Get Free VMware 3V0-24.25 Dumps PDF Questions

Why risk failure? Download updated VMware Advanced VMware Cloud Foundation 9.0 vSphere Kubernetes Service exam PDF questions today. Practice with real 3V0-24.25 dumps and verified answers designed to help you ace your certification quickly using PrepBolt 3V0-24.25 exam pdf questions and answers.

Thank you for Downloading 3V0-24.25 exam PDF Demo

https://prepbolt.com/3V0-24.25.html

Question 1

Question Type: MultipleChoice

Which two types of groups can be created to collect and manage objects in an Istio-based service mesh environment? (Choose two.)

Options:

A- Security
B- API
C- Node
D- Service
E- Cluster

Answer:

C, D

Explanation:

Comprehensive and Detailed 150 to 250 words of Explanation From Exact Extract of VMware Cloud Foundation (VCF) 9.0 + vSphere Supervisor + vSphere Kubernetes Service documents :

In VMware Cloud Foundation 9.0, the documented ''group'' construct for collecting and managing Kubernetes objects is implemented as generic groups with Kubernetes member types that can be used for policy-driven operations (for example, securing traffic between infrastructure workloads and Kubernetes workloads). The documentation explicitly states that you can ''create generic groups with Kubernetes member types in dynamic membership criteria'' and then use these groups in firewall rules to secure traffic involving Kubernetes clusters.

The same section provides a table of Kubernetes member types available for group membership criteria, and it explicitly lists Kubernetes Node (cluster scope) and Kubernetes Service (namespace scope) as supported member types. This maps directly to the answer choices Node and Service as the two valid ''types'' that can be used to build logical collections of Kubernetes objects for consistent management and policy enforcement.

The other answer options do not match the documented Kubernetes member types. ''Security'' and ''API'' are not Kubernetes member types in the group criteria model, and while ''Kubernetes Cluster'' is also a listed member type, the question asks for two, and Node and Service are the most direct object types for grouping runtime endpoints and service front-ends.

Question 2

Question Type: OrderList

An administrator had deployed a Supervisor cluster on vSphere in a multi-zone-enabled environment and now wants to create a zonal vSphere Namespace so that workloads can be scheduled across zones.

Drag and drop the six actions into the correct order from Configuration Option list on the left and place them Into the Configuration Sequence on the right.

(Choose six.)

Answer:

Create the vSphere NamespaceAssign the zonesSelect workload networkingAssign the zonal storage policyDefine resource quotas / limitsGrant RBAC / permissions

Question 3

Question Type: MultipleChoice

A VKS administrator is tasked to leverage day-2 controls to monitor, scale, and optimize Kubernetes clusters across multiple operating systems and workload characteristics.

What two steps should the administrator take? (Choose two.)

Options:

A- Configure namespace quotas to set resource limits for CPU, memory, and storage.
B- Disable Cluster Autoscaler to ensure resources in the pool are not depleted.
C- Deploy Prometheus and Grafana to collect and display scrapeable metrics on nodes, pods, and applications.
D- Set all VM Class limits to Compute Heavy to ensure worker nodes get all the resources needed.
E- Ensure all node pools use the same Machine Deployment configuration for different workload characteristics.

Answer:

A, C

Explanation:

VCF 9.0 describes a vSphere Namespace as the control point where administrators defineresource boundariesfor workloads, explicitly stating that vSphere administrators can create namespaces and ''configure them with specified amount ofmemory, CPU, and storage,'' and that you can ''set limits forCPU, memory, storage'' for a namespace. This directly supports stepAas a day-2 control to keep multi-tenant clusters governed and prevent resource contention across different teams and workload types.

For monitoring and optimization, VCF 9.0 explains that day-2 operations include visibility into utilization and operational metrics for VKS clusters, noting that application teams can use day-2 actions and gain insights intoCPU and memory utilizationand advanced metrics (including contention and availability) for VKS clusters. In addition, VCF 9.0 monitoring guidance for VKS clusters states thatTelegraf and Prometheusmust be installed and configured on each VKS cluster before metrics and object details are sent for monitoring, and that VCF Operations supports metrics collection for Kubernetes objects (namespaces, nodes, pods, containers) via Prometheus. Since the Prometheus stack commonly includes Grafana dashboards for visualization, deployingPrometheus + Grafanamatches the required monitoring/optimization outcome inC.

Question 4

Question Type: MultipleChoice

How would an administrator obtain the kubectl config file in the VMware vSphere Kubernetes Service (VKS)?

Options:

A- Download from VCF Operations.
B- Download from the vSphere UI.
C- Download from the Supervisor Services webpage.
D- Use the command kubectl vsphere login.

Answer:

D

Explanation:

VCF 9.0 documents that you obtain access to the Supervisor (and the contexts you're entitled to,

which commonly correspond to namespaces and workload clusters) by authenticating withkubectl vsphere login. The procedure ''Get and Use the Supervisor Context'' explains that after you log in, a Kubernetes configuration context is generated and you can view it in$HOME/.kube/config, which is ''commonly called the kubeconfig file.''

