foundation (VCF) solution must align with the hardware specifications, the latency-sensitive nature of the applications, and VMware best practices for performance optimization. To justify the decisions limiting VMs to 10 vCPUs and 256 GB RAM, we need to analyze the ESXi host configuration and the implications of NUMA (Non-Uniform Memory Access) architecture, which is critical for latency-sensitive workloads.
ESXi Host Configuration:
CPU:2 sockets, each with 10 cores (20 cores total, or 40 vCPUs with hyper-threading, assuming it??s enabled).
RAM:512 GB total, divided evenly between sockets (256 GB per socket).
Each socket represents a NUMA node, with its own local memory (256 GB) and 10 cores. NUMA nodes are critical because accessing local memory is faster than accessing remote memory across nodes, which introduces latency.
Design Decisions:
Maximum 10 vCPUs per VM:Matches the number of physical cores in one socket (NUMA node).
Maximum 256 GB RAM per VM:Matches the memory capacity of one socket (NUMA node).
Latency-sensitive applications:These workloads (e.g., research applications) require minimal latency, making NUMA optimization a priority.
NUMA Overview (VMware Context):In vSphere (a core component of VCF), each physical CPU socket and its associated memory form a NUMA node. When a VM??s vCPUs and memory fit within a single NUMA node, all memory access is local, reducing latency. If a VM exceeds a NUMA node??s resources (e.g., more vCPUs or memory than one socket provides), it spans multiple nodes, requiring remote memory access, which increases latency—a concern for latency-sensitive applications. VMware??s vSphere NUMA scheduler optimizes VM placement, but the architect can enforce performance by sizing VMs appropriately.
Option Analysis:
* A. The maximum resource configuration will ensure efficient use of RAM by sharing memory pages between virtual machines:This refers to Transparent Page Sharing (TPS), a vSphere feature that allows VMs to share identical memory pages, reducing RAM usage. While TPS improves efficiency, it is not directly tied to the decision to cap VMs at 10 vCPUs and 256 GB RAM. Moreover, TPS has minimal impact on latency-sensitive workloads, as it??s a memory-saving mechanism, not a performance optimization for latency. The VMware Cloud Foundation Design Guide and vSphere documentation note that TPS is disabled by default in newer versions (post-vSphere 6.7) due to security concerns, unless explicitly enabled. This justification does not align with the latency focus or the specific resource limits, making it incorrect.
* B. The maximum resource configuration will ensure the virtual machines will cross NUMA node boundaries:If VMs were designed to cross NUMA node boundaries (e.g., more than 10 vCPUs or 256 GB RAM), their vCPUs and memory would span both sockets. For example, a VM with 12 vCPUs would use cores from both sockets, and a VM with 300 GB RAM would require memory from both NUMA nodes. This introduces remote memory access, increasing latency due to inter-socket communication over the CPU interconnect (e.g., Intel QPI or AMD Infinity Fabric). For latency-sensitive applications, crossing NUMA boundaries is undesirable, as noted in the VMware vSphere Resource Management Guide. This option contradicts the goal and is incorrect.
* C. The maximum resource configuration will ensure the virtual machines will adhere to a single NUMA node boundary:By limiting VMs to 10 vCPUs and 256 GB RAM, the architect ensures each VM fits within one NUMA node (10 cores and 256 GB per socket). This means all vCPUs and memory for a VM are allocated from the same socket, ensuring local memory access and minimizing latency. This is a critical optimization for latency- sensitive workloads, as remote memory access is avoided. The vSphere NUMA scheduler will place each VM on a single node, and since the VM??s resource demands do not exceed the node??s capacity, no NUMA spanning occurs. The VMware Cloud Foundation 5.2 Design Guide and vSphere best practices recommend sizing VMs to fit within a NUMA node for performance-critical applications, making this the correct justification.
* D. The maximum resource configuration will ensure each virtual machine will exclusively consume a whole CPU socket:While 10 vCPUs and 256 GB RAM match the resources of one socket, this option implies exclusive consumption, meaning no other VM could use that socket. In vSphere, multiple VMs can share a NUMA node as long as resources are available (e.g., two VMs with 5 vCPUs and 128 GB RAM each could coexist on one socket). The architect??s decision does not mandate exclusivity but rather ensures VMs fit within a node??s boundaries. Exclusivity would limit scalability (e.g., only two VMs per host), which isn??t implied by the design or required by the scenario. This option overstates the intent and is incorrect.
Conclusion:The architect should record thatthe maximum resource configuration will ensure the virtual machines will adhere to a single NUMA node boundary (C). This justification aligns with the hardware specs, optimizes for latency-sensitive workloads by avoiding remote memory access, and leverages VMware??s NUMA-aware scheduling for performance.
References:
VMware Cloud Foundation 5.2 Design Guide (Section: Workload Domain Design) VMware vSphere 8.0 Update 3 Resource Management Guide (Section: NUMA Optimization)
VMware Cloud Foundation 5.2 Planning and Preparation Workbook (Section: Host Sizing) VMware Best Practices for Performance Tuning Latency-Sensitive Workloads (White Paper)

NEW QUESTION 53
An architect is tasked with updating the design for an existing VMware Cloud Foundation (VCF) deployment to include four vSAN ESA ready nodes. The existing deployment comprises the following:
Four homogenous vSAN ESXi ready nodes in the management domain.
Four homogenous ESXi nodes with iSCSI principal storage in workload domain A. What should the architect recommend when including this additional capacity for application workloads?

