Exam Code: 5V0-23.20
Exam Name: VMware vSphere with Tanzu Specialist

QUESTION 1
Where is a storage policy applied to enable Persistent Volumes?

A. Namespace
B. Datastore
C. Virtual Machine
D. Cluster

Correct Answer: A
Section:
Explanation:
The vSphere administrator defines and assigns VM storage policies to a namespace:
• 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.
Developers can list the available storage classes in their namespace by running the kubectl describe ns <namespace-name> command.

QUESTION 2
The virtualization team supports many development teams on a Supervisor cluster. For a specific development team, they would like to limit persistent volumes that can be created on Tanzu Kubernetes clusters to only an NFS based storage array.
Which action should be taken to accomplish this goal?

A. Use kubectl to create a storage class in the Supervisor cluster.
B. Set a resource quota limiting the number of PVCs for that development team.
C. Add a storage policy to that development team's Supervisor Namespace containing only the NFS datastore
D. Disconnect non-NFS datastores from the ESXi hosts that make up the Supervisor cluster.

Correct Answer: C
Section:
Explanation:
The storage policy would become a storage class in Kubernetes.

QUESTION 3
Which command displays the storage limits that have been set together with the amount of resources consumed?

A. kubect1 get resourcequotas
B. kubect1 config get-resourcequotas limits
C. kubect1 list resourcequotas
D. kubect1 describe resourcequotas

Correct Answer: A
Section:
Explanation:

Create the ResourceQuota:

kubectl apply -f https://k8s.io/examples/admin/resource/quota-mem-cpu.yaml --namespace=quotamem- cpu-example

View detailed information about the ResourceQuota:

kubectl get resourcequota mem-cpu-demo --namespace=quota-mem-cpu-example --output=yaml

QUESTION 4

Which is a valid version change for a Tanzu Kubernetes cluster running Kubernetes version 1.16.7?

A. Upgrade one major version (e.g.. 2.0.1)

B. Upgrade two minor versions (e.g., 1.18.0)

C. Downgrade one patch version (e.g.. 1.16.5)

D. Upgrade one minor version (e.g.. 1.17.0)

Correct Answer: D

Section:

Explanation:

Be aware of the following constraints when upgrading your cluster.

You can upgrade a cluster up to one minor version of Kubernetes from its current version. If necessary, you can perform subsequent upgrades to move the version forward.

Upgrading your version of Kubernetes is a one-way operation. You cannot subsequently downgrade the Kubernetes version, or undo an upgrade.

QUESTION 5

What is the minimum number of portgroups needed, in addition to the management portgroup. to provide connectivity for external services on a Supervisor Cluster?

A. 1

B. 4

C. 3

D. 2

Correct Answer: A

Section:

QUESTION 6

A developer is trying to deploy a Kubernetes Application into a namespace within a Supervisor Cluster The deployment must utilize the latest assets that have been pushed into the Registry Service.

What should the developer add to the YAML file to ensure that the deployment is successful?

A. image: /<namespace>/<image name>:latest

B. template: <image registry url>/<namespace name>/<image name> : latest

C. image: <image registry url>/<namespace name>/<image name>:latest

D. template: /<namespace name>/<image name>:latest

Correct Answer: C

Section:

Explanation:

Create an example pod spec with the details about the private registry.

apiVersion: v1

kind: Pod

metadata:

name: <workload-name>

namespace: <kubernetes-namespace>

spec:

containers:

- name: private-reg-container

image: <Registry-IP-Address>/<vsphere-namespace>/<image-name>:<version>

imagePullSecrets:

- name: <registry-secret-name>

Replace <workload-name> with the name of the pod workload.

Replace <kubernetes-namespace> with the Kubernetes namespace in the cluster where the pod will be created. This must be the same Kubernetes namespace where the Registry Service image pull secret is stored in the Tanzu Kubernetes cluster (such as the default namespace).

Replace <Registry-IP-Address> with the IP address for the embedded Harbor Registry instance running on the Supervisor Cluster.

Replace <vsphere-namespace> with the vSphere Namespace where the target Tanzu Kubernetes is provisioned.

Replace <image-name> with an image name of your choice.

