VMware

Exam Questions 2V0-13.24

VMware Cloud Foundation 5.2 Architect

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
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 2
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 3
Which statement defines the purpose of Business Requirements?

A. Business requirements define which audience needs to be involved.
B. Business requirements define how the goals and objectives can be achieved.
C. Business requirements define which goals and objectives can be achieved.
D. Business requirements define what goals and objectives need to be achieved.

Answer: D

Explanation:
In the context of VMware Cloud Foundation (VCF) 5.2 and IT architecture design,business requirementsarticulate the high-level needs and expectations of the organization that the solution must address. They serve as the foundation for the architectural design process, guiding the development of technical solutions to meet specific organizational goals. According to VMware??s architectural methodology and standard IT frameworks (e.g., TOGAF, which aligns with VMware??s design principles), business requirements focus onwhatthe organization aims to accomplish rather thanhowit will be accomplished orwhowill be involved. Let??s evaluate each option:
Option A: Business requirements define which audience needs to be involved.This statement is incorrect. Identifying the audience or stakeholders (e.g., end users, IT staff, ormanagement) is part of stakeholder analysis or requirements gathering, not the purpose of business requirements themselves. Business requirements focus on the goals and objectives of the organization, not the specific people involved in the process. This option misaligns with the role of business requirements in VCF design.
Option B: Business requirements define how the goals and objectives can be achieved.This statement is incorrect. Thehowaspect—detailing the methods, technologies, or processes to achieve goals—falls under the purview offunctional requirementsor technical design specifications, not business requirements. For example, in VCF 5.2, deciding to use vSAN for storage or NSX for networking is a technical decision, not a business requirement. Business requirements remain agnostic to implementation details, making this option invalid.
Option C: Business requirements define which goals and objectives can be achieved.This statement is misleading. Business requirements do not determinewhich goals are achievable (implying a feasibility assessment); rather, they statewhatthe organization intends or needs to achieve. Assessing feasibility comes later in the design process (e.g., during risk analysis or solution validation). In VCF, business requirements might specify the need for high availability or scalability, but they don??t evaluate whether those are possible—that??s a technical consideration. Thus, this option is incorrect.

Option D: Business requirements define what goals and objectives need to be achieved.This is the correct answer. Business requirements articulatewhatthe organization seeks to accomplish with the solution, such as improving application performance, ensuring disaster recovery, or supporting a specific number of workloads. In the context of VMware Cloud Foundation 5.2, examples might include ??the solution must support 500 virtual machines?? or ??the environment must provide 99.99% uptime.?? These statements define the goals and objectives without specifying how they will be met (e.g., via vSphere HA or vSAN) or who will implement them. This aligns with VMware??s design methodology, where business requirements drive the creation of subsequent functional and non-functional requirements.
In VMware Cloud Foundation 5.2, the architectural design process begins with capturing
business requirements to ensure the solution aligns with organizational needs. The VMware Cloud Foundation Planning and Preparation Guide emphasizes that business requirements establish the ??what?? (e.g., desired outcomes like cost reduction or workload consolidation), which then informs the technical architecture, such as the sizing of VI Workload Domains or the deployment of management components.
References:
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Requirements Gathering)
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Design Methodology Overview)
VMware Validated Design Documentation (Business Requirements Definition, applicable to VCF 5.2 principles)

NEW QUESTION 4
An architect is preparing a VI Workload Domain design with a dedicated NSX instance. The workload domain is planned to grow up to 300 ESXi hosts within the next six months. Which is the minimum NSX Manager form factor that should be recommended by the architect for this VI Workload Domain to support the forecasted growth?

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

Answer: A

Explanation:
Reference:NSX-T 3.2 Reference Design Guide (VCF 5.2 compatible), Section on NSX Manager Sizing; VMware Cloud Foundation 5.2 Deployment Guide, Workload Domain Sizing.

