



CertsMania

Free Questions for 2V0-13.25

Shared by **Aitor** on **Dec 2, 2025**

For More Free Questions and Preparation Resources

Check the Links on Last Page



CertsMania

Questions # 1:

As part of an initial stakeholder meeting, one of the stakeholders has stated the following:

The initial design must be completed within the next 3 months so that hardware can be ordered within the current budget cycle.

How would the architect classify and record this statement?

Options:

A.

A constraint

B.

A risk

C.

An assumption

D.

A requirement

Answer

A

Explanation

This is a **constraint**, as it defines a non-negotiable time limit imposed by the customer's budgeting timeline. It restricts the design phase's schedule and deliverables. In VMware conceptual modeling, timing constraints are explicitly captured as constraints rather than requirements or assumptions.

[Reference: VMware Cloud Foundation Conceptual Design Guide - Project Timeline and Constraints, ,]

Questions # 2:

When designing a backup and recovery solution for VKS clusters, which tool can be leveraged to back up and restore workloads?

Options:

A.

Site Recovery Manager

B.

Velero

C.

Restic

D.

VMware Live Recovery



CertsMania

Answer

B

Explanation

According to the **VMware Cloud Foundation 9.0.4 Architecture and Design Guide**, the officially supported solution for backing up and restoring workloads running on **VMware Kubernetes Service (VKS)** clusters is **Velero**, specifically the **Velero Plugin for vSphere**. The documentation explicitly states:

“You can backup and restore workloads running on VKS clusters by using Velero. The Velero Plugin for vSphere provides a solution for backing up and restoring workloads running on vSphere Pods and VKS clusters.”

Velero integrates with vSphere through a plugin that manages **persistent volume snapshots** and **metadata backups** to an **S3-compatible object store**. This provides data protection for both **stateless** and **stateful Kubernetes workloads**.

While **Restic** is mentioned in conjunction with Velero for additional support in some environments, it cannot operate standalone in VKS environments. **Site Recovery Manager** and **VMware Live Recovery** are used for virtual machine-based disaster recovery, not containerized workloads. Therefore, Velero is the **correct and VMware-supported tool** for VKS workload backup and recovery.

References (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.4 — *“Backing Up and Restoring VKS Clusters and*

Workloads” (pp. 5879–5887)

VMware Cloud Foundation 9.0.4 — “*Install and Configure the Velero Plugin for vSphere on a Supervisor*” (pp. 5880–5885)

Questions # 3:

An architect is working on a VMware Cloud Foundation (VCF) architecture design and identified the following requirements:

- The organization is using a third-party virtual appliance that does not support overlay networks.
- The virtual appliance must reside on the same L2 domain as an external physical firewall.
- The virtual appliance also needs access to workloads that are currently hosted on overlay segments provided by NSX.

Which design decision should the architect make to meet these requirements?

Options:

A.

Request the third-party vendor to certify the virtual appliance for NSX Overlay segments.

B.

Connect the virtual appliance to a VLAN-backed segment and configure NSX bridging to allow access to overlay segments.

C.

Place the virtual appliance and all workloads on VLAN-backed segments.

D.

Connect the virtual appliance to an overlay-backed segment and use static routes to the firewall.

Answer

B

Explanation

According to the **VMware Cloud Foundation 9.0.2 Design Guide (NSX Bridging and Layer 2 Connectivity)**, when workloads or appliances require Layer 2 adjacency with physical network devices and do not support overlay encapsulation (Geneve), **NSX Edge bridging** is the validated solution.

The documentation states:

“The bridge functionality extends an overlay segment into a VLAN identified by a VLAN ID on an uplink of the NSX Edge. This enables Layer 2 connectivity between virtual machines on NSX overlay segments and physical workloads on VLAN-backed networks.”

By connecting the third-party appliance to a **VLAN-backed segment** and configuring **NSX bridging**, the appliance gains direct L2 connectivity to the firewall, while the workloads on overlay segments can still communicate seamlessly through the bridge.

This approach maintains the logical separation of NSX networks while satisfying L2 adjacency requirements, a common design pattern for integrating non-overlay-capable systems.

References (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.2 Architecture and Design Guide — *“Edge Bridging: Extending Overlay Segments to VLAN.”*

VMware Cloud Foundation 9.0.2 NSX Design Recommendations — *“Bridge Profiles and VLAN Uplink Configuration.”*

Questions # 4:

During a VMware Cloud Foundation (VCF) architectural design workshop, one of the stakeholders made the following comment:

“The company has just used the remaining budget to purchase eight vSAN Ready Nodes for this project.”

How would the architect classify this statement within the conceptual model document?

Options:

A.

Requirement

B.

Risk

C.

Assumption

D.

Constraint



CertsMania

Answer

D

Explanation

This statement expresses a financial limit — “Remaining budget” — constraining future expenditures on hardware. This is clearly a **constraint**, as it restricts the design options (e.g., can't procure new hardware). In the VMware Conceptual Model framework, constraints are factors that limit design choices without introducing risk or goal definitions.

