



VMware VCP-VCF Architect 2V0-13.25 Practice Questions

Shared by Walker on 17-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

What open source project does vSphere Supervisor use to automate the lifecycle management of VMware Kubernetes Service (VKS) clusters?

Options:

- A- Cluster API
- B- Grafana
- C- Contour
- D- Kubeadm



Answer:

A

Explanation:

According to the VMware Cloud Foundation 9.0.4 Architecture and Design Guide, the vSphere Supervisor leverages the Cluster API open-source project to provide declarative, Kubernetes-style APIs for the creation, configuration, and lifecycle management of VMware Kubernetes Service (VKS) clusters.

The documentation states:

"The Cluster API provides declarative, Kubernetes-style APIs for cluster creation, configuration, and management. The inputs to Cluster API include a resource describing the cluster, a set of resources describing the virtual machines that make up the cluster, and a set of resources describing cluster add-ons."

This API-driven model allows for automated provisioning and scaling of Kubernetes clusters by defining the desired state through YAML manifests. Cluster API (CAPI) is a CNCF open-source project that VMware has extended and integrated into vSphere Supervisor to deliver infrastructure-level automation, ensuring consistency, repeatability, and lifecycle management for VKS clusters.

In contrast, Grafana is used for monitoring, Contour for ingress control, and Kubeadm is a bootstrapping tool for standalone Kubernetes clusters --- none of which provide cluster lifecycle automation in VKS.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.4 Architecture and Design Guide --- "VKS Architecture and Components." (pp. 5635--5637)

VMware Cloud Foundation 9.0.4 Supervisor Components --- "Cluster API Integration for VKS Lifecycle Management."

Question 2

Question Type: MultipleChoice

An architect is tasked with designing a VMware Cloud Foundation (VCF) solution for a financial services organization to modernize its core banking applications and high-frequency trading systems using vSAN.

The following requirements were gathered:

- * For critical transactional database workloads, the solution must provide low-latency and high performance storage.
- * For all non-critical workloads, the solution must provide the most efficient capacity utilization.

Which three design decisions would the architect make to meet the requirements for the workload domain cluster? (Choose three.)

Options:

- A- Configure vSAN Policies (RAID-5) for all critical transactional database workloads.
- B- Deploy a vSAN OSA (All-NVMe) cluster with a minimum of 4 nodes.
- C- Deploy a vSAN ESA cluster with a minimum of 6 nodes.
- D- Configure vSAN Policies (RAID-5/6) for all non-critical workloads.
- E- Configure vSAN Policies (RAID-1) for all workloads.
- F- Configure vSAN Policies (RAID-1) for all critical transactional database workloads.

Answer:

C, D, F

Explanation:

RAID-1 policies offer low latency and high performance, ideal for critical workloads such as transactional databases.

RAID-5/6 provides efficient capacity usage, which is preferred for non-critical workloads.

VMware recommends the vSAN Express Storage Architecture (ESA) with 6 nodes minimum for performance and redundancy.

The ESA supports better performance with fewer overheads than OSA. VMware Cloud Foundation 9.0 promotes using vSAN ESA for mission-critical workloads when latency is a priority.

Question 3

Question Type: MultipleChoice

An architect is designing a solution with Istio Service Mesh.

What two types of groups can collect and manage objects? (Choose two.)

Options:

- A- Service
- B- Cluster
- C- Security
- D- API
- E- Node

Answer:

A, B

Explanation:

The VMware Cloud Foundation 9.0.4 Architecture and Design Guide, in the section covering Istio Service Mesh integration, defines that objects and configurations in a service mesh are grouped and managed under Service and Cluster scopes.

"In an Istio Service Mesh deployment, objects are collected into groups to simplify policy management and routing. These groups can represent logical services or clusters to which policies apply, allowing consistent control across the environment."

Service groups organize workloads that share a common application or microservice. They enable the consistent application of routing, observability, and security policies (e.g., mTLS, rate limiting).

Cluster groups define the boundary for multi-cluster management, enabling federated control and workload discovery across multiple Kubernetes clusters within the VMware Cloud Foundation ecosystem.

These group types allow centralized configuration, policy enforcement, and telemetry

aggregation across workloads deployed in VCF-managed Kubernetes environments.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.4 Architecture and Design Guide --- "Istio Service Mesh Architecture and Grouping Constructs." (pp. 5798--5802).

VMware Cloud Foundation 9.0.4 --- "Service and Cluster Group Management in Multi-Cluster Istio Deployments."

Question 4

Question Type: MultipleChoice

An architect has compiled a list of statements following a workshop with the business stakeholders.

Which statement would be included in a conceptual model?

Options:

- A- The solution must meet a Mean Time To Recovery (MTTR) of 6 hours.
- B- Sites A and B will each have a stretched Layer-2 for their management network.
- C- The `das.isolationshutdowntimeout` setting will be configured to 120 seconds.
- D- Users will connect to the application servers via the NSX Advanced Load Balancer.