NEW QUESTION 5
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 6
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 7
An administrator is documenting the design for a new VMware Cloud Foundation (VCF) solution. During discovery workshops with the customer, the following information was shared with the architect:
All users and administrators of the solution will need to be authenticated using accounts in the corporate directory service.
The solution will need to be deployed across two geographically separate locations and run in an Active/Standby configuration where supported.
The management applications deployed as part of the solution will need to be recovered to the standby location in the event of a disaster.
All management applications will need to be deployed into a management tooling zone of the network, which is separated from the corporate network zone by multiple firewalls.
The corporate directory service is deployed in the corporate zone.
There is an internal organization policy that requires each application instance (management or end user) to detail the ports that access is required on through the firewall separately.
Firewall rule requests are processed manually one application instance at a time and typically take a minimum of 8 weeks to complete.
The customer also informed the architect that the new solution needs to be deployed and ready to start the organization??s acceptance into service process within 3 months, as it is a dependency in the deployment of a business-critical application. When considering the design for the Cloud Automationand Operations products within the VCF solution, which three design decisions should the architect include based on this information? (Choose three.)

A. The Cloud Automation and Operations products will be reconfigured to integrate with the Identity Broker solution instance at the standby site in case of a Disaster Recovery incident.
B. The Identity Broker solution will be deployed at both the primary and standby site.
C. The Identity Broker solution will be connected with the corporate directory service for user authentication.
D. The Identity Broker solution will be deployed at the primary site and failed over to the standby site in case of a disaster.
E. The Cloud Automation and Operations products will be integrated with a single instance of an Identity Broker solution at the primary site.
F. The Cloud Automation and Operations products will be integrated directly with the corporate directory service.

Answer: BCE

Explanation:
In VMware Cloud Foundation (VCF) 5.2, Cloud Automation (e.g., Aria Automation) and Operations (e.g., Aria Operations) products rely on identity management for authentication. The customer??s requirements—corporate directory authentication, Active/Standby across two sites, disaster recovery (DR), network zoning, slow firewall processes, and a 3-month deployment timeline—shape the design decisions. The architect must ensure authentication works efficiently across sites while meeting the timeline and DR
needs. Let??s evaluate:
Key Constraints and Context:
Authentication: All users/administrators use the corporate directory (e.g., Active Directory in the corporate zone).
Deployment: Active/Standby across two sites, with management apps in a separate tooling zone behind firewalls.
DR: Management apps must recover to the standby site.
Firewall Delays: 8-week minimum per rule, but deployment must occur within 12 weeks (3 months).
Identity Broker: In VCF, VMware Workspace ONE Access (or similar) acts as an identity broker, bridging VCF components with external directories (e.g., AD via LDAP/S). Evaluation of Options:
Option A: The Cloud Automation and Operations products will be reconfigured to integrate with the Identity Broker solution instance at the standby site in case of a Disaster Recovery incident
This implies a single Identity Broker at the primary site, with reconfiguration to a standby instance post-DR. Reconfiguring products (e.g., updating SSO endpoints) during DR adds complexity and downtime, contradicting the Active/Standby goal of seamless failover. It??s feasible but not optimal given the need for continuous operation and the 3-month timeline. Option B: The Identity Broker solution will be deployed at both the primary and standby site
This is correct. Deploying Workspace ONE Access (or equivalent) at both sites supports Active/Standby by ensuring authentication availability at the primary site and immediate usability at the standby site post-DR. It aligns with VCF??s multi-site HA capabilities and avoids reconfiguration delays, addressing the DR requirement efficiently within the timeline. Option C: The Identity Broker solution will be connected with the corporate directory service for user authentication
This is correct. The requirement states all users/administrators authenticate via the corporate directory (in the corporate zone). An Identity Broker (e.g., Workspace ONE Access) connects to AD via LDAP/S, acting as a proxy between the management tooling zone and corporate zone. This satisfies the authentication need and simplifies firewall rules (one broker-to-AD connection vs. multiple app connections), critical given the 8-week delay.
Option D: The Identity Broker solution will be deployed at the primary site and failed over to the standby site in case of a disaster
This suggests a single Identity Broker with DR failover. While possible (e.g., via vSphere Replication), it risks authentication downtime during failover, conflicting with Active/Standby continuity. The 8-week firewall rule delay for the standby site??s broker connection post-DR also jeopardizes the 3-month timeline and DR readiness, making this less viable than dual- site deployment (B).
Option E: The Cloud Automation and Operations products will be integrated with a single instance of an Identity Broker solution at the primary site
This is correct. Integrating Aria products with one Identity Broker instance at the primary site during initial deployment simplifies setup and meets the 3-month timeline. It leverages the broker deployed at the primary site (part of B) for authentication, minimizing firewall rules (one broker vs. multiple apps). Pairing this with a standby instance (B) ensures DR readiness without immediate complexity.
Option F: The Cloud Automation and Operations products will be integrated directly with the corporate directory service
This is incorrect. Direct integration requires each product (e.g., Aria Automation, Operations) to connect to AD across the firewall, necessitating multiple rule requests. With an 8-week minimum per rule and several products, this exceeds the 3-month timeline. It also complicates DR, as each app would need re-pointing to a standby AD, violating efficiency and zoning policies.
Conclusion:
The three design decisions are:
B: Identity Broker at both sites ensures Active/Standby and DR readiness.
C: Connecting the broker to the corporate directory fulfills the authentication requirement and simplifies firewall rules.
E: Integrating products with a primary-site broker meets the 3-month deployment goal while leveraging B and C for DR.This trio balances timeline, security, and DR needs in VCF 5.2. References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Identity and Access Management)
VMware Aria Automation 8.10 Documentation (integrated in VCF 5.2): Authentication Design
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Multi-Site and DR Considerations)