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

Questions # 5:

An architect is responsible for designing a VMware Cloud Foundation (VCF)-based private cloud for a customer. During the customer requirements gathering workshop, the customer has stated the following:

All Platinum applications/services must have an availability SLA of 99.99%.

All Gold applications/services must have an availability SLA of 99.9%.

All Silver applications/services must have an availability SLA of 99%.

The private cloud must have an availability SLA of 99.9%.

What should the architect recommend to meet the stated requirements?

Options:

A.

The private cloud must only be used to host Silver and Gold services.

B.

The private cloud SLA can only be met using multiple VMware Cloud Foundation instances configured as a single VCF Fleet.

C.

The Platinum service availability requirements must be met by the application.

D.

The Platinum service availability requirements must be met by configuring Proactive High Availability (HA) on the workload domain.

Answer

C

Explanation

The overall private cloud SLA is **99.9%**, but **Platinum applications require 99.99%**, which exceeds the infrastructure SLA. VMware design methodology states:

If an application's SLA is **higher than the infrastructure SLA**, then **the application itself must provide the additional resilience** (e.g., clustering, active-active deployment, or application-level HA).

Options A and B reduce flexibility or impose unnecessary complexity.

Option D (Proactive HA) improves resiliency but cannot guarantee **99.99%**, since Proactive HA still depends on underlying infrastructure SLA.

Therefore, the only valid recommendation is that **Platinum workloads must implement resilience at the application layer**.

[Reference: VMware Cloud Foundation 9.0 – Availability and SLA Design Guide., ,]

Questions # 6:

An architect has compiled a list of design choices following a design workshop with the

business stakeholders.

Which statement represents a **logical design decision**?

Options:

A.

Synchronous data replication will be used to meet the stated Recovery Point Objective (RPO) between site A and B.

B.

Users must experience application availability in under 2 seconds.

C.

Sites A and B will each have a /16 subnet for their networks.

D.

Users must connect to the application servers via a shared Global Load Balancer.

Answer

A

Explanation

Logical design decisions define **how conceptual requirements are realized** by technology features and configurations.

A is a logical design choice because it specifies **how replication will be configured (synchronous)** to meet an RPO requirement.

B is a **requirement**, not a design decision.

C is a **physical design choice**, as it specifies subnetting (/16 networks).

D is an **application-level requirement**, not a VCF infrastructure logical decision.

Thus, the **logical design decision** is **using synchronous replication** to satisfy RPO.

[Reference: VMware Cloud Foundation 9.0 – Logical vs Physical Design Distinctions., ,]

Questions # 7:

An architect is designing a VMware Cloud Foundation (VCF)-based solution. The company policy mandates that all VCF patches and upgrades must be tested in a development environment before applying to production.

Which VCF construct design decision would comply with this mandate?

Options:

A.

Deploy two VCF vSphere Clusters within a VCF Domain.

B.

Deploy two VCF Instances within a VCF Fleet.

C.

Deploy two VCF Domains within a VCF Instance.

D.

Deploy two VCF Fleets within a VCF Private Cloud.

Answer

B

Explanation

Deploying **two VCF Instances** allows isolation between environments. One instance can serve as **development/test** and the other as **production**.

This separation enables patch/upgrade validation without impacting production. VCF architecture supports managing multiple VCF instances in a **Fleet** for centralized visibility and policy enforcement.

VMware explicitly recommends separating environments for lifecycle testing in environments with strict change control policies.

Questions # 8:

An organization is evacuating their current datacenter and moving all workloads to a new datacenter. The organization has a total of 800 workloads to move, and the migration must be

completed with no downtime within a planned change window that is scheduled to occur in four weeks.

What migration method will meet the requirements?

Options:

A.

Cross vCenter vMotion

B.

HCX OS Assisted Migration

C.

HCX Bulk Migration

D.

HCX Replication Assisted vMotion



CertsMania

Answer

D

Explanation

The requirements specify:

Large-scale migration (800 workloads).

No downtime.

Strict change window (4 weeks).

HCX Replication Assisted vMotion (RAV) is the only method that combines **large-scale migration efficiency** with **zero downtime**:

Workloads are replicated in advance.

During the cutover, a **vMotion-like switchover** completes with no outage.

Other methods:

Cross vCenter vMotion works for live moves but is not efficient for **hundreds of VMs** within a short timeline.

HCX OS Assisted Migration is for physical-to-virtual or non-vSphere workloads (not applicable).

HCX Bulk Migration is efficient but introduces downtime during cutover.

Thus, **HCX RAV** best meets the no-downtime, large-scale requirement.

[Reference: VMware HCX Migration Guide - Replication Assisted vMotion Use Cases., ,]

Questions # 9:

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**.

References (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.”

Questions # 10:

Requirement: The solution must include high security hardening levels to meet military compliance standards.