Replace <version> with an appropriate version of the image, such as "latest".

Replace <registry-secret-name> with the name of the Registry Service image pull secret that you created previously.

QUESTION 7
Which two container network interfaces (CNIs) are supported with Tanzu Kubernetes clusters created by the Tanzu Kubernetes Grid Service? (Choose two)

A. NSX-T
B. Weave Net
C. Flannel
D. Antrea
E. Calico

Correct Answer: D, E
Section:
Explanation:
https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-A7756D67-0B95-447D-A645-E2A384BF8135.html
A Tanzu Kubernetes cluster provisioned by the Tanzu Kubernetes Grid Service supports two CNI options: Antrea (default) and Calico. Both are open-source software that provide networking for cluster pods, services, and ingress.
Tanzu Kubernetes clusters provisioned by the Tanzu Kubernetes Grid Service support the following Container Network Interface (CNI) options:
Antrea
Calico
Explanation

Tanzu Kubernetes Grid Service CNI

Tanzu Kubernetes Grid Service supports Antrea and Calico as container network interfaces (CNI).

The default CNI in vSphere 7 Update 1 is Antrea.

Antrea is a VMware-supported, open source, Kubernetes-native project that implements the container network interface (CNI) and Kubernetes network policy, providing network connectivity and security for pod workloads. Antrea extends the benefit of programmable networks from Open vSwitch (OVS) to Kubernetes.

For more information about Antrea, see https://antrea.io/

QUESTION 8

Where are the virtual machine images stored that are used to deploy Tanzu Kubernetes clusters?

A. Content Library

B. Supervisor Cluster

C. Harbor Image Registry

D. Namespace

Correct Answer: A
**Section:**
Explanation:

The vSphere administrator configures a Subscribed Content Library on the Supervisor Cluster. The virtual machine image that is used for the Tanzu Kubernetes cluster nodes is pulled from this library.

A Subscribed Content Library originates from a Published Content Library. After the subscription is created, the system synchronizes it with the published library. To create the Tanzu Kubernetes cluster nodes, VMware publishes a Photon OS OVA library to which you subscribe. After the subscriber is synchronized with the publisher, you associate the content library with the Supervisor Cluster.

QUESTION 9
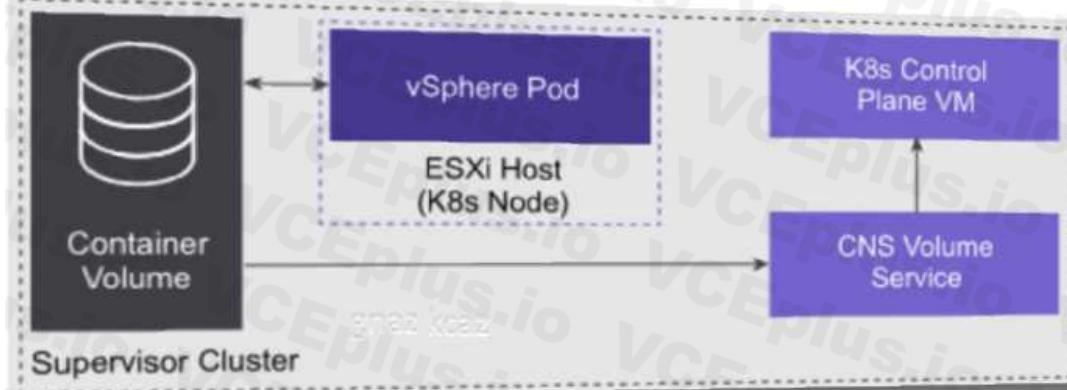
Which capability do persistent volumes provide to containerized applications?

A. Automated disk archival

B. Support for in-memory databases

C. Support for ephemeral workloads

D. Retention of application state and data

Correct Answer: D

Section:

Explanation:

Certain Kubernetes workloads require persistent storage to store data permanently. To provision persistent storage for Kubernetes workloads, vSphere with Tanzu integrates with Cloud Native Storage (CNS), a vCenter Server component that manages persistent volumes.

