
Exam : 3V0-12.26

Title : Advanced VMware Cloud
Foundation 9 Architect

Version : DEMO

1. An architect is designing a VMware Cloud Foundation (VCF) single sign-on solution with the following requirements:

- The provider should be able to handle failure of a single node.
- Some VCF components will not be able to access the management domain vCenter.
- The solution should be deployed using the fewest nodes possible, while still addressing all other requirements.

Which deployment mode for the single sign-on provider should be used?

- A. VCF LDAP Services (Embedded Mode)
- B. VCF Identity Broker (Embedded Mode)
- C. VCF LDAP Services (Appliance Mode)
- D. VCF Identity Broker (Appliance Mode)

Answer: D

Explanation:

D is correct. VMware Cloud Foundation 9 uses VCF Identity Broker as the SSO provider, and it supports two deployment modes: Embedded and Appliance. The appliance deployment is specifically intended for environments requiring greater availability or where VCF components cannot reliably reach the management-domain vCenter.

The Embedded Mode runs the Identity Broker inside the management-domain vCenter. Although this minimizes infrastructure footprint, it ties authentication availability to that vCenter. This conflicts directly with the requirement that some VCF components cannot access the management-domain vCenter.

Broadcom's design guidance explicitly identifies this condition as a reason to choose the appliance deployment model. (VMware Blogs)

The Appliance Mode deploys VCF Identity Broker as an external, resilient cluster. Broadcom describes the appliance implementation as a three-node Identity Broker cluster, providing resiliency and allowing authentication services to survive failure of a single Identity Broker node. It can also provide identity services independently of the management-domain vCenter's availability. (VMware Blogs)

Because three nodes are the minimum appliance topology needed to satisfy the stated single-node-failure requirement, this also meets the requirement to deploy the fewest nodes possible while retaining resiliency.

Options A and C are incorrect because VCF LDAP Services is not the VCF 9 SSO provider deployment model. B lacks the required independence and node-level resiliency.

Study Guide References/Topics: VCF Single Sign-On; VCF Identity Broker Deployment Models; Embedded vs Appliance Mode; Identity Broker High Availability; Management Domain Dependency; Identity and Access Architecture.

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

- The solution must support secure communication.

Which design decision should be included in the logical design for the workload domain?

- A. Verify all physical components used for the deployments are on the hardware compatibility list.
- B. Ensure the host servers have TPM 2.0 hardware.
- C. Use a SHA-2 algorithm or higher for signed certificates.
- D. Set promiscuous mode port group security policy to Reject.

Answer: C

Explanation:

C is correct. Secure communication between VCF components depends heavily on trusted certificates and cryptographically strong signing algorithms. VMware Cloud Foundation design guidance explicitly recommends using SHA-2 or higher for signed certificates. The official VCF design documentation states: "Use a SHA-2 algorithm or higher for signed certificates." It further explains that SHA-1 is considered less secure and has been deprecated. (Adobe Tech Docs)

This directly maps to the stated requirement because certificates are used to establish authenticated and encrypted TLS communication between VCF management and workload-domain components. A sufficiently strong certificate-signing algorithm protects certificate integrity and reduces exposure to cryptographic weaknesses associated with legacy algorithms.

A concerns hardware compatibility and supportability rather than secure communications.

B provides hardware-rooted security capabilities such as measured boot and attestation but does not directly implement secure component-to-component communication.

D is a valid Layer 2 security hardening setting that prevents unauthorized traffic inspection, but it addresses virtual-switch security rather than the broader requirement for authenticated secure communications throughout the VCF platform.

The VMware Cloud Foundation Architect exam guide explicitly includes workload-domain logical-design decisions within the Plan and Design objective. (Broadcom Docs)

Study Guide References/Topics: VCF Workload Domain Logical Design; Certificate Management Design Recommendations; SHA-2 Certificate Signing; TLS/SSL Security; Security Architecture.

3.A technology provider offers standardized digital workspaces to hundreds of corporate customers. Their internal provisioning engine must deliver a uniform setup experience, free of manual intervention, regardless of customer size or complexity. The team also mentions concerns about inconsistent identity handling, but these are secondary to the challenge of keeping tenant deployments repeatable and automated. When Clients request custom integrations, the engineering team rejects these because deviations from the standard pattern would break long-term operational stability.

Based on the above scenario, what are the two design traits the architect needs to focus on in the VMware Cloud Foundation (VCF) design? (Choose two.)

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

Answer: B C

Explanation:

Manageability (B) is the dominant design trait. The scenario emphasizes standardized deployment, repeatability, automation, minimal manual intervention, and rejection of one-off custom integrations. These are classic manageability concerns because the objective is to make the platform consistently deployable, maintainable, and operationally sustainable across hundreds of tenants. VMware Cloud Foundation explicitly uses automation and orchestration to standardize infrastructure deployment, provisioning, patching, and lifecycle operations, reducing configuration drift and administrative effort. (VMware) Security (C) is the second applicable trait because the scenario explicitly identifies inconsistent identity handling as a concern. Identity consistency, authentication, access control, credential governance, and centralized identity federation are security-design considerations. Current

VCF guidance emphasizes centralized identity services, least-privilege access, and consistent identity governance as core security controls. (VMware Blogs) Availability is not the primary concern because the scenario does not specify uptime, redundancy, or failure tolerance. Performance is not driven by latency, throughput, or response-time requirements. Recoverability is also not stated because there are no RPO, RTO, backup, or restoration requirements.

Study Guide References/Topics: AMPRS Design Qualities; Manageability; Security; Automation and Standardization; Identity Management; Configuration Consistency; Operational Governance.

