



## Get Free VMware 2V0-13.25 Dumps PDF Questions

Why risk failure? Download updated VMware Cloud Foundation 9.0 Architect exam PDF questions today. Practice with real 2V0-13.25 dumps and verified answers designed to help you ace your certification quickly using PrepBolt 2V0-13.25 exam pdf questions and answers.

Thank you for Downloading 2V0-13.25 exam PDF Demo

<https://prepbolt.com/2V0-13.25.html>

QUESTIONS & ANSWERS  
**DEMO VERSION**  
*(LIMITED CONTENT)*

# Question 1

---

Question Type: MultipleChoice

---

An architect responsible for creating the automation design for a VMware Cloud Foundation (VCF) Private Cloud is reviewing the notes from a customer design workshop. The customer has provided the following information:

- \* The customer's existing fleet management instance will be upgraded to maintain the existing process for virtual machine deployments.
- \* The customer would like to limit the total active resource consumption per VCF Automation user.
- \* The customer would like to ensure requests meet company requirements prior to deployment for certain users.

A combination of which two VCF Automation policies should the architect recommend to meet the customer's stated requirements? (Choose two.)

## Options:

---

- A- IaaS Policy
- B- Approval Policy
- C- Resource Quota Policy
- D- Deployment Limit Policy
- E- Lease Policy

## Answer:

---

B, C

## Explanation:

---

In VMware Cloud Foundation 9.0.4, automation and governance are enforced through VCF Automation policies built on VMware Aria Automation Service Broker.

Two policies directly address the customer's requirements:

Resource Quota Policy --- Controls total active resource consumption by defining limits on CPU, memory, and storage at the organization, project, or user level.

"Resource quota policies control the amount of resources that are available to your users. You define these policies to limit the resources consumed by each user, project, or the organization."

Approval Policy --- Ensures governance and validation before resource deployment or modification.

"Approval policies are used to control deployment and day-2 action requests before they are

executed. Requests are evaluated and approved by designated users or roles to ensure they meet organizational policies."

Deployment Limit Policies (option D) apply to per-deployment limits but do not govern total user-level resource consumption. Therefore, the correct combination of policies that satisfy both governance and consumption control requirements is Resource Quota Policy and Approval Policy.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.4 Architecture and Design Guide --- Automation Governance: Resource Quota Policies and Approval Policies.

VMware Cloud Foundation 9.0.4 Operations Guide --- Policy Configuration in VCF Automation for All Apps.

VMware Cloud Foundation 9.0.4 Policy Enforcement Reference --- Approval Policy Enforcement Examples.

## Question 2

---

Question Type: MultipleChoice

---

An architect is designing for a VMware Cloud Foundation (VCF) Instance. The following requirements and constraints were documented:

- \* The management domain cluster utilizes vSAN stretched as the principal storage.
- \* Company policy states that compute and storage capacity utilization must not exceed 90% at all times.

Which three statements should the architect consider when designing the solution to satisfy the requirements? (Choose three.)

### Options:

---

- A- Use a homogenous cluster configuration.
- B- Size and monitor the cluster for a maximum compute peak utilization of < 45%.
- C- Use a heterogeneous cluster configuration.
- D- Size and monitor the cluster for a maximum storage utilization of 40%.
- E- Size and monitor the cluster for a maximum compute peak utilization of < 90%.
- F- Size and monitor the cluster for a maximum storage utilization of 90%.

### Answer:

---

A, B, D

## Explanation:

---

The VCF 9.0.1 Design Guide provides sizing best practices for vSAN stretched clusters, emphasizing capacity planning to accommodate failover conditions. It states:

"When sizing a stretched vSAN cluster, consider that in case of a site failure, all resources from the failed site must be supported by the surviving site. Therefore, compute and storage should not exceed 50% utilization per site."

This translates into a maximum target utilization of ~45% compute and ~40% storage per site, ensuring compliance with the company's overall 90% capacity policy when combined across both sites.

A homogeneous cluster configuration (A) ensures consistent performance and predictable recovery during failover.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.1 Design Guide --- "vSAN Stretched Cluster Capacity Planning and Sizing." (pp. 874--885)

VMware vSAN Design and Sizing Reference --- "Cluster Utilization Planning for Failover Scenarios."

## Question 3

---

Question Type: MultipleChoice

---

An architect is designing a Business Continuity Disaster Recovery (BCDR) strategy for a Virtual Cloud Foundation (VCF) environment with a management domain and multiple workload domains deployed in two datacenters located in the same city.