NEW QUESTION 8
A customer is deploying VCF at a new datacenter location. They will migrate their workloads from the existing datacenter to the new VCF platform over six months.

Both datacenters will run simultaneously for six months during the migration. Which of the following should be a documented risk?

A. Six months may not be enough time to complete the migration.
B. There will be connectivity between the two locations.
C. Bandwidth between the two locations is sufficient to accommodate the workload migration.
D. Workloads will be powered off during migration.

Answer: A

Explanation:
Reference:VMware Cloud Foundation 5.2 Planning and Preparation Guide, Chapter 5: Risk Assessment; VMware Migration Best Practices for VCF.

NEW QUESTION 9
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 10
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 10
The following requirements were identified in an architecture workshop for a virtual infrastructure design project.
REQ001: All virtual machines must meet the Recovery Time Objective (RTO) of twenty- four hours or less in a disaster recovery (DR) scenario.
Which two test cases will verify these requirements?

A. Simulate or trigger an outage of the primary datacente
B. All virtual machines must be restored within four hours or less.
C. Simulate or trigger an outage of the primary datacente
D. All virtual machines must be restored within twenty-four hours or less.
E. Simulate or trigger an outage of the primary datacente
F. All virtual machines must not lose more than twenty-four hours of data prior to the outage.
G. Simulate or trigger an outage of the primary datacente
H. All virtual machines must not lose more than four hours of data prior to the outage.

Answer: BC

Explanation:
Reference:VMware Cloud Foundation 5.2 Disaster Recovery Guide, RTO Validation; VMware SRM 8.6 Documentation, Test Case Scenarios.

NEW QUESTION 15
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 17
During a requirements gathering workshop, several Business and Technical requirements were captured from the customer. Which requirement is classified as a Technical Requirement?

A. Reduce system processing time for service requests by 25%.
B. The system must support 5,000 concurrent users.
C. Increase customer satisfaction by 15%.
D. Expand market reach to include new geographical regions.

Answer: B

Explanation:
In VMware Cloud Foundation (VCF) architecture, requirements are categorized as Business or Technical based on their focus. Technical requirements specify measurable system capabilities or constraints, directly influencing design decisions for infrastructure components like compute, storage, or networking. Business requirements, conversely, focus on organizational goals or outcomes that IT supports. Option B, "The system must support 5,000 concurrent users," is a technical requirement because it defines a specific system capacity metric (concurrent users), which directly impacts scalability and resource allocation in VCF design, such as the sizing of workload domains or NSX configurations. Option A, "Reduce system processing time for service requests by 25%," could be technical but is often a derivative of a business goal (efficiency), making it less explicitly technical in this context. Options C and D, focusing on customer satisfaction and market reach, are clearly business-oriented, tied to organizational outcomes rather than system specifications.
Reference: VMware Cloud Foundation 5.2 Architect Study Guide, Chapter 2: Requirements
Gathering and Analysis, Section on Classifying Requirements.