A. Commission the four new nodes into the existing workload domain A cluster.
B. Create a new vLCM image workload domain with the four new nodes.
C. Create a new vLCM baseline cluster in the existing workload domain with the four new nodes.
D. Create a new vLCM baseline workload domain with the four new nodes.

Answer: D

Explanation:
The task involves adding four vSAN ESA (Express Storage Architecture) ready nodes to an existing VCF 5.2 deployment for application workloads. The current setup includes a vSAN-based Management Domain and a workload domain (A) using iSCSI storage. In VCF, workload domains are logical units with consistent storage and lifecycle management via vSphere Lifecycle Manager (vLCM). Let??s analyze each option: Option A: Commission the four new nodes into the existing workload domain A clusterWorkload domain A uses iSCSI storage, while the new nodes are vSAN ESA ready. VCF 5.2 doesn??t support mixing principal storage types (e.g., iSCSI and vSAN) within a single cluster, as per theVCF 5.2 Architectural Guide. Commissioning vSAN nodes into an iSCSI cluster would require converting the entire cluster to vSAN, which isn??t feasible with existing workloads and violates storage consistency, making this impractical.
Option B: Create a new vLCM image workload domain with the four new nodesThis phrasing is ambiguous. vLCM manages ESXi images and baselines, but ??vLCM image workload domain?? isn??t a standard VCF term. It might imply a new workload domain with a custom vLCM image,but lacks clarity compared to standard options (C, D). TheVCF 5.2 Administration Guideuses ??baseline?? or ??image-based?? distinctly, so this is less precise. Option C: Create a new vLCM baseline cluster in the existing workload domain with the four new nodesAdding a new cluster to an existing workload domain is possible in VCF, but clusters within a domain must share the same principal storage (iSCSI in workload domain A). TheVCF 5.2 Administration Guidestates that vSAN ESA requires a dedicated cluster and can??t coexist with iSCSI in the same domain configuration, rendering this option invalid.
Option D: Create a new vLCM baseline workload domain with the four new nodesA new workload domain with vSAN ESA as the principal storage aligns with VCF

5.2 design principles. vLCM baselines ensure consistent ESXi versioning and firmware for the new nodes. TheVCF 5.2 Architectural Guiderecommends separate workload domains for different storage types or workload purposes (e.g., application capacity). This leverages the vSAN ESA nodes effectively, isolates them from the iSCSI-based domain A, and supports application workloads seamlessly.
Conclusion:Option D is the best recommendation, creating a new vSAN ESA-based workload domain managed by vLCM, meeting capacity needs while adhering to VCF 5.2 storage and domain consistency rules.References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Workload Domain Design and vSAN ESA.
VMware Cloud Foundation 5.2 Administration Guide(docs.vmware.com): vLCM and Cluster Expansion.
vSAN ESA Planning and Deployment Guide(docs.vmware.com): Storage Requirements.

NEW QUESTION 54
An architect has been asked to recommend a solution for a mission-critical application running on a single virtual machine to ensure consistent performance. The virtual machine operates within a vSphere cluster of four ESXi hosts, sharing resources with other production virtual machines. There is no additional capacity available. What should the architect recommend?

A. Use CPU and memory reservations for the mission-critical virtual machine.
B. Use CPU and memory limits for the mission-critical virtual machine.
C. Create a new vSphere Cluster and migrate the mission-critical virtual machine to it.
D. Add additional ESXi hosts to the current cluster.