Persistent storage is used by vSphere Pods, Tanzu Kubernetes clusters, and VMs. The following example illustrates how persistent storage is used by a vSphere Pod. vSphere Pods use different types of storage depending on the objects that are stored. The types of storage are ephemeral virtual machine disks (VMDKs), persistent volume VMDKs, and containers image VMDKs:

• Storage policies for container image and ephemeral disks are defined at the cluster level.

• Storage policies for persistent volumes are defined at the namespace level.

• Networking for vSphere Pods uses the topology provided by NSX.

QUESTION 10
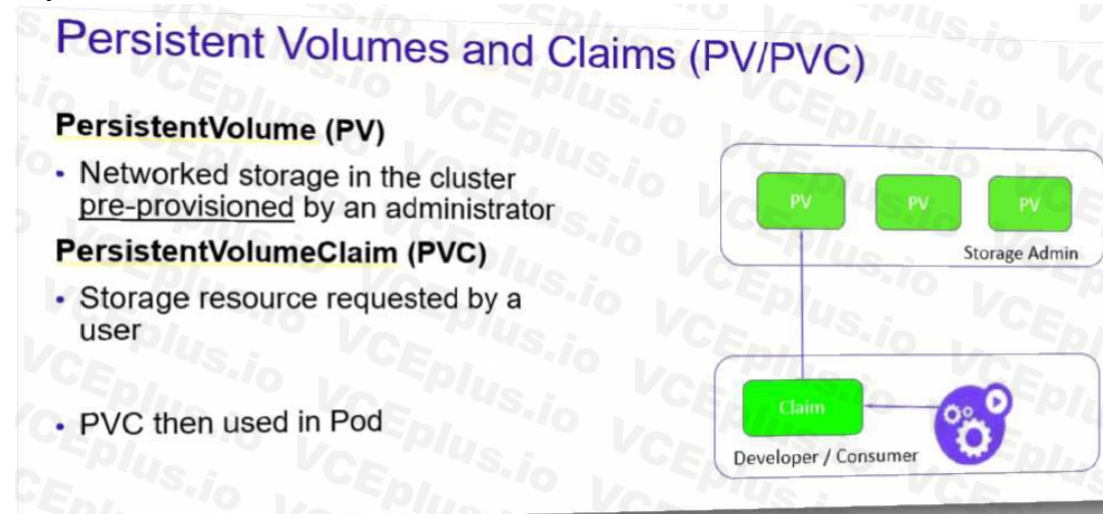
What is the proper way to delete a Persistent Volume Claim?

A. By using the kubectl delete persistentvolumeclaim command

B. By using the kubectl remove pvc command

C. Through the SPBM policy engine using the vSphere Client

D. By unmounting the volume from the VM and deleting it from the vSphere datastore

Correct Answer: A

Section:

Explanation:

Also, kubectl delete pvc, which is much shorter.

DevOps engineers create persistent volume claims to request persistent storage resources. The request provisions a persistent volume object and a matching virtual disk. In the vSphere Client, the persistent volume claim manifests as an FCD virtual disk that can be monitored by vSphere administrators.

The claim is bound to the persistent volume. The workloads can use the claim to mount the persistent volumes and access storage.

When the DevOps engineers delete the claim, the corresponding persistent volume object and the provisioned virtual disk are also deleted.

QUESTION 11

Which command will show the Tanzu Kubernetes cluster versions available in the vSphere content library?