NEW QUESTION 21
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 26
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 31
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 tenants will be configured for Development and Production.
B. Users?? access to resources will be controlled by tenant membership.
C. Users?? access to resources will be controlled by project membership.
D. Separate cloud zones will be configured for Development and Production.

Answer: CD

Explanation:
In VMware Cloud Foundation (VCF) 5.2, Aria Automation (formerly vRealize Automation) manages resource provisioning and access control. The requirements involve role-based access, environment isolation, and workload placement flexibility. Let??s analyze each option:
Option A: Separate tenants will be configured for Development and ProductionAria Automation in VCF 5.2 operates as a single-tenant application by default, integrated with SDDC Manager and vCenter. Multi-tenancy (separate tenants) is an advanced configuration typically used for service providers, not standard VCF private cloud designs. TheVMware Aria Automation Installation Guidenotes that multi-tenancy adds complexity and isn??t required for environment segregation within a single organization. Instead, projects and cloud zones handle these needs, making this unnecessary.
Option B: Users?? access to resources will be controlled by tenant membership
Tenant membership applies in multi-tenant setups, where users are assigned to distinct tenants (e.g., Dev vs. Prod). Since VCF 5.2 typically uses a single tenant, and the requirements can be met with projects (group-based access), this isn??t a must-have decision. TheVCF 5.2 Architectural Guidefavors project-based access over tenant separation for organizational control, rendering this optional.
Option C: Users?? access to resources will be controlled by project membership Projects in Aria Automation group users and define their access to resources (e.g., cloud zones, policies). To meet the first requirement (access based on company organization) and the second (developers provisioning only to Development), projects can restrict developers to a ??Dev?? project linked to a Development cloud zone, while other teams (e.g., ops) access Production/DMZ via

separate projects. TheVMware Aria Automation Administration Guideconfirms projects as the primary mechanism for role-based access in VCF, making this a required decision.
Option D: Separate cloud zones will be configured for Development and Production Cloud zones in Aria Automation map to vSphere clusters or resource pools (e.g., Development, Production, DMZ clusters). To satisfy the second requirement (developers limited to Development) and the third (Production workloads on DMZ or Production clusters), separate cloud zones ensure environment isolation and placement flexibility. The VCF 5.2 Architectural Guidemandates cloud zones for workload segregation, tying them to projects for access control, making this essential.
Conclusion:
C: Project membership enforces user access per organization and restricts developers to Development, meeting the first two requirements.
D: Separate cloud zones isolate Development from Production/DMZ, enabling precise workload placement per the third requirement.These decisions align with Aria Automation??s design in VCF 5.2.References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Aria Automation Design and Cloud Zones.
VMware Aria Automation Administration Guide(docs.vmware.com): Projects and Access Control.
VMware Aria Automation Installation Guide(docs.vmware.com): Tenancy Options in VCF.

NEW QUESTION 34
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 35
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.
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.
Based on the information provided, which two requirements satisfy the stated design decisions? (Choose two.)