Answer: A

Explanation:
In VMware vSphere, ensuring consistent performance for a mission-critical virtual machine (VM) in a resource-constrained environment requires guaranteeing that the VM receives the necessary CPU and memory resources, even when the cluster is under contention. The scenario specifies that the VM operates in a four-host vSphere cluster with no additional capacity available, meaning options that require adding resources (like D) or creating a new cluster (like C) are not feasible without additional hardware, which isn??t an option here.
Option A: Use CPU and memory reservationsReservations in vSphere guarantee a minimum amount of CPU and memory resources for a VM, ensuring that these resources are always available, even during contention. For a mission-critical application, this is the most effective way to ensure consistent performance because it prevents other VMs from consuming resources allocated to this VM. According to theVMware Cloud Foundation 5.2 Architectural Guide, reservations are recommended for workloads requiring predictable performance, especially in environments where resource contention is a risk (e.g., 90% utilization scenarios). This aligns with VMware??s best practices for mission-critical workloads.
Option B: Use CPU and memory limitsLimits cap the maximum CPU and memory a VM
can use, which could starve the mission-critical VM of resources when it needs to scale up to meet demand. This would degrade performance rather than ensure consistency, making it an unsuitable choice. ThevSphere Resource Management Guide(part of VMware??s documentation suite) advises against using limits for performance-critical VMs unless the goal is to restrict resource usage, not guarantee it.
Option C: Create a new vSphere Cluster and migrate the mission-critical virtual machine to itCreating a new cluster implies additional hardware or reallocation of existing hosts, but the question states there is no additional capacity. Without available resources, this option is impractical in the given scenario.
Option D: Add additional ESXi hosts to the current clusterWhile adding hosts would increase capacity and potentially reduce contention, the lack of additional capacity rules this out as a viable recommendation without violating the scenario constraints.
Thus,Ais the best recommendation as it leverages vSphere??s resource management capabilities to ensure consistent performance without requiring additional hardware. References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Section on Resource Management for Workload Domains.
vSphere Resource Management Guide(docs.vmware.com): Chapter on Configuring Reservations, Limits, and Shares.

NEW QUESTION 56
A company will be expanding their existing VCF environment for a new application. The existing VCF environment currently has a management domain and two separate VI workload domains with different hardware profiles. The new application has the following requirements:
• The application will use significantly more memory than current workloads today.
• The application will have a limited number of licenses to run on hosts.
• Additional VCF and hardware costs have been approved for the application.
• The application will contain confidential customer information that requires isolation from other workloads.
What design recommendation should the administrator document?

A. Deploy a new consolidated VCF instance and deploy the new application into it.
B. A new Workload domain with hardware supporting the memory requirements of the new application should be implemented.
C. Enough identical hardware for the management domain should be ordered to accommodate the new application requirements and a new workload domain should be designed for the application.
D. Purchase enough matching hardware to accommodate the new application??s memory requirements and expand an existing cluster to accommodate the new applicatio
E. Use host affinity rules to manage the new licensing.

Answer: B

Explanation:
Reference:VMware Cloud Foundation 5.2 Architecture and Deployment Guide, Workload Domain Design; VMware vSphere 7.0 Documentation, DRS Affinity Rules.

NEW QUESTION 59
Which Operating System (OS) is not supported by Aria Operations for OS and Application Monitoring?

A. Windows Server 2012 R2
B. CentOS
C. Windows Server 2012
D. MacOS

Answer: D

Explanation:
Reference:VMware Aria Operations 8.10 Product Documentation, Supported Operating Systems for Monitoring; VMware Cloud Foundation 5.2 Release Notes.

NEW QUESTION 60
An administrator is designing a new VMware Cloud Foundation instance that has to support management, VDI, DB, and general workloads. The DB workloads will stay the same in terms of resources over time. However, the general workloads and VDI environments are expected to grow over the next 3 years. What should the architect include in the documentation?

A. An assumption that the DB workload resource requirements will remain static.
B. A constraint of including the management, DB, and VDI environments.
C. A requirement consisting of the growth of the general workloads and VDI environment.
D. A risk that the VCF instance may not have enough capacity for growth.

Answer: A

Explanation:
In VMware Cloud Foundation (VCF) 5.2, design documentation includes assumptions, constraints, requirements, and risks to define the solution??s scope and address potential challenges. The scenario provides specific information about workload types and their behavior over time, which the architect must categorize appropriately. Let??s evaluate each option:
Option A: An assumption that the DB workload resource requirements will remain staticThis is the correct answer. Anassumptionis a statement taken as true without proof, often based on customer-provided information, to guide design planning. The customer explicitly states that ??the DBworkloads will stay the same in terms of resources over time.?? Documenting this as an assumption reflects this fact and allows the architect to size the VCF instance with a fixed resource allocation for DB workloads, while planning scalability for other workloads. This aligns with VMware??s design methodology for capturing stable baseline conditions.
Option B: A constraint of including the management, DB, and VDI environmentsThis is incorrect. Aconstraintis a limitation or restriction imposed on the design, such as existing hardware or policies. The need to support management, VDI, DB, and general workloads is arequirement(what the solution must do), not a limitation. Labeling it a constraint misrepresents its role—it??s a design goal, not a restrictive factor. Constraints might include budget or rack space, but this scenario doesn??t indicate such limits.
Option C: A requirement consisting of the growth of the general workloads and VDI environmentThis is a strong contender but incorrect in this context. Arequirementdefines what the solution must achieve, and the customer??s statement that ??general workloads and VDI environments are expected to grow over the next 3 years?? could be a requirement (e.g., ??The solution must support growth????). However, the question asks for a single item, and Option A better captures a foundational planning element (static DB workloads) that directly informs sizing. Growth could be a requirement, but it??s less immediate than the assumption about DB stability for initial design documentation.
Option D: A risk that the VCF instance may not have enough capacity for growthThis is incorrect as the primary answer. Ariskidentifies potential issues that could impact success, such as insufficient capacity for growing workloads. While this is a valid concern given VDI and general workload growth, the scenario doesn??t provide evidence of immediate capacity limitations—only an expectation of growth. Risks are typically documented after sizing, not as the sole initial inclusion. The assumption about DB workloads is more fundamental to start the design process.
Conclusion:The architect should includean assumption that the DB workload resource requirements will remain static(Option A). This reflects the customer??s explicit statement, establishes a baseline for sizing the Management Domain and Workload Domains, and allows planning for growth elsewhere. While growth (C) and risk (D) are relevant, the assumption is the most immediate and appropriate single item for initial documentation in VCF 5.2.
References:
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Design Assumptions and Requirements)
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Workload Domain Sizing)