Answer:

D

Explanation:

As per the VMware Cloud Foundation 9.0.1 Design Guide (Conceptual, Logical, and Physical Model Framework), a conceptual model defines what the solution will deliver at a high level --- focusing on major components and interactions without specifying implementation details or configurations.

The statement "Users will connect to the application servers via the NSX Advanced Load Balancer" fits this definition because it identifies a functional relationship between solution components (users and applications) and introduces the NSX ALB as a key service enabler --- without going into configuration or performance specifics.

In contrast:

Option A (MTTR 6 hours) defines a service level target, part of non-functional design.

Option B (stretched L2 networks) and C (vSphere HA parameter tuning) belong to the logical and physical design stages respectively, as they involve implementation specifics.

Therefore, D represents a conceptual model statement, describing service delivery intent and architectural interaction.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.1 Design Guide --- "Conceptual, Logical, and Physical Design Layers Explained."

VMware Cloud Foundation 9.0.2 Architecture Design Standards --- "Conceptual Model: High-Level Service Definition and Interaction."

Question 5

Question Type: MultipleChoice

Which VMware Cloud Foundation (VCF) Storage Model can be deployed to scale storage capacity independent of compute and network?

Options:

- A- vSAN Compute Cluster Storage Model
- B- vSAN ESA Storage Model
- C- vSAN Capacity Cluster Model
- D- vSAN ESA Storage Cluster Model

Answer:

A

Explanation:

As defined in the VMware Cloud Foundation 9.0.1 Architecture and Design Guide, the vSAN Compute Cluster Storage Model is designed to allow compute and storage scaling to occur independently. The guide specifies:

"A vSAN Compute Cluster is a vSphere cluster that does not contribute storage resources to the vSAN datastore. Instead, it consumes storage from a remote vSAN cluster through HCI Mesh. This architecture enables compute and network resources to scale independently of storage."

This design provides flexibility in disaggregated architectures, allowing organizations to expand compute clusters without needing to expand vSAN capacity nodes and vice versa. Such a model is particularly useful for data-intensive workloads or environments with unpredictable growth patterns in compute versus storage.

Other options, such as vSAN ESA Storage Model or vSAN ESA Storage Cluster Model, represent Enhanced Storage Architecture configurations but do not inherently separate compute scaling. Hence, the correct answer is the vSAN Compute Cluster Storage Model, as it uniquely supports independent scalability while maintaining centralized management under the VCF architecture.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.1 Architecture and Design Guide --- "vSAN Compute Cluster Model (Disaggregated Storage Architecture)" (pp. 870--875).

VMware Cloud Foundation 9.0.2 Design Guide --- "Compute and Storage Independent Scaling using HCI Mesh."

Question 6

Question Type: MultipleChoice

During an initial design workshop with stakeholders, the architect was provided with an overview of the current state and other information required to proceed to the design phase.

The architect has assumed that the solution will need to support high availability for workloads.

Given the assumption, which statement should the architect document as a risk?

Options:

- A- The solution supports the separation of management components from production workloads.
- B- BGP is the dynamic routing protocol on the physical fabric and cannot be changed.
- C- The solution supports a recovery point objective (RPO) of 24 hours for infrastructure components.
- D- The entire infrastructure is hosted on a single physical site.

Answer:

D

Explanation:

Comprehensive and Detailed Explanation (Based on VMware Cloud Foundation 9.0.1 Design

Guide):

The VMware Cloud Foundation 9.0.1 Design Guide clearly states that a single-site deployment introduces inherent risks when availability requirements extend beyond single host or rack failures. It specifies:

"A single-site VCF deployment tolerates host, rack, and component failure, but does not provide protection against total site failure. This design profile is not well-suited for disaster recovery or high availability across sites."

Since the architect assumed that the solution must support high availability for workloads, but the environment is hosted entirely on a single physical site, the solution would not meet this assumption in the event of a site outage. This discrepancy represents a design risk --- a potential event or condition that could prevent the design from fulfilling the availability requirement.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.1 Design Guide --- VCF Fleet in a Single Site Blueprint and Risk Considerations for HA and DR.

VMware Cloud Foundation 9.0.2 Architecture Overview --- Availability and Recoverability Design Qualities.

Question 7

Question Type: MultipleChoice

An architect is designing a new VMware Cloud Foundation (VCF) solution. They are meeting with the key stakeholders and subject matter experts (SMEs) for the first time as part of the requirements gathering process. The following information has been shared with the architect prior to the meeting:

Names and job titles of the attendees

Project timelines and budget

What step should the architect perform as part of this initial requirements gathering workshop?

Options:

- A- Open the meeting with a diagram of the VCF topology that must meet the customer requirements in order to start a discussion.
- B- Ask questions to agree on the key product features the SMEs want from the design.
- C- Open the meeting with a list of the VCF design decisions from the public documentation to agree on any required changes.
- D- Ask questions to start a discussion on the business objectives and desired outcomes.