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

Answer: CE

Explanation:
In VMware Cloud Foundation (VCF) 5.2, VMware Aria Automation (formerly vRealize Automation) enhances the platform by providing self-service, automation, and multi-site management capabilities. The architect must determine which requirements (REQ01-REQ05) are directly satisfied by the design decisions (DD01 and DD02). Let??s evaluate each requirement against the decisions:
Design Decisions:
DD01: Clustered deployment of Aria Automation
A clustered deployment ensures high availability and scalability of Aria Automation, supporting multiple users and workloads with resilience.
DD02: Integration between Aria Automation and multiple geo-located vCenter Servers
This enables centralized management of distributed vSphere environments (e.g., across availability zones or regions), facilitating network and resource orchestration. Evaluation of Requirements:
Option A: REQ01 - The solution must support the private cloud cybersecurity industry and local standards and controls
This requirement focuses on cybersecurity and compliance (e.g., encryption, access controls, auditing). While Aria Automation supports role-based access control (RBAC) and integrates with secure VCF components, neither DD01 nor DD02 directly addresses cybersecurity standards or local controls. These are typically met by VCF??s baseline security features (e.g., NSX, vSphere hardening), not specifically by Aria Automation??s clustering or vCenter integration. Thus, REQ01 is not directly satisfied by the stated decisions.
Option B: REQ02 - The solution must ensure that the cloud services are transitioned to operation teams
This requirement implies operational handoff, training, or automation to enable operations teams to manage services. Aria Automation??s clustering (DD01) improves reliability, and vCenter integration (DD02) centralizes management, but neither explicitly ensures a transition process (e.g., documentation, runbooks). This is more about operational processes than the technical decisions provided, so REQ02 is not directly satisfied. Option C: REQ03 - The solution must provide a self-service portal
This is correct. Aria Automation??s primary function in VCF 5.2 is to provide a self-service portal for users to provision and manage resources (e.g., VMs, applications). A clustered deployment (DD01) ensures theportal??s availability and scalability, supporting multiple users concurrently. Integration with vCenter Servers (DD02) enhances its capability to deploy resources across sites, but DD01 alone directly satisfies REQ03 by enabling a robust self-service experience.

Thus, REQ03 is satisfied.
Option D: REQ04 - The solution must provide the ability to consume storage based on policies
This requirement involves policy-driven storage management (e.g., vSAN storage policies).
Aria Automation supports storage policies via integration with vSphere/vSAN, allowing users to define storage profiles (e.g., performance, capacity). However, this capability is inherent to vSphere/vSAN integration, not uniquely tied to clustering (DD01) or geo-located vCenter integration (DD02). While Aria Automation facilitates this, the design decisions don??t specifically address storage policy consumption as a primary outcome, making REQ04 less directly satisfied compared to others.
Option E: REQ05 - The solution should provide the ability to extend networks between different availability zones
This is correct. Integrating Aria Automation with multiple geo-located vCenter Servers (DD02) enables management of distributed environments, including network extension across availability zones. In VCF 5.2, this leverages NSX-T for Layer 2 stretching (e.g., via HCX or NSX Federation), orchestrated through Aria Automation. DD02 directly supports this by connecting disparate vCenters, allowing network policies and extensions to be applied across zones. Clustering (DD01) supports scalability but isn??t the key factor—DD02 is the primary enabler. Thus, REQ05 is satisfied.
Conclusion:
The two requirements satisfied by the design decisions are:
REQ03 (C): A clustered Aria Automation deployment (DD01) directly provides a reliable self-service portal.
REQ05 (E): Integration with multiple geo-located vCenter Servers (DD02) enables network extension across availability zones.While REQ04 is partially supported, REQ03 and REQ05 are the most directly tied to the stated decisions in the VCF 5.2 context. 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): Self-Service Portal and Multi-Site Management
VMware NSX-T 3.2 Reference Design (integrated in VCF 5.2): Network Extension Capabilities

NEW QUESTION 39
A VMware Cloud Foundation design incorporates the following technical requirements:
All management components must have their login sessions timeout after 2 minutes of inactivity.
Communication between management components should be limited to required ports only.
Modifications required by compliancy should not impact the management components?? functionality.
What would be the recommendation from a design perspective that would aid in achieving the above requirements?

A. Consult the vSphere Security Configuration kit
B. Leverage the results of a vulnerability assessment and apply the recommendations
C. Consult the Compliance Kit for VMware Cloud Foundation
D. Apply NSX DFW (Distributed Firewall) to achieve zero-trust

Answer: C

Explanation:
Reference:VMware Cloud Foundation 5.2 Compliance Kit Documentation, Security Configuration Section; VMware Cloud Foundation 5.2 Security Guide.