NEW QUESTION 61
During a transformation project kick-off meeting, an architect highlights specific areas on which to focuswhile developing the new conceptual design. Which statement is the business requirement?

A. The solution must continue to operate even in case of an entire datacenter failure.
B. The project should use the existing storage devices within the data center.
C. Sites must support a network latency of less than 12 ms RTT.
D. There is no budget specifically assigned for disaster recovery.

Answer: A

Explanation:

Reference:VMware Cloud Foundation 5.2 Architect Study Guide, Chapter 1: Conceptual Design, Section on Identifying Business Requirements.

NEW QUESTION 64
A customer has a requirement to improve bandwidth and reliability for traffic that is routed through the NSX Edges in VMware Cloud Foundation. What should the architect recommend satisfying this requirement?

A. Configure a Load balanced Group for NSX Edges
B. Configure a TEP Group for NSX Edges
C. Configure a TEP Independent Group for NSX Edges
D. Configure a LAG Group for NSX Edges

Answer: D

Explanation:
Reference:NSX-T 3.2 Administration Guide (included in VCF 5.2), Section on Edge Networking and Link Aggregation; VMware Cloud Foundation 5.2 Networking Guide.

NEW QUESTION 67
An architect is designing a VMware Cloud Foundation (VCF)-based Private Cloud solution. During the requirements gathering workshop with customer stakeholders, the following information was captured:
The solution must be capable of deploying 50 concurrent workloads.
The solution must ensure that once submitted, each service does not take longer than 6 hours to provision.
When creating the design documentation, which design quality should be used to classify the stated requirements?

A. Availability
B. Recoverability

C. Performance
D. Manageability

Answer: C

Explanation:
In VMware Cloud Foundation (VCF) 5.2, design qualities (or non-functional requirements) categorize how the solution meets its objectives. The requirements—??deploying 50 concurrent workloads?? and??provisioning each service within 6 hours??—must be classified under a quality that reflects their intent. Let??s evaluate each option:
Option A: AvailabilityAvailability ensures the solution is accessible and operational when needed (e.g., uptime percentage). While deploying workloads and provisioning services assume availability, the requirements focus onspeedandcapacity(50 concurrent workloads, 6-hour limit), not uptime or fault tolerance. This quality doesn??t directly address the stated needs, making it incorrect.
Option B: RecoverabilityRecoverability addresses the ability to restore services after a failure (e.g., disaster recovery). The requirements don??t mention failure scenarios, backups, or restoration—they focus on provisioning speed and concurrency during normal operation. Recoverability is unrelated to these operational metrics, so this is incorrect.
Option C: PerformanceThis is the correct answer. Performance measures how well the solution executes tasks, including speed, throughput, and capacity. In VCF 5.2:
??Deploying 50 concurrent workloads?? is a throughput requirement, ensuring the system can handle multiple deployments simultaneously.
??Each service does not take longer than 6 hours to provision?? is a latency or response time requirement, setting a performance boundary.Both align with theperformancequality, which governs resource efficiency and user experience in provisioning workflows (e.g., via SDDC Manager or Aria Automation). This classification fits VMware??s design framework.
Option D: ManageabilityManageability focuses on ease of administration, monitoring, and maintenance (e.g., automation, UI simplicity). While provisioning workloads involves management, the requirements emphasizehow fastandhow many—performance metrics—not the ease of managing the process. Manageability might apply to tools enabling this, but it??s not the primary quality here.
Conclusion:The design quality to classify these requirements isPerformance(Option C). It directly reflects the solution??s ability to handle 50 concurrent workloads and provision services within 6 hours, aligning with VCF 5.2??s focus on operational efficiency. References:
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Design Qualities) VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Performance Considerations)

NEW QUESTION 70
As part of a new VMware Cloud Foundation (VCF) deployment, a customer is planning to implement the vSphere IaaS control plane. What component could be installed and enabled to implement the solution?