Answer:

D

Explanation:

The first workshop focuses on understanding business drivers, objectives, and outcomes before diving into design specifics. VMware design methodology stresses that architects should start with business requirements and then move toward conceptual, logical, and physical designs.

A and C are premature --- topology and design decisions should not be presented before business outcomes are captured.

B focuses on product features, which is technical and should come later.

D is correct because identifying business objectives ensures the architecture aligns with customer goals.

Question 8

Question Type: MultipleChoice

An architect is responsible for designing a new VMware Cloud Foundation (VCF)-based Private Cloud solution. During the requirements gathering workshop with key customer stakeholders, the following information was captured:

* The solution must support the monitoring of up to 10,000 objects.

* The solution must support 24 months retention for all monitoring data.

When creating the design document, which design quality should be used to classify the stated requirements?

Options:

A- Performance

B- Manageability

C- Availability

D- Recoverability

Answer:

B

Explanation:

The VMware Cloud Foundation 9.0.1 Design Guide defines Manageability as the design quality that focuses on monitoring, configuration, reporting, and operational management capabilities. It states:

"Manageability encompasses the ability to monitor, configure, and control the solution effectively, ensuring operational insights and compliance with business objectives."

The requirement to support 10,000 monitored objects and retain monitoring data for 24 months directly relates to the system's ability to manage and retain operational metrics, not to its performance or availability. These parameters ensure long-term visibility, compliance, and data analytics consistency---core elements of Manageability.

In contrast:

Performance focuses on throughput and latency.

Availability ensures uptime and fault tolerance.

Recoverability deals with restoration after failure.

Therefore, the classification of these requirements falls under Manageability, as they directly define how the system manages and retains monitoring data.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.1 Architecture Guide --- "Design Qualities: Availability, Manageability, Performance, Recoverability, and Security." (pp. 72--74)

VMware Cloud Foundation 9.0.2 Design Framework --- "Classification of Monitoring and Retention Requirements under Manageability."

Question 9

Question Type: MultipleChoice

An architect was in an architectural workshop and noted the following business objectives:

The solution must prioritize optimal end-to-end user shopping experience for customers accessing the website.

The website must be available 24 x 7 x 365.

Which three conceptual model items relate to these business objectives? (Choose three.)

Options:

- A- A requirement to have 99.99% availability uptime measured at the front-end application layer
- B- A risk that the external internet network provider does not meet the service level agreement (SLA) requirements
- C- A requirement to have 99.99% availability uptime measured at the network infrastructure layer
- D- An assumption that site performance is not a key performance indicator (KPI) for the customer
- E- A constraint of any planned changes limited to outside of business hours only
- F- An assumption that there is sufficient budget for the design to meet the performance requirements

Answer:

A, B, F

Explanation:

Conceptual model items are categorized as requirements, assumptions, risks, and constraints:

A is a clear requirement, aligned with the business objective of continuous uptime and availability.

B is a valid risk, as the availability of the external internet provider impacts the user experience.

F is an assumption, implying that budget will not be a limitation in achieving the desired level of performance.

Items like C are too technically specific for the conceptual layer. D directly contradicts the business objective, and E relates more to operational constraints rather than high-level design goals.

VMware Cloud Foundation Conceptual Design Guide -- Requirements, Assumptions, Constraints, Risks (RACR) Framework

=====

Question 10

Question Type: MultipleChoice

While designing a solution, an architect is tasked with defining limits for a vSphere Namespace.

What three limits are available? (Select three.)

Options:

- A- The amount of storage
- B- The amount of containers
- C- The amount of services
- D- The amount of memory
- E- The amount of CPU

Answer:

A, D, E

Explanation:

According to the VMware Cloud Foundation 9.0.4 Documentation (p.5802--5804), administrators can configure resource limits for vSphere Namespaces, which define how much CPU, memory, and storage can be consumed.

The guide specifies under "Set Resource Limits to a vSphere Namespace":

"As a vSphere administrator, you can set resource limits and container defaults on a vSphere Namespace. Options include CPU, Memory, and Storage. Container defaults set CPU limits, CPU requests, memory requests, and memory limits for containers within the Namespace."

These limits provide governance and ensure that workloads deployed within a namespace do not exceed predefined consumption levels, maintaining cluster resource balance and preventing noisy-neighbor scenarios.

The document further clarifies that storage limits are applied per storage policy and CPU and memory limits define the total reserved compute capacity available to workloads and pods running in the namespace.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.4 --- "Set Resource Limits to a vSphere Namespace" (p.5802-5803)

VMware Cloud Foundation 9.0.4 --- "Self-Service Namespace Template Provisioning Workflow."

To Get Premium Files for 2V0-13.25 Visit

<https://www.p2pexams.com/products/2v0-13.25>

For More Free Questions Visit

<https://www.p2pexams.com/vmware/pdf/2v0-13.25>

20%
DISCOUNT

P2P
exams