NEW QUESTION 40
An architect is tasked with designing a new VMware Cloud Foundation environment and has identified the following customer-provided requirements:
REQ01: The application server must handle at least 30,000 transactions per second. REQ02: The design must meet ISO 27001 information security standards.
REQ03: The storage network should maintain a minimum latency of 12 milliseconds before path failover.
REQ04: The staging environment should utilize a secondary third-party data center. REQ05: Planned maintenance must be performed outside the hours of 8 AM to 8 PM GMT. What are the two functional requirements? (Choose two.)

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

Answer: AD

Explanation:
In VMware Cloud Foundation (VCF) 5.2, requirements are classified as functional(what the system must do) ornon-functional(how the system performs or operates). Functional requirements describe specific capabilities or behaviors, while non-functional requirements address qualities like performance, security, or constraints. Let??s classify each:
Option A: REQ01 - The application server must handle at least 30,000 transactions per second
This is correct. This is afunctional requirementbecause it specifies what the application server (a component of the solution) must do—process a defined transaction volume. It??s a capability the system must deliver, directly tied to workload performance within the VCF environment.
Option B: REQ02 - The design must meet ISO 27001 information security standards This is anon-functional requirement. ISO 27001 addresses security qualities (e.g., confidentiality, integrity), defininghowthe system should operate securely, not what it does. It??s a compliance and operational constraint, not a functional capability.
Option C: REQ03 - The storage network should maintain a minimum latency of 12 milliseconds before path failover
This is anon-functional requirement. It specifies a performance threshold (latency) and reliability behavior (failover), describinghowthe storage network should perform, not a specific function it must provide.
Option D: REQ04 - The staging environment should utilize a secondary third-party data center
This is correct. This is afunctional requirementbecause it defines what the solution must include—a staging environment located in a specific secondary data center. It??s a capability or structural requirement of the VCF deployment, dictating a functional aspect of the system.
Option E: REQ05 - Planned maintenance must be performed outside the hours of 8 AM to 8 PM GMT
This is anon-functional requirement. It??s an operational constraint onwhenmaintenance occurs, affecting availability and manageability, not a specific function the system must perform.
Conclusion:The two functional requirements areREQ01 (A)andREQ04 (D). They define what the VCF solution must do (handle transactions, include a staging environment), aligning with VMware??s design methodology for functional specifications.
References:
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Functional vs. Non-Functional Requirements)
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Requirements Classification)

NEW QUESTION 43
During a requirements gathering workshop, several Business and Technical requirements were captured from the customer. Which requirement will be classified

as a Business Requirement?

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

Answer: A

Explanation:
In VMware??s design methodology (aligned with VCF 5.2), requirements are categorized asBusiness Requirements(goals tied to organizational outcomes, often non- technical) orTechnical Requirements(specific system capabilities or constraints). Let??s classify each option:
Option A: Reduce processing time for service requests by 30%This is a Business Requirement. It focuses on a business outcome—improving service request efficiency by a measurable percentage—without specifying how the system achieves it. TheVMware Cloud Foundation 5.2 Architectural Guideclassifies such high-level, outcome-driven goals as business requirements, as they reflect the customer??s operational or strategic priorities rather than technical implementation details.
Option B: The system must support 10,000 concurrent usersThis is a Technical Requirement. It specifies a measurable system capability (supporting 10,000 concurrent users), directly tied to performance and capacity. VMware documentation treats such quantifiable system behaviors as technical, focusing on ??what?? the system must do functionally.
Option C: Data must be encrypted using AES-256 encryptionThis is a Technical Requirement. It mandates a specific technical implementation (AES-256 encryption) for security, a non-functional attribute. TheVCF 5.2 Design Guidecategorizes encryption standards as technical constraints or requirements, not business goals.
Option D: The application must be compatible with Windows, macOS, and Linux operating systemsThis is a Technical Requirement. It defines a functional capability—cross-platform compatibility—specifying technical details about the system??s operation. VMware classifies such compatibility needs as technical, per the design methodology.
Conclusion:Option A is the Business Requirement, as it aligns with a business goal (efficiency improvement) rather than a technical specification.References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Section on Requirements Gathering and Classification.
VMware Cloud Foundation 5.2 Design Guide(docs.vmware.com): Business vs. Technical Requirements.