A. Storage DRS
B. Aria Automation
C. Aria Operations
D. NSX Edge networking

Answer: B

Explanation:
In VMware Cloud Foundation (VCF) 5.2, the vSphere IaaS (Infrastructure as a Service) control plane extends vSphere to provide cloud-like provisioning and automation, typically through integration with higher-level tools. The question asks which component enables this capability. Let??s evaluate:
Option A: Storage DRS
Storage DRS (Distributed Resource Scheduler) automates storage management (e.g., load balancing) within vSphere. It??s a vSAN/vSphere feature, not an IaaS control plane, as it lacks broad provisioning or orchestration capabilities. This is incorrect.
Option B: Aria Automation
This is correct. VMware Aria Automation (formerly vRealize Automation) integrates with VCF via SDDC Manager to provide an IaaS control plane on vSphere. It enables self- service provisioning of VMs, applications, and infrastructure (e.g., via blueprints), extending vSphere into a cloud model. In VCF 5.2, Aria Automation??s vSphere IaaS control plane feature (introduced in vSphere 7.0+) allows direct management of vSphere resources as an IaaS platform, making it the key component for this solution.
Option C: Aria Operations
Aria Operations (formerly vRealize Operations) provides monitoring and analytics for VCF. It tracks performance and health, not provisioning or IaaS control. While valuable, it doesn??t implement an IaaS control plane, so this is incorrect.
Option D: NSX Edge networking
NSX Edge provides advanced networking (e.g., load balancing, gateways) in VCF. It supports IaaS by enabling network services but isn??t the control plane itself—control planes orchestrate resources, not just network them. This is incorrect.
Conclusion:The component to install and enable for the vSphere IaaS control plane is Aria Automation (B). It transforms vSphere into an IaaS platform within VCF 5.2, meeting the customer??s deployment goal.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Aria Automation Integration)
VMware Aria Automation 8.10 Documentation (integrated in VCF 5.2): vSphere IaaS Control Plane
VMware vSphere 7.0U3 Documentation (integrated in VCF 5.2): IaaS Features

NEW QUESTION 72
An architect has come up with a list of design decisions after a workshop with the business stakeholders. Which design decision describes a logical design decision?

A. Asynchronous storage replication that satisfies a recovery point objective (RPO) of 15min between site A and B
B. Both sites A and B will have a /16 dedicated network subnets.
C. End users will interact with application server hosted in Site A
D. End users should always experience instantaneous application response

Answer: A

Explanation:

Reference:VMware Cloud Foundation 5.2 Planning and Preparation Guide, Chapter 4: Logical Design Decisions.

NEW QUESTION 75

The following requirements were identified in an architecture workshop for a virtual infrastructure design project.
REQ001: All virtual machines must satisfy the Recovery Point Objective (RPO) of fifteen
(15) minutes or less in a disaster recovery (DR) situation
REQ002: Service level availability must satisfy 99.999% measured yearly. Which two test cases will validate these requirements?

A. Simulate or invoke an outage of the primary datacente
B. All virtual machines must be restored within fifteen (15) minutes or less.
C. Simulate or invoke an outage of the primary datacente
D. All virtual machines must not lose more than one (1) hour of data prior to the outage.
E. Simulate or invoke an outage of the primary datacente
F. All virtual machines must not lose more than fifteen (15) minutes of data prior to the outage.
G. Simulate or invoke an outage of the primary datacente
H. All virtual machines must be restored within one (1) hour or less.

Answer: AC

Explanation:
Reference:VMware Cloud Foundation 5.2 Disaster Recovery Guide, Section on RPO and RTO Validation; VMware Site Recovery Manager 8.6 Documentation, Test Case Design.

NEW QUESTION 76
An architect is documenting the design for a new VMware Cloud Foundation-based solution. Following the requirements gathering workshops held with customer stakeholders, the architect has made the following assumptions:
The customer will provide sufficient licensing for the scale of the new solution.
The existing storage array that is to be used for the user workloads has sufficient capacity to meet the demands of the new solution.
The data center offers sufficient power, cooling, and rack space for the physical hosts required by the new solution.
The physical network infrastructure within the data center will not exceed the maximum latency requirements of the new solution.
Which two risks must the architect include as a part of the design document because of these assumptions? (Choose two.)

A. The physical network infrastructure may not provide sufficient bandwidth to support the user workloads.
B. The customer may not have sufficient data center power, cooling, and physical rack space available.
C. The customer may not have licensing that covers all of the physical cores the design requires.
D. The assumptions may not be approved by a majority of the customer stakeholders before the solution is deployed.