Which two physical design decisions will meet this security requirement in the workload domain? (Choose two.)

Options:

A.

The vSAN storage policy will be configured as Secondary Failures to Tolerate = 1.

B.

VCF Operations will be configured to renew the SSL certificate for vCenter Server per security policies.

C.

NTP will be configured to the internal NTP servers of 192.168.12.1 and 192.168.24.1.

D.

The advanced setting UserVars.SuppressShellWarning will be configured to 0 across all ESXi hosts.

E.

The certificate of the VI workload domain vCenter Server will be issued by RootCA.Military.Domain.Com.

Answer

D, E

Explanation

D: Setting UserVars.SuppressShellWarning = 0 ensures **warnings are shown when shell access is used**, improving auditability and compliance.

E: Using **certificates signed by a trusted military Root CA** aligns with security compliance standards for certificate trust and chain of custody.

Options A and C are operational best practices, but not **military hardening**. B is about operations automation, not compliance-level hardening.

[Reference: VMware Cloud Foundation 9.0 Security Hardening Guide., ,]

Questions # 11:

As part of the VMware Cloud Foundation (VCF) logical design, the architect has determined that the VCF Private Cloud will encompass multiple VCF instances contained within a single VCF Fleet. The architect documented the following requirements when using VCF Operations:

Monitoring downtime must be minimized.

Alerting downtime must be minimized.

Which design decision supports these requirements?

Options:

A.

Deploy two VCF Operations instances and configure the Aggregator Management Pack.

B.

Deploy VCF Operations using the Simple model with Collector nodes at remote sites.

C.

Deploy VCF Operations using the High Availability model with Collector nodes at remote sites.

D.

Deploy a single VCF Operations instance across a multi-VCF instance fleet.

Answer

C

Explanation

The **High Availability (HA) deployment model** of VCF Operations ensures that both **monitoring and alerting services are resilient** to node failure. Deploying **Collector nodes** at remote sites enables **local data collection**, reducing WAN dependency and ensuring data is not lost during network interruptions.

This configuration aligns perfectly with the need to minimize **monitoring and alerting downtime**, which is critical in distributed, multi-instance environments.

[Reference:VMware Aria Operations for VCF Design and Deployment Guide – HA and Remote Collection Models, ,]

Questions # 12:

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.

References (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*.”

Questions # 13:

A customer has a new initiative to build a private cloud based on VMware Cloud Foundation (VCF). The customer technical team is presenting an overview of the current state of the infrastructure as well as describing what the expectations are for the private cloud.

Based on the notes captured by the architect, which statement should be documented as a constraint?

Options:

A.

The existing storage is out of hardware vendor maintenance.

B.

No funding exists for a new storage array. Therefore, existing storage hardware must be used.

C.

The design must address security zone requirements for management, production, dev/test, and QA workloads.

D.

The design must provide a centralized management console to manage both data centers.

Answer

B

Explanation

Constraints are **design limitations** that cannot be changed and must be worked around.

B (no funding for new storage, so existing must be used) is a **clear constraint**, as it restricts the architect from proposing new storage hardware.

A (out of maintenance) represents a **risk** (unsupported hardware may fail).

C and **D** are **requirements**, not constraints, because they describe desired functionality of the solution.

Thus, the correct constraint is that **existing storage must be used due to funding limitations**.

[Reference: VMware Cloud Foundation 9.0 – Conceptual Design, RACR Framework: Constraints Section.,]

Questions # 14:

During the design workshop, the customer stated the following requirement:

- The solution must comply with the organization's security standards.

Which two design decisions should be included in the logical design for the workload domain?
(Choose two.)

Options:

A.

Use large-size NSX Edge virtual appliances to account for the additional firewall rules.

B.

Enable VM Monitoring for each workload within the cluster.

C.

Enable Inter-SR iBGP routing.

D.

Use an SHA-2 algorithm or higher when signing certificates.

E.

Establish an operations practice to capture and update the thumbprint of the NSX Local Manager certificate on the NSX Global Manager every time the certificate is updated.

Answer

D, E

Explanation

SHA-2 or higher certificate signing ensures cryptographic compliance with modern security standards.

Maintaining updated certificate thumbprints is a crucial operational security task to prevent man-in-the-middle attacks between NSX managers.

These measures directly align with security hardening guidance outlined in the VCF design practices.

[Reference: VCF 9.0 Security Design Guide – Certificate Management and NSX Security Practices, =====]

Questions # 15:

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.

[Reference:VMware Cloud Foundation Conceptual Design Guide - Requirements, Assumptions, Constraints, Risks (RACR) Framework, =====, ,]



CertsMania

To Get Premium Files for 2V0-13.25 Visit

<https://www.certsmania.com/vmware/2v0-13-25-practice>

For More Free Questions Visit

<https://www.certsmania.com/vmware/pdf/2v0-13-25>



CertsMania