4. An architect is designing the network architecture for a multi-site VMware Cloud Foundation (VCF) deployment using NSX Federation. The environment spans two geographically dispersed data centers (Site A and Site B) to support disaster recovery, workload mobility, and consistent security policies across sites. During the design workshop, the architect determined that the solution must leverage NSX Federation for stretched networking to support overlay transport, routing, and high availability. Which five design decisions are required when implementing NSX Federation in a multi-site VCF environment? (Choose five.)

- A. Configure a stretched Global Tier-0 Gateway for North/South connectivity of spanning segments with BGP.
- B. Limit round-trip latency between sites to 150 ms or less to enable optimal workload mobility and policy synchronization.
- C. Ensure end-to-end MTU of at least 1,700 bytes on the underlay network to support Geneve encapsulation for federated overlays.
- D. Require NSX Edge nodes at both sites for all federated constructs, even when using Distributed Transit Gateways for VPCs.
- E. Configure separate local Tier-1 gateways for site-specific segments, in addition to a spanning Tier-1 for cross-site communication.
- F. Deploy a Global Manager in one site to centrally configure and synchronize spanning Tier-0 and Tier-1 gateways.
- G. Use DHCP IP pools for host and Edge Tunnel Endpoints (TEPs).
- H. Ensure end-to-end MTU of at least 1,500 bytes on the underlay network to support Geneve encapsulation for federated overlays.

Answer: A B C E F

Explanation:

NSX Federation uses a Global Manager to centrally define and synchronize global networking and security objects across the Local Managers at participating sites, making F fundamental to the architecture. Broadcom confirms that configuration is performed through the active Global Manager and synchronized to the applicable locations. (Broadcom TechDocs)

A is required for stretched North/South routing: a Global Tier-0 can span locations, with site-local service routers and BGP connectivity to the physical network. Federation also uses RTEPs on Edge Nodes for cross-site Geneve transport.

B reflects the important 150-ms cross-location Edge-to-Edge RTT limit for stretched networking. Newer NSX releases permit higher latency between some manager components, but the inter-site Edge data path remains subject to the tighter latency consideration.

C is correct because NSX overlay networking requires additional MTU for Geneve encapsulation.

Broadcom specifies 1600 bytes as the minimum and 1700 bytes as the recommended value, making

1700 the appropriate design target.

E separates site-local routing from stretched workloads: local Tier-1 gateways serve site-specific segments, while a Global/spanning Tier-1 provides cross-site connectivity. Broadcom confirms that local Tier-1 gateways can connect to a Global Tier-0 in Federation.

D is too absolute, particularly with distributed VPC connectivity models. G is incorrect because Edge TEP addressing is normally based on NSX IP pools/static allocation rather than requiring DHCP. H provides insufficient MTU for Geneve overlays.

Study Guide References/Topics: NSX Federation Architecture; Global and Local Managers; Stretched Tier-0/Tier-1 Gateways; RTEPs; Multi-Site BGP Routing; Geneve MTU; Local versus Global Networking Objects; Multi-Site Availability.

5. An architect is responsible for designing a VMware Cloud Foundation (VCF) Private Cloud solution. During a requirements gathering workshop, the customer supplied the following information:

- The solution must support the monitoring of all existing workloads.
- There are currently 15,000 virtual machine workloads running within the existing environment.
- All infrastructure components will also be monitored by the solution.

The solution must ensure that the 99.9% uptime Service Level Agreement can be met.

The solution must be resilient to a single fault domain failure.

The following logical design decisions have been made within the design:

- Deploy the VCF Operations cluster following the High Availability Model.

The following table is from the VCF Operations sizing guide:

Metrics	Small	Medium
Single-Node Object Maximum	10000	30000
Single-Node Max Collected Metrics	1600000	5000000
Maximum Nodes in a cluster	2	8
Multi-Node Object Maximum	6000	17000
Multi-Node Max Collected Metrics	1400000	4000000
Maximum Objects in a Cluster	12000	136000
Maximum Metrics in a Cluster	2800000	32000000
Maximum Nodes for Continuous Availability per Fault Domain	1	4

Given the information above, which three physical design decisions meet the stated requirements? (Choose three.)

A. Enable Continuous Availability in VCF Operations.

B. Apply a vSphere Distributed Resource Scheduler (DRS) affinity rule to the VCF Operations cluster nodes.

C. Enable High Availability in VCF Operations.

D. Apply a vSphere Distributed Resource Scheduler (DRS) anti-affinity rule to the VCF Operations cluster nodes.

E. Deploy four medium VCF Operations nodes.

F. Deploy two small VCF Operations nodes.

Answer: A D E

Explanation:

A, D, and E meet the availability, sizing, and fault-domain requirements.

A is required because Continuous Availability (CA), rather than standard HA, provides fault-domain resilience. Broadcom states that VCF Operations CA stretches cluster nodes across two fault domains and can tolerate the loss of an entire fault domain without cluster downtime. Standard VCF Operations HA protects only against one node failure, so C does not satisfy the stated requirement for a single fault domain failure. CA requires equal node distribution across the two fault domains and a witness in a third location.

E is correct because 15,000 VMs plus all monitored infrastructure objects exceed the practical capacity of the Small configuration. Two Small nodes support only 12,000 maximum cluster objects before CA replication considerations. Four Medium nodes provide sufficient capacity and can be distributed as two analytics nodes per fault domain. Broadcom also notes that CA replication effectively reduces available object capacity compared with a non-CA deployment.

D ensures Operations analytics nodes are separated across physical hosts rather than sharing a common failure point. Broadcom specifically recommends DRS anti-affinity to keep Operations nodes on separate hosts.

Study Guide References/Topics: VCF Operations Sizing; Continuous Availability; Fault Domains; Analytics Node Placement; DRS Anti-Affinity; High Availability versus Continuous Availability.