Answer: AC

Explanation:
In VMware Cloud Foundation (VCF) 5.2, assumptions are statements taken as true for design purposes, but they introduce risks if unverified. The architect must identify risks—potential issues that could impact the solution??s success—stemming from these assumptions and include them in the design document. Let??s evaluate each option against the assumptions:
Option A: The physical network infrastructure may not provide sufficient bandwidth to support the user workloadsThis is correct. The assumption states that the physical network infrastructure ??will not exceed the maximum latency requirements,?? but it doesn??t address bandwidth. In VCF, user workloads (e.g., in VI Workload Domains) rely on network bandwidth for performance (e.g., vSAN traffic, VM communication). Insufficient bandwidth could degrade workload performance or scalability, despite meeting latency requirements. This is a direct risk tied to an unaddressed aspect of the network assumption, making it a necessary inclusion.
Option B: The customer may not have sufficient data center power, cooling, and physical rack space availableThis is incorrect as a mandatory risk in this context. The assumption explicitly states that ??the data center offers sufficient power, cooling, and rack space?? for the required hosts. While it??s possible this could be untrue, the risk is already implicitly covered by questioning the assumption??s validity. Including this risk would be redundant unless specific evidence (e.g., unverified data center specs) suggests doubt, which isn??t provided. Other risks (A, C) are more immediate and distinct.
Option C: The customer may not have licensing that covers all of the physical cores the design requiresThis is correct. The assumption states that ??the customer will provide sufficient licensing for the scale of the new solution.?? In VCF 5.2, licensing (e.g., vSphere, vSAN, NSX) is core-based, and misjudging the number of physical cores (e.g., due to host specs or scale) could lead to insufficient licenses. This riskdirectly challenges the assumption??s accuracy—if the customer??s licensing doesn??t match the design??s core count, deployment could stall or incur unplanned costs. It??s a critical risk to document.
Option D: The assumptions may not be approved by a majority of the customer stakeholders before the solution is deployedThis is incorrect. While stakeholder approval is important, this is a process-related risk, not a technical or operational risk tied to the assumptions?? content. The VMware design methodology focuses risks on solution impact (e.g., performance, capacity), not procedural uncertainties like consensus. This risk is too vague and outside the scope of the assumptions?? direct implications. Conclusion:The two risks the architect must include are:
A: Insufficient network bandwidth (not covered by the latency assumption).
C: Inadequate licensing for physical cores (directly tied to the licensing assumption).These align with VCF 5.2 design principles, ensuring potential gaps in network performance and licensing are flagged for validation or mitigation.
References:
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Risk Identification)
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Network and Licensing Considerations)

NEW QUESTION 81
As part of the requirement gathering phase, an architect identified the following requirement for the newly deployed SDDC environment:
Reduce the network latency between two application virtual machines.
To meet the application owner's goal, which design decision should be included in the design?

A. Configure a Storage DRS rule to keep the application virtual machines on the same datastore.
B. Configure a DRS rule to keep the application virtual machines on the same ESXi host.
C. Configure a DRS rule to separate the application virtual machines to different ESXi hosts.
D. Configure a Storage DRS rule to keep the application virtual machines on different datastores.

Answer: B

Explanation:
The requirement is to reduce network latency between two application virtual machines (VMs) in a VMware Cloud Foundation (VCF) 5.2 SDDC environment. Network latency is influenced by the physical distance and network hops between VMs. In a vSphere environment (core to VCF), VMs on the same ESXi host communicate via the host??s virtual switch (vSwitch or vDS), avoiding physical network traversal, which minimizes latency. Let??s evaluate each option:
Option A: Configure a Storage DRS rule to keep the application virtual machines on the same datastoreStorage DRS manages datastore usage and VM

placement based on storage I/O and capacity, not network latency. ThevSphere Resource Management Guide notes that Storage DRS rules (e.g., VMaffinity) affect storage location, not host placement. Two VMs on the same datastore could still reside on different hosts, requiring network communication over physical links (e.g., 10GbE), which doesn??t inherently reduce latency. Option B: Configure a DRS rule to keep the application virtual machines on the same ESXi hostDRS (Distributed Resource Scheduler) controls VM placement across hosts for load balancing and can enforce affinity rules. A ??keep together?? affinity rule ensures the two VMs run on the same ESXi host, where communication occurs via the host??s internal vSwitch, bypassing physical network latency (typically <1µs vs. milliseconds over a LAN). TheVCF 5.2 Architectural GuideandvSphere Resource Management Guiderecommend this for latency-sensitive applications, directly meeting the requirement.
Option C: Configure a DRS rule to separate the application virtual machines to different ESXi hostsA DRS anti-affinity rule forces VMs onto different hosts, increasing network latency as traffic must traverse the physical network (e.g., switches, routers). This contradicts the goal of reducing latency, making it unsuitable.
Option D: Configure a Storage DRS rule to keep the application virtual machines on different datastoresA Storage DRS anti-affinity rule separates VMs across datastores, but this affects storage placement, not host location. VMs on different datastores could still be on different hosts, increasing network latency over physical links. This doesn??t address the requirement, per thevSphere Resource Management Guide.
Conclusion:Option B is the correct design decision. A DRS affinity rule ensures the VMs share the same host, minimizing network latency by leveraging intra-host communication, aligning with VCF 5.2 best practices for latency-sensitive workloads.References: VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Section on DRS and Workload Placement.
vSphere Resource Management Guide(docs.vmware.com): DRS Affinity Rules and Network Latency Considerations.
VMware Cloud Foundation 5.2 Administration Guide(docs.vmware.com): SDDC Design for Performance.

