

# VMware

## 3V0-24.25

Advanced VMware Cloud Foundation 9.0 vSphere  
Kubernetes Service

For More Information – Visit link below:

<https://www.examsempire.com/>

**Product Version**

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/3v0-24-25>

# Latest Version: 8.1

## Question: 1

An administrator must create a multi-zone vSphere Supervisor deployment in a VMware Cloud Foundation (VCF) environment. What is the primary purpose of this configuration?

- A. To create isolated security domains using NSX micro-segmentation.
- B. To enable cross-site vSAN stretched clusters for data replication between data centers.
- C. To provide high availability for the Supervisor Cluster and vSphere Kubernetes clusters.
- D. To simplify the management of network pools and IP address ranges.

**Answer: C**

Explanation:

A multi-zone Supervisor in VCF 9.0 is designed to deliver platform resiliency and high availability at the vSphere cluster (zone) failure-domain level. The VCF 9.0 documentation states that a multi-zone Supervisor “leverages three vSphere clusters” (each mapped to a vSphere Zone) and that these zones are used by both “workloads and Supervisor management components to deliver high availability,” exposing “each cluster as an independent, consumable availability zone,” resulting in a “resilient, HA-capable platform.”

This is reinforced in the vSphere Zones guidance: deploying the Supervisor on three vSphere Zones spreads the control plane VMs across three zones, providing “cluster-level high availability” that protects the Supervisor control plane against a single cluster-level failure (one control plane VM per management zone).

Because VKS (vSphere Kubernetes Service) runs on Supervisor, distributing Supervisor control plane and workload placement across zones improves overall availability of Supervisor services and Kubernetes consumption in that Supervisor instance.

## Question: 2

An administrator runs several critical workloads on vSphere Kubernetes Service (VKS). An audit identified an outdated container image with a known CVE that exposed internal APIs to unauthorized access. To mitigate this risk and enhance image security, the administrator enabled Harbor as a Supervisor Service.

Which two Harbor registry capabilities help the organization prevent a recurrence of this type of security incident? (Choose two.)

- A. Image signing
- B. Automatic image update
- C. Deploy both container and virtual machine images
- D. Automatic image validation

## E. Vulnerability scanning

**Answer: A,E**

### Explanation:

Harbor reduces the risk of running vulnerable or tampered images primarily through vulnerability scanning and image signing. Vulnerability scanning (E) detects known CVEs in image layers (OS packages and application dependencies, depending on the scanner configuration). This allows teams to identify—and gate the use of—images that contain high/critical vulnerabilities before those images are deployed to Kubernetes clusters. Enforcing scanning as part of the image promotion process helps prevent outdated images with known CVEs from being pulled into production. Image signing (A) provides integrity and provenance controls by enabling consumers to verify that an image was produced and approved by a trusted publisher and has not been altered. When combined with admission controls/policies (for example, only allowing signed images from specific projects), signing helps block unauthorized or unapproved images from being deployed, which is critical when the incident involves exposed internal APIs and supply-chain risk.

The other choices do not directly prevent recurrence: automatic image update (B) is not a core Harbor registry control, deploy both container and VM images (C) is a content capability rather than a security control, and automatic image validation (D) is not a standard Harbor registry capability distinct from signing/scanning.

## Question: 3

A company standardized on the following configurations:

- vSphere Kubernetes Service (VKS) upgrade is separate from vCenter upgrades.
- A private registry will be utilized.

How should an administrator adhere to these standards?

- A. Issue a PowerCLI command to point to the private registry.
- B. Issue a kubectl command pointing the service definition to the private registry.
- C. When uploading the service definition, choose Asynchronous Private.
- D. When uploading the service definition, choose Asynchronous Public.

**Answer: C**

### Explanation:

VCF 9.0 documentation explicitly indicates that vCenter upgrades and the Supervisor/cluster (Workload Management) upgrade are distinct, noting that “if you have only upgraded vCenter and not the cluster” then DevOps engineers have reduced permissions until the cluster is upgraded. This supports the stated standard that VKS/Workload Management lifecycle can be treated separately from vCenter. For the private registry requirement, VCF 9.0 provides an operational mechanism to authenticate and pull artifacts from private registries: “Registry secrets allow package and repository consumers to authenticate to and pull images from private registries,” implemented via a standard Kubernetes Secret of type `kubernetes.io/dockerconfigjson`.

Taken together, the standard implies (1) asynchronous handling (separate lifecycle from vCenter) and

(2)privatesourcing (images pulled from an internal registry with registry secrets). Therefore, selectingAsynchronous Privatebest matches both requirements in a single configuration choice, aligning with the documented separation of upgrades and the documented need to use authenticated access to private registries.

## Question: 4

An administrator is deploying vSphere Kubernetes Service (VKS) to support containerized workloads across multiple regions. Each region hosts a dedicated Workload Domain with Supervisor instances deployed on vSphere Distributed Switch (VDS) networking. The organization’s security policy requires that pod-to-pod and pod-to-service communications be fully observable and controllable at the Kubernetes layer, without introducing additional licensing or overlay complexity. When deploying a Supervisor, which CNI should the administrator select as the default supported option?

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

**Answer: A**

Explanation:

VCF 9.0 explicitly documents thatVKS supports two CNI options: Antrea and Calico, and that thesystem-defined default CNI is Antrea. This directly eliminates Flannel and Cilium as default supported options for VKS clusters on Supervisor in this context. VCF 9.0 also describes how a vSphere administrator can view or change this setting in the vSphere Client underSupervisor Management → Configure → Kubernetes Service → Default CNI, further reinforcing that Antrea is the baseline/default choice.

From a policy perspective in the question, the requirement is Kubernetes-layer observability and control of pod communications “without additional licensing or overlay complexity.” Antrea is presented in VCF 9.0 as the default CNI and is implemented usingOpen vSwitch, with networking and network policy capabilities provided at the Kubernetes layer for pods and services. Because it is the documented default (and supported) option for new VKS clusters, selectingAntreabest aligns with the “default supported option” requirement.

## Question: 5

What tool can be used to back up and restore workloads on clusters provisioned by vSphere Supervisor?

- A. Velero
- B. VMware Live Recovery
- C. Restic

#### D. Site Recovery Manager

**Answer: A**

##### Explanation:

VMware Cloud Foundation 9.0 documents a dedicated backup-and-restore approach for Workload Management where different components use different tools. For workloads running on vSphere Supervisor-based Kubernetes (both vSphere Pods and VKS cluster workloads), the documented solution is Velero, specifically the Velero Plugin for vSphere installed and configured on the Supervisor. The “Considerations for Backing Up and Restoring Workload Management” table explicitly lists: “Backup and restore vSphere Pods — Velero Plugin for vSphere” and “Backup stateless and stateful workloads on a VKS cluster and restore to a cluster provisioned by VKS — Velero Plugin for vSphere.” The same section also clarifies that Supervisor backups (via vCenter file-based backup) are for restoring the Supervisor control plane/state and VKS node VMs—not for restoring workloads themselves—so workloads must be backed up separately. Therefore, the correct tool for backing up and restoring workloads on clusters provisioned by vSphere Supervisor is Velero (using the Velero Plugin for vSphere).

**Thank You for Trying Our Product**  
**Special 16 USD Discount Coupon: NSZUBG3X**

**Email:** [support@examsempire.com](mailto:support@examsempire.com)

**Check our Customer Testimonials and ratings  
available on every product page.**

**Visit our website.**

**<https://examsempire.com/>**