During one of the initial workshops with stakeholders, the following information was identified:

The Recovery Time Objective (RTO) for workloads is 24 hours.

The management domain must remain continuously available with Recovery Point Objective (RPO) of 0.

Hardware overhead should be minimized by utilizing standby resources that host test workloads during normal operation.

Operational overhead should be minimized.

Latency between both datacenters is 2 ms.

Which design decision should the architect document to satisfy provided requirements?

### Options:

---

- A- Use VCF Automation to redeploy the entire environment in case of a failure.
- B- Implement vSAN stretched cluster for the management domain and Live Recovery for the workload domains.
- C- Back up all workloads daily and store them in a central repository to meet RTO expectations.
- D- Use asynchronous replication for both management and workload domains.

### Answer:

---

B

### Explanation:

---

To ensure zero RPO and high availability for the management domain, the best-fit design is to use a vSAN Stretched Cluster. With a 2ms latency and shared witness site, this design enables synchronous replication and automatic failover, ensuring no data loss (RPO 0) and minimal downtime.

For the workload domains where 24-hour RTO is acceptable, Live Recovery (leveraging replication and automation like Site Recovery Manager or Aria Automation Orchestrator) can be used to minimize operational effort and still meet recovery timelines. This strategy also aligns with minimizing hardware overhead by using the standby test infrastructure as failover capacity.

VMware Cloud Foundation Multisite Design Guide -- Stretched Clusters and Disaster Recovery

VMware vSAN Stretched Cluster Architecture Design

=====

## Question 4

---

Question Type: MultipleChoice

---

As part of the VMware Cloud Foundation (VCF) logical design, the architect documented the following requirement:

\* The solution must be able to support latency-sensitive workloads.

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

### Options:

---

- A- Intel TDX and AMD's SEV-SNP integration

- B- Advanced Memory Tiering with NVMe: Enabled
- C- vSAN Global Deduplication: Enabled
- D- NSX Enhanced Data Path: Enabled
- E- vSAN Deep Snapshots: Enabled

**Answer:**

---

B, D

**Explanation:**

---

Advanced Memory Tiering with NVMe optimizes memory performance, essential for workloads that are latency-sensitive.

NSX Enhanced Data Path (EDP) is specifically designed for high-performance, low-latency packet processing, critical in real-time applications.

Options like Global Deduplication or Deep Snapshots can negatively impact latency due to additional processing overhead.

=====

## Question 5

---

Question Type: MultipleChoice

---

Which Container Network Interface (CNI) is selected by default in a VMware Kubernetes Service (VKS) workload cluster?

**Options:**

---

- A- Antrea
- B- Flannel
- C- Calico
- D- Cilium

**Answer:**

---

A

**Explanation:**

---

According to the VMware Cloud Foundation 9.0.4 Architecture and Design Guide, under the section "VKS Cluster Networking", VMware Kubernetes Service (VKS) supports two Container Network Interface (CNI) options --- Antrea and Calico --- for pod and service networking. The documentation explicitly states:

"A VKS cluster provisioned by the VKS supports two CNI options: Antrea (default) and Calico."

It further explains that Antrea is automatically selected as the default and requires no explicit configuration unless the user specifies Calico in the cluster YAML file.

Antrea, developed by VMware, leverages Open vSwitch (OVS) for networking and supports advanced features such as network policy enforcement, encapsulation, and traffic visibility, making it the optimal default CNI for VKS clusters. Calico, using Linux bridge and BGP, remains an alternative for organizations requiring specific network topologies or integrations.

Thus, unless otherwise defined, all new VKS workload clusters in VMware Cloud Foundation are deployed using Antrea as the default network provider to maintain consistency, integration with NSX networking, and full compatibility with VCF lifecycle management.

Reference (VMware Cloud Foundation documents):

VMware Cloud Foundation 9.0.4 Design Guide -- "VKS Cluster Networking: Antrea (default) and Calico options."

VMware Cloud Foundation 9.0.4 -- "Set the Default CNI for VKS Clusters" (System-defined default: Antrea).

# Thank You for trying 2V0-13.25 PDF Demo

To try our 2V0-13.25 practice exam software  
visit link below

<https://prepbolt.com/2V0-13.25.html>

## Start Your 2V0-13.25 Preparation

Use Coupon "SAVE50" for extra 50% discount on the purchase of Practice Test Software. Test your 2V0-13.25 preparation with actual exam questions.