NEW QUESTION 85
An Architect has been tasked with reviewing a VMware Cloud Foundation design document. Observe the following requirements:
REQ01: The solution must support the private cloud cybersecurity industry and local standards and controls.
REQ02: The solution must ensure that the cloud services are transitioned to operation teams.
REQ03: The solution must provide a self-service portal.
REQ04: The solution must provide the ability to consume storage based on policies. REQ05: The solution should provide the ability to extend networks between different
availability zones.
REQ06: The solution should allow only supported versions of management solutions to be deployed.
Observe the following design decisions:
DD01: There will be a clustered deployment of Aria Automation.
DD02: There will be an integration between Aria Automation and multiple geo-located vCenter Servers.
DD03: Aria Suite Lifecycle will be deployed to provide lifecycle management of Aria Suite components.
Based on the stated requirements, what are the three implications for taking the stated design decisions? (Choose three.)

A. Aria Automation must have network access to all vCenter Servers.
B. Aria Suite Lifecycle should be deployed through the SDDC Manager.
C. An external database is required for Aria Automation clustering.
D. A load balancer is required for Aria Automation high availability.
E. The latency between the Aria Automation Appliances must be less than 2ms.
F. The vCenter Servers must have network access to each other.

Answer: ACD

Explanation:
The design decisions (DD01, DD02, DD03) must align with the requirements (REQ01-REQ06) in a VMware Cloud Foundation (VCF) 5.2 context, and the implications must reflect architectural necessities or dependencies introduced by these decisions. Let??s evaluate each option based on the requirements and decisions:
Option A: Aria Automation must have network access to all vCenter Servers Relevance:DD02 states integration between Aria Automation and multiple geo-located vCenter Servers, supporting REQ03 (self-service portal), REQ04 (policy-based storage), and REQ05 (network extension across availability zones).
Implication:Aria Automation (formerly vRealize Automation) requires network connectivity to manage vCenter Servers for workload provisioning, policy enforcement (e.g., vSphere Storage Profiles), and network extension (e.g., via NSX). TheVMware Aria Automation Installation Guidemandates that Aria Automation appliances have TCP/IP access to vCenter instances over specific ports (e.g., 443). This is a direct implication of DD02 and is critical for multi-site integration.
Conclusion:This is a necessary implication.
Option B: Aria Suite Lifecycle should be deployed through the SDDC Manager Relevance:DD03 involves deploying Aria Suite Lifecycle for lifecycle management, aligning with REQ06 (supported versions of management solutions).
Implication:While SDDC Manager in VCF can deploy and manage Aria Suite components, theVMware Cloud Foundation 5.2 Administration Guideindicates that Aria Suite Lifecycle can be deployed standalone or via SDDC Manager, depending on the design. It??s not a strict requirement (implication) of DD03—rather, it??s a deployment choice. REQ06 is satisfied by Aria Suite Lifecycle??s version control, regardless of deployment method. Conclusion:This is not a mandatory implication, as it??s not enforced by the design
decisions.
Option C: An external database is required for Aria Automation clustering Relevance:DD01 specifies a clustered deployment of Aria Automation, supporting REQ03 (self-service portal) and REQ02 (transition to operations via a robust platform). Implication:For high availability (HA) clustering, Aria Automation requires an external PostgreSQL database to synchronize state across appliances. TheVMware Aria Automation Installation Guideexplicitly states that clustering (three-node HA) mandates an external database (e.g., PostgreSQL 13) rather than the embedded one used in single- node setups. This ensures data consistency and failover, making it a direct implication of DD01.
Conclusion:This is a necessary implication.
Option D: A load balancer is required for Aria Automation high availability Relevance:DD01 involves a clustered deployment, supporting REQ03 and REQ02.
Implication:Aria Automation clustering for HA requires a load balancer (e.g., VMware NSX Advanced Load Balancer or third-party) to distribute traffic across the three appliances and provide a single access point. TheVMware Aria Automation Installation Guidemandates a load balancer for HA configurations to ensure availability and seamless failover, directly tied to DD01. This also supports operational transition (REQ02) by ensuring a reliable self- service portal (REQ03).
Conclusion:This is a necessary implication.
Option E: The latency between the Aria Automation Appliances must be less than 2ms
Relevance:DD01 (clustered deployment).
Implication:Aria Automation clustering requires low latency between appliances for database replication and cluster health. However, theVMware Aria Automation Installation Guidespecifies a maximum latency of10msbetween nodes (not 2ms), with 2ms being a recommendation for optimal performance, not a strict requirement. In a VCF context, this isn??t a mandated implication unless specified by additional constraints not present here. Conclusion:This is not a precise implication based on standard requirements.
Option F: The vCenter Servers must have network access to each other Relevance:DD02 (integration with multiple geo-located vCenter Servers).
Implication:While Aria Automation integrates with vCenter Servers, there??s no requirement in VCF or Aria Automation for vCenter Servers to communicate directly with each other across sites unless Enhanced Linked Mode or a specific multi-site feature (e.g., stretched clusters) is in use, which isn??t indicated by the requirements or decisions. REQ05 (network extension) is managed by NSX, not vCenter-to-vCenter connectivity. TheVCF 5.2 Architectural Guideconfirms vCenter Servers can operate independently under Aria Automation.

