

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : 2V0-13.25

**Title : VMware Cloud Foundation
9.0 Architect**

Version : DEMO

1. An architect is responsible for designing a VMware Cloud Foundation (VCF)-based private cloud for a customer.

The architect noted the following requirements during a design workshop:

Co-locate application workloads with VCF management component workloads within the same vSphere cluster.

Shared storage data is always available and 100% current in the event of a single site outage.

Have two sites available no more than 10 miles apart (10ms latency) connected with high-speed network technology to host their virtual infrastructure.

Protect against outages of a single site designated as an availability zone.

Which two storage technologies could meet the stated requirements? (Choose two.)

- A. NVMe over TCP
- B. NVMe over Fibre Channel (FC)
- C. VMFS on Fibre Channel (FC)
- D. vSAN
- E. vSphere Virtual Volumes (vVols)

Answer: D, E

Explanation:

According to VMware Cloud Foundation 9.0 documentation, when a stretched cluster is deployed across sites with sub-5ms latency and high-speed interconnects, vSAN can be configured for zero RPO (Recovery Point Objective), ensuring 100% data consistency and availability in the event of a site failure. vSAN supports co-locating management and application workloads and provides the shared storage functionality with automatic failover capabilities.

Additionally, vSphere Virtual Volumes (vVols) provide granular control of virtual machine storage, and when backed by a storage system that supports replication and failover across sites (with support for VASA 3.0 or later), vVols can meet the same requirements for data availability and disaster recovery.

Reference: VMware Cloud Foundation 9.0.1 Reference Architecture Guide – Multi-Site Design and Availability Zones

VMware vSAN 8 ESA/OSA Architecture – Stretched Cluster Requirements

2. 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 a yearly workload growth of up to 10%.

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

- A. Performance
- B. Availability
- C. Manageability
- D. Security

Answer: A

Explanation:

The requirement specifying "yearly workload growth of up to 10%" relates directly to the system's ability to handle increased demand over time, which falls under the design quality of Performance.

Performance in VMware Cloud Foundation design includes considerations for scalability and the ability

to sustain projected growth. This requirement addresses the system's capacity to manage future workload expansion without degradation in service levels.

Reference: VMware Cloud Foundation Architecture and Design Guide 9.0 – Design Qualities Section: Performance and Scalability

3. An architect is responsible for designing a new VMware Cloud Foundation (VCF)-based private cloud. During the discovery workshops, the following information was captured from key customer stakeholders:

The private cloud will operate with three different monitoring levels:

Self-Managed Service: VM construct only

OS Managed Service: OS metrics and VM construct

Fully Managed Service: Infrastructure apps + OS metrics + VM construct

Approved infrastructure applications include: Microsoft IIS, SQL Server, MySQL, PostgreSQL, Tomcat Server, and Apache HTTPD

Each workload must subscribe to a monitoring level

Minimal management overhead is required for agent operations

Which two design decisions should the architect make to meet the stated monitoring requirements? (Choose two.)

- A. Configure the Service Discovery for all workloads that subscribe to the Self-Managed service
- B. Deploy the Managed Telegraf Agent for all workloads that subscribe to the Fully Managed service
- C. Deploy the Managed Telegraf Agent for all workloads that subscribe to the OS Managed service
- D. Deploy the Managed Telegraf Agent for all workloads that subscribe to the Self-Managed service
- E. Deploy the Open Source Telegraf Agent for all workloads that subscribe to the Fully Managed service

Answer: B, C

Explanation:

The Managed Telegraf Agent provides full observability and is centrally managed through VMware Aria Operations (formerly vRealize Operations). It is ideal for scenarios where both OS and infrastructure application metrics are needed, such as in the Fully Managed and OS Managed services. This meets the requirement for monitoring while minimizing manual overhead. The Self-Managed service does not require agent deployment, as only VM-level metrics are needed.

Reference: VMware Aria Operations for VCF 9.0 Deployment Guide – Monitoring Levels and Telegraf Agent Management

VMware Cloud Foundation Operations Monitoring Levels Explained

4. During a design workshop, the security team provides the following requirement for the VMware Cloud Foundation (VCF) Automation deployment:

All Virtual Machine images must be reviewed and vetted by the security team prior to consumption.

Which Content Library type supports the requirement?

- A. Subscribed Content Library
- B. Tenant-managed Content Library
- C. Local Content Library
- D. Provider-managed Content Library

Answer: D

Explanation:

A Provider-managed Content Library is curated and controlled centrally by the VCF administrator or provider. It enables vetting and version control over VM templates and OVF's, ensuring compliance with organizational policies, including security reviews. This directly supports the requirement that all images must be reviewed and approved before being used in deployments.

Reference: VMware Cloud Foundation 9.0 Planning and Preparation Guide – Content Library Management
VMware Aria Automation Content Management Best Practices

5. A VMware Cloud Foundation (VCF) architect is planning for the expansion of an existing VCF instance. The existing VCF instance is deployed with a single workload domain. The number of ESXi hosts has grown to the maximum number the existing vCenter can support.

Which design decision would the architect need to make to allow the existing VCF Instance to add more ESXi hosts?

- A. Deploy a second vCenter server appliance within the existing workload domain
- B. Deploy a second workload domain within the existing VCF Instance
- C. Deploy a second cluster within the existing vCenter
- D. Deploy a second VCF Instance within the existing VCF Fleet

Answer: B

Explanation:

A single workload domain in VCF maps to a single vCenter instance. When the host limit for that vCenter is reached (typically ~1000 hosts per vCenter), the correct and supported scale-out design is to deploy a second VI workload domain, which comes with its own dedicated vCenter instance, allowing continued expansion without affecting the existing domain.

Reference: VMware Cloud Foundation 9.0 Design Guide – Workload Domain Scaling Considerations
VMware Configuration Maximums for vCenter Server