The documentation is explicit that running kubectl vsphere login --server=<KUBERNETES-CONTROL-PLANE-IP> --vsphere-username <VCENTER-SSO-USER> ''creates a configuration file with the JSON Web Token (JWT) to authenticate to the Kubernetes API,'' and then lists how to view and switch contexts using kubectl config get-contexts and kubectl config use-context.

This matches optionDprecisely: you don't ''download'' kubeconfig from VCF Operations or a Supervisor Services page; instead, you generate/update your kubeconfig locally by logging in withkubectl's vSphere plugin, which writes the required cluster/user/context details into .kube/config.

Question 5

Question Type: MultipleChoice

What Kubernetes component is responsible for workload creation?

Options:

A- etcd
B- API Server
C- Scheduler
D- Kubelet

Answer:

D

Explanation:

In Kubernetes, the component that actuallycreates and runs workloads on a nodeis thekubelet. The kubelet is the node agent that ensures the containers described by PodSpecs are running on that node. VCF 9.0 maps this concept directly into vSphere Supervisor by describingSphereletas ''a kubelet that is ported natively to ESXi and allows the ESXi host to become part of the Kubernetes cluster,'' showing that kubelet functionality is responsible for running workloads on worker nodes (ESXi hosts in the Supervisor case).

The other options have different roles:etcdis the control plane data store,API Serveris the front-end for Kubernetes API operations, and theSchedulerdecides placement (which node should run a pod). VCF 9.0 even calls out that ''the Kubernetes scheduler... cannot place pods intelligently'' without visibility

into vCenter inventory---reinforcing that scheduling is about placement decisions, not the act of creating/running the workload on the node.

So, while the scheduler selects where a pod should run, thekubeletis the component responsible for actually instantiating and maintaining the workload on the target node.

Question 6

Question Type: MultipleChoice

What are three resource limitations defined on a vSphere Namespace? (Choose three.)

Options:

A- Containers
B- Services
C- Memory
D- CPU
E- Storage

Answer:

C, D, E

Explanation:

In VCF 9.0 Workload Management, avSphere Namespaceis the construct that ''sets the resource boundaries'' for workloads running on a Supervisor, includingCPU, memory, and storage. The documentation explicitly states that a vSphere Namespace ''sets the resource boundaries forCPU, memory, storage, and also the number of Kubernetes objects that can run within the namespace.'' In the operational procedure ''Set Resource Limits to a vSphere Namespace,'' VMware further lists the configurable limits as:CPU(''set a limit to the CPU consumption''),Memory(''set a limit to the memory consumption''), andStorage(''set a limit on the storage consumption... per storage policy that is used'').

By contrast,Containersare not a namespace ''resource limit'' category; VMware documents ''Container Defaults'' separately (defaults for container CPU/memory requests and limits) rather than a top-level resource limit type. Similarly,Servicesare governed under ''Object Limits'' (how many Kubernetes objects like Services can exist), which is distinct from resource limits. Therefore, the three resource limitations defined on a vSphere Namespace areCPU, Memory, and Storage.

Question 7

Question Type: MultipleChoice

An administrator is upgrading to VKS 3.4 and encounters the following error during cluster creation using workload, yami:

How should the administrator resolve this issue to successfully complete the upgrade"?

Options:

A- Verify workload cluster versions to ensure compatibility
B- Remove the deprecated variables and apply the new workload, yaml.
C- Rename the vSphere storage policy and apply the new workload.yami.
D- Restart the Kubernetesservices and restart the upgrade

Answer:

B

Explanation:

The error shows an admission webhook denial wherevariable validation failedand multiple entries under spec.topology.variables3...] are reported as''variable is not defined''. That message indicates the manifest is supplying variables that arenot part of the current Cluster API / topology schemaenforced by the Supervisor during cluster creation. In VKS, cluster provisioning isdeclarative: you invoke the VKS API withkubectl + a YAML file, and ''after the cluster is created, you update the YAML to update the cluster.'' When the API/schema changes between releases, older manifests can contain fields/variables that are no longer recognized, and the admission webhook blocks them to prevent creating an invalid cluster spec.

This aligns with VMware's broader direction that the olderTanzuKubernetesCluster (TKC) API was deprecatedand customers are encouraged to useCluster APIfor bootstrap/config/lifecycle management. In practice, to complete the upgrade/creation successfully, you must update the cluster manifest to match the supported schema:remove the deprecated/unknown topology variablesshown in the error (for example, the undefined storage-policy and trust variables) and re-apply the correctedworkload.yaml.

Question 8

Question Type: Hotspot

The vSphere Admin creates a vSphere Namespace harbor-project and grants the DevOps Engineer edit permissions. The DevOps Engineer deploys the Harbor service in thisnamespace with the command: kubectl apply -f harbor-svc.yml

The Harbor service deploys successfully, but the database data harbor database 0 pods does not come into a Running state.

Click the two locations where the administrator should verify the Access Modes forthis pod's PerslstentVolumeClaim (PVC). (Choose two.)

Answer:

See the Answer in the Premium Version!

Thank You for trying 3V0-24.25 PDF Demo

To try our 3V0-24.25 practice exam software visit link below

https://prepbolt.com/3V0-24.25.html

Start Your 3V0-24.25 Preparation

Use Coupon “SAVE50” for extra 50% discount on the purchase of Practice Test Software. Test your 3V0-24.25 preparation with actual exam questions.