Conclusion:This is not an implication of the stated decisions.
Conclusion:The three implications are:
A: Network access from Aria Automation to vCenter Servers is required for DD02.
C: An external database is mandatory for Aria Automation clustering per DD01.
D: A load balancer is essential for HA in Aria Automation clustering per DD01.These align with the requirements and design decisions in a VCF 5.2 context.References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Aria Suite Integration and Multi-Site Design.
VMware Aria Automation Installation Guide(docs.vmware.com): Clustering Prerequisites (Database, Load Balancer, Latency).
VMware Cloud Foundation 5.2 Administration Guide(docs.vmware.com): Aria Suite Lifecycle Deployment Options.

NEW QUESTION 86
An architect decided to deploy an NSX Edge cluster using SDDC Manager. These Edges will be used by a Tier-0 Gateway configured with BGP to provide North-South connectivity in the Management Domain. Which statement justifies this design decision?

A. NSX Edges deployed via SDDC Manager can be updated separately in the future.
B. VPN service in NSX will be available and configurable via SDDC Manager with NSX Edges deployed using this method.
C. Extra Large form factor is available only when edges are deployed using SDDC Manager.
D. This deployment method will automatically configure dynamic routing.

Answer: B

Explanation:
In VMware Cloud Foundation 5.2, NSX Edge clusters provide critical networking services, such as North-South connectivity via Tier-0 Gateways, often using BGP for dynamic routing. Deploying NSX Edges via SDDC Manager integrates them into the VCF lifecycle management framework, which impacts their configuration and operational capabilities. Let??s analyze each option:
Option A: NSX Edges deployed via SDDC Manager can be updated separately in the futureIn VCF, SDDC Manager manages the lifecycle (deployment, upgrades, etc.) of NSX components, including Edge nodes. However, updates are not performed ??separately?? from the VCF stack; they are part of a coordinated upgrade process across the management domain. TheVCF 5.2 Administration Guidenotes that Edge updates are tied to NSX Manager and SDDC Manager workflows, contradicting the idea of independent updates. This doesn??t justify the design decision.
Option B: VPN service in NSX will be available and configurable via SDDC Manager with NSX Edges deployed using this methodWhen NSX Edges are deployed via SDDC
Manager in the Management Domain, they are fully integrated into the VCF architecture. This enables advanced NSX features, such as VPN services (L2VPN, IPsec VPN), to be configured and managed through SDDC Manager or NSX Manager UIs. TheVMware Cloud Foundation 5.2 Networking Guideconfirms that deploying Edges via SDDC Manager supports North-South connectivity (e.g., via Tier-0 with BGP) and additional services like VPN, providing operational flexibility. This justifies the decision by aligning with VCF??s integrated management capabilities.
Option C: Extra Large form factor is available only when edges are deployed using SDDC ManagerNSX Edge form factors (Small, Medium, Large, Extra Large) are determined by resource requirements and deployment method, but the Extra Large form factor is available whether Edges are deployed manually via NSX Manager or through SDDC Manager in VCF. TheNSX-T Data Center Installation Guide(part of VMware docs) clarifies that form factor selection is independent of the deployment tool, making this statement inaccurate and not a justification.
Option D: This deployment method will automatically configure dynamic routing Deploying Edges via SDDC Manager automates some aspects of setup (e.g., cluster creation, basicnetworking), but dynamic routing (e.g., BGP) requires manual configuration of peers, ASNs, and route maps via NSX Manager. TheVCF 5.2 Networking Guidestates that while SDDC Manager streamlines deployment, BGP configuration remains a post- deployment task, disproving ??automatic?? configuration as a justification. Conclusion:Option B is the correct justification because deploying NSX Edges via SDDC Manager ensures integration with VCF??s management plane, enabling features like VPN services alongside BGP-based North-South connectivity in the Management Domain. This aligns with the architect??s goal of leveraging VCF??s centralized management strengths. References:
VMware Cloud Foundation 5.2 Networking Guide(docs.vmware.com): Section on NSX Edge Deployment and Tier-0 Gateway Configuration.
VMware Cloud Foundation 5.2 Administration Guide(docs.vmware.com): SDDC Manager Workflows for NSX Edge Clusters.
NSX-T Data Center Installation Guide(docs.vmware.com): Edge Node Deployment Options.

NEW QUESTION 89
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 2V0-13.24 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 2V0-13.24 Product From:

https://www.2passeasy.com/dumps/2V0-13.24/

Money Back Guarantee

2V0-13.24 Practice Exam Features:

* 2V0-13.24 Questions and Answers Updated Frequently

* 2V0-13.24 Practice Questions Verified by Expert Senior Certified Staff

* 2V0-13.24 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* 2V0-13.24 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year