NEW QUESTION 47
An architect is designing a new VMware Cloud Foundation (VCF)-based Private Cloud solution. During the requirements gathering workshop, a stakeholder from the network team stated that:
The solution must ensure that any physical networking component is redundant to N+N. The solution must ensure inter-datacenter network links are diversely routed.
When writing the design documentation, how should the architect classify the stated requirement?

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

Answer: A

Explanation:
In VMware Cloud Foundation (VCF) 5.2, design qualities (non-functional requirements) categorizehowthe system operates. The network team??s requirements focus on redundancy and routing diversity, which the architect must classify. Let??s evaluate: Option A: Availability
This is correct. Availability ensures the solution remains operational and accessible. ??N+N redundancy?? (e.g., dual active components where N failures are tolerated by N spares) for physical networking components eliminates single points of failure, ensuring continuous network uptime. ??Diversely routed inter-datacenter links?? prevents outages from a single path failure, enhancing availability across sites. In VCF, these align with high-availability network design (e.g., NSX Edge uplink redundancy), makingavailabilitythe proper classification.
Option B: Performance
Performance addresses speed, throughput, or latency (e.g., ??10 Gbps links??). Redundancy and diverse routing might indirectly support performance by avoiding bottlenecks, but the primary intent is uptime, not speed. This doesn??t fit the stated requirements?? focus.
Option C: Recoverability
Recoverability focuses on restoring service after a failure (e.g., backups, failover time). N+N redundancy and diverse routingpreventdowntime rather than recover from it. While related, the requirements emphasize proactive uptime (availability) over post-failure recovery, making this incorrect.
Option D: Manageability
Manageability concerns ease of administration (e.g., monitoring, configuration). Redundancy and routing diversity are infrastructure design choices, not management processes. This quality doesn??t apply.
Conclusion:The architect should classify the requirement asAvailability (A). It ensures the VCF solution??s network remains operational, aligning with VCF 5.2??s focus on resilient design.
References:
VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Design Qualities) VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Network Availability)

NEW QUESTION 51
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 56
The following design decisions were made relating to storage design:
• A storage policy that would support failure of a single fault domain being the server rack
• Two vSAN OSA disk groups per host each consisting of four 4TB Samsung SSD capacity drives
• Two vSAN OSA disk groups per host each consisting of a single 300GB Intel NVMe cache drive
• Encryption at rest capable disk drives
• Dual 10Gb or faster storage network adapters
Which two design decisions would an architect include within the physical design? (Choose two.)

A. A storage policy that would support failure of a single fault domain being the server rack
B. Two vSAN OSA disk groups per host each consisting of a single 300GB Intel NVMe cache drive
C. Encryption at rest capable disk drives
D. Dual 10Gb or faster storage network adapters
E. Two vSAN OSA disk groups per host each consisting of four 4TB Samsung SSD capacity drives

Answer: DE

Explanation:

Reference:VMware Cloud Foundation 5.2 vSAN Design Guide, Physical Storage Design; VMware vSAN 7.0 Planning and Deployment Guide.

NEW QUESTION 57
A design requirement has been specified for a new VMware Cloud Foundation (VCF) instance. All managed workload resources must be lifecycle managed with the following criteria:
• Development resources must be automatically reclaimed after two weeks
• Production resources will be reviewed yearly for reclamation
• Resources identified for reclamation must allow time for review and possible extension What capability will satisfy the requirements?

A. Aria Suite Lifecycle Content Management
B. Aria Operations Rightsizing Recommendations
C. Aria Automation Lease Policy
D. Aria Automation Project Membership

Answer: C

Explanation:

Reference:VMware Aria Automation 8.10 Administration Guide, Section on Lease Policies;
VMware Cloud Foundation 5.2 Architect Study Guide, Automation Features.

NEW QUESTION 62
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 64
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 66
.....

Relate Links

100% Pass Your 2V0-13.24 Exam with Exambible Prep Materials

https://www.exambible.com/2V0-13.24-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/