



VMware Specialist - vSphere with Tanzu 5V0-23.20 Practice Questions

Shared by Booker on 17-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

An administrator is tasked with increasing the amount of CPU and memory in an existing Tanzu Kubernetes cluster.

Which change must the administrator complete to ensure the cluster scales successfully when updating the YAML definition?

Options:

- A- Manually update the CPU and memory of the nodes
- B- Update the Virtual Machine Class Type
- C- Increase the number of worker nodes
- D- Increase the number of control plane nodes

Answer:

B

Explanation:

Virtual Machine Class Types for Tanzu Kubernetes Clusters

A virtual machine class defines the resource sizing for Tanzu Kubernetes cluster VMs: * CPU * Memory *

Storage

Virtual machine class types range from extra small (xsmall) to extra large (xlarge). Class types are categorized as guaranteed or best effort:

* Guaranteed: Reserve all CPU and memory allocations. * Best effort: Allocate the same CPU and memory but do not reserve the resources.

The class type guaranteed-small allocates 2 CPU, 4 GB of memory, and 16 GB of storage and reserves CPU and memory allocations. Custom virtual machine class types cannot be defined.

Question 2

Question Type: MultipleChoice

The creation of which object by an administrator in the vSphere client automatically results in the creation of a new segment within NSX -T?

Options:

- A- Service
- B- Pod
- C- Network policy
- D- Namespace

Answer:

D

Explanation:

NSX Container Plugin (NCP) runs as a pod on the control plane VMs. It listens for requests for network objects to the API server and interfaces with the NSX Manager to create, update, or delete those objects:

- * A request to create a namespace results in a new NSX segment.
- * A request to deploy a pod results in a segment port request and IP assignment.
- * A request to create a service results in a new virtual server.
- * A request to create a network policy results in a new distributed firewall rule.

Question 3

Question Type: MultipleChoice

Which requirement is needed to deploy Tanzu Kubernetes clusters but is unnecessary when deploying vSphere Pods?

Options:

- A- Harbor Image Registry
- B- vSphere Single Sign-On
- C- vSphere Content Library

D- Namespaces

Answer:

A

Question 4

Question Type: MultipleChoice

Which step is necessary in order to allocate storage to a newly created namespace?

Options:

- A- Assign a VM Storage Policy
- B- Assign permissions to users
- C- Define resource limits
- D- Add a content library

Answer:

A

Question 5

Question Type: MultipleChoice

The Gold storage policy has been assigned to the Web namespace, and the DevOps engineer wants to place a persistent volume for the Web application in the Gold storage class.

How should this goal be accomplished?

Options:

- A- Indicate the Gold storage class in the persistent volume claim specification
- B- Assign the persistent volume to the Gold storage class
- C- Indicate the Gold storage class in the persistent volume specification
- D- Configure tag-based placement for the persistent volume

Answer:

A

Explanation:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: my-pvc
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: silver-storage-policy
resources:
  requests:
    storage: 3Gi
```

* VM storage policies are translated into Kubernetes storage classes.

* Developers can access all assigned VM storage policies in the form of storage classes. *

Developers cannot manage storage classes.

Question 6

Question Type: MultipleChoice

Which external load balancer is supported in vSphere 7 U1 using the vSphere networking stack?

Options:

- A- Nginx
- B- Seesaw
- C- Loadmaster
- D- HAProxy

Answer:

D

Explanation:

When using vSphere with Tanzu with vDS networking, HAProxy provides load balancing for developers accessing the Tanzu Kubernetes control plane, and for Kubernetes Services of Type Load Balancer. Review the possible topologies that you can implement for the HAProxy load balancer.

<https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-1F885AAE-92FF-41E6-BF04-0F0FD4173BD9.html>

The HAProxy appliance is an open-source solution developed by HAProxy Technologies and chosen by VMware as the first supported open-source load balancer for use with vSphere with Tanzu. With the HAProxy, external network traffic is routed to Kubernetes pods running in the vSphere with Tanzu environment.



Question 7

Question Type: MultipleChoice

To which network are HA Proxy virtual server IP addresses issued when using the vSphere networking stack default configuration?

Options:

- A- vMotion
- B- Overlay
- C- Primary workload
- D- Management

Answer:

C

Explanation:

The HAProxy virtual IP range where external services and DevOps users connect. In this configuration, HAProxy is deployed with two virtual NICs (Default configuration), one connected to the management network, and a second one connected to the Primary Workload Network. You must plan for allocating Virtual IPs on a separate subnet from the Primary Workload Network.

<https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-1F885AAE-92FF-41E6-BF04-0F0FD4173BD9.html>

Question 8

Question Type: MultipleChoice

A company needs its administrators to be able to manage access, apply security policies, and inspect Kubernetes clusters for security and configuration risks across multiple platforms.

Which solution will meet these requirements?

Options:

- A- Kubernetes Security Contexts
- B- Kubernetes Role Based Access Control (RBAC)
- C- Tanzu Kubernetes Grid Service
- D- Tanzu Mission Control

Answer:

B

Question 9

Question Type: MultipleChoice

The application development team is pushing a Kubernetes application into production. It consists of an application server and a database. The team wants to ensure that only the production application server can access the production database.

Can the development team meet this requirement using Kubernetes Network Policy?

Options:

- A- Yes, by using kubectl to create a Network Policy that only allows pods on the same network segment to talk to each other.
- B- Yes, by logging in to NSX Manager and creating a firewall rule to only allow the production application server pod to talk to the database
- C- Yes, by using kubectl to create a policy that disables pod to pod communication in the Namespace
- D- No, Kubernetes Network Policy does not support this action.

Answer:

A

Explanation:

If you want to control traffic flow at the IP address or port level (OSI layer 3 or 4), then you might consider using Kubernetes NetworkPolicies for particular applications in your cluster. NetworkPolicies are an application-centric construct which allow you to specify how a pod is allowed to communicate with various network 'entities' (we use the word 'entity' here to avoid overloading the more common terms such as 'endpoints' and 'services', which have specific Kubernetes connotations) over the network. NetworkPolicies apply to a connection with a pod on one or both ends, and are not relevant to other connections.

Question 10

Question Type: MultipleChoice

Which two actions are orchestrated by the NSX Container Plugin (NCP)? (Choose two.)

Options:

- A- Creation of an embedded registry
- B- Creation of an NSX Edge
- C- Creation of a portgroup
- D- Creation of a service
- E- Creation of a network policy

Answer:

A, D

Question 11

Question Type: MultipleChoice

What can be used to define whether a vSphere pod can be scaled?

Options:

- A- Deployment
- B- Namespaces
- C- Persistent Volume
- D- Network Policies

Answer:

A

To Get Premium Files for 5V0-23.20 Visit

<https://www.p2pexams.com/products/5v0-23.20>

For More Free Questions Visit

<https://www.p2pexams.com/vmware/pdf/5v0-23.20>

20%
DISCOUNT

P2P
exams