A. kubect1 get rc,services

B. kubect1 get contentlibrary

C. kubect1 get tanzukubernetesreleases

D. kubect1 get tanzuimages

Correct Answer: C
Section:
Explanation:
VMware Tanzu distributes Kubernetes software versions as Tanzu Kubernetes releases. To consume these releases, you configure a vSphere Content Library and synchronize the available releases. You can do so using a subscription-based model, or on-demand. If you want to provision Tanzu Kubernetes in an internet restricted environment, you can create a local library and manually import the releases.

```
ubuntu@cli-vm:~$
ubuntu@cli-vm:~$ kubectl get tanzukubernetesreleases
NAME                               VERSION
v1.16.12---vmware.1-tkg.1.da7afe7  1.16.12+vmware.1-tkg.1.da7afe7
.1-tkg.1.ada4837]
v1.16.14---vmware.1-tkg.1.ada4837  1.16.14+vmware.1-tkg.1.ada4837
v1.16.8---vmware.1-tkg.3.60d2ffd   1.16.8+vmware.1-tkg.3.60d2ffd
.1-tkg.1.ada4837]
v1.17.11---vmware.1-tkg.1.15f1e18  1.17.11+vmware.1-tkg.1.15f1e18
.1-tkg.1.d44d45a]
v1.17.11---vmware.1-tkg.2.ad3d374  1.17.11+vmware.1-tkg.2.ad3d374
.1-tkg.1.d44d45a]
v1.17.13---vmware.1-tkg.2.2c133ed  1.17.13+vmware.1-tkg.2.2c133ed
.1-tkg.1.d44d45a]
v1.17.17---vmware.1-tkg.1.d44d45a  1.17.17+vmware.1-tkg.1.d44d45a
v1.17.7---vmware.1-tkg.1.154236c   1.17.7+vmware.1-tkg.1.154236c
.1-tkg.1.d44d45a]
v1.17.8---vmware.1-tkg.1.5417466   1.17.8+vmware.1-tkg.1.5417466
.1-tkg.1.d44d45a]
v1.18.10---vmware.1-tkg.1.3a6cd48  1.18.10+vmware.1-tkg.1.3a6cd48
1-tkg.2.ebf6117]
```

QUESTION 12
Which object helps maintain copies of a vSphere pod?

A. ReplicaSets
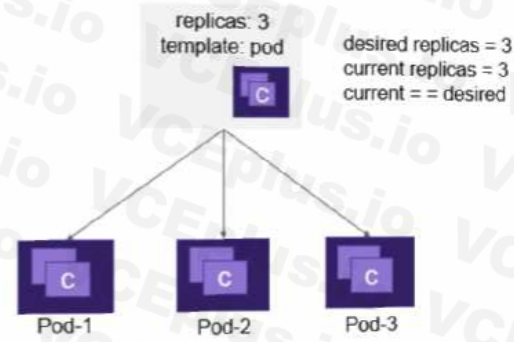B. Network Policies
C. Namespaces
D. Persistent Volume

Correct Answer: A
Section:
Explanation:

Replica Set

- A **Replica Set (rs)** is the next-generation **Replication Controller**.

- Replica Sets support both equality- and set-based selectors, whereas Replication Controllers only support equality-based Selectors.

replicas: 3
template: pod

desired replicas = 3
current replicas = 3
current == desired

Pod-1   Pod-2   Pod-3

A ReplicaSet declares how the functionality of a pod is made scalable and resilient through redundancy. The ReplicaSet ensures that a specified number of pods is kept running. Example: Deploy a ReplicaSet.
• The ReplicaSet name is nginx-replica-demo.
• Two replicas are expected to be running.
• The ReplicaSet applies to pods with label nginx.
For more information about Kubernetes replica sets, see
https://kubernetes.io/docs/concepts/workloads/controllers/replicaset/

QUESTION 13
On which network are TKG clusters deployed in vSphere with Tanzu when using the vSphere networking stack?

A. Workload

B. Backend

C. Edge

D. Frontend

Correct Answer: A
Section:
Explanation:
The Workload Network, such as TKGS-VLAN1000, is where the Tanzu Kubernetes clusters run.
A workload network is a network construct that is used by supervisor control plane VMs and vSphere namespaces:
• The workload network is supported by a vSphere Distributed Switch port group.
• An IP range is defined to allocate an IP address for VMs attached to the workload network.
• A primary workload network must be selected.
• The supervisor control plane VMs attach to the primary workload networks port group.
A workload network can be used by multiple namespaces. A namespace can be assigned only one workload network.

QUESTION 14
What is the correct process to store images in a project on the Registry Service?

A. Use the kubect1 push command

B. Use the docker push command

C. Use the vSphere Client to upload the image the content library

D. Use the vSphere Client to upload the image to the Registry Service

Correct Answer: B
Section:
Explanation:
https://docs.docker.com/docker-hub/repos/

Command	Description
docker image history	Show the history of an image
docker image import	Import the contents from a tarball to create a filesystem image
docker image inspect	Display detailed information on one or more images
docker image ls	List images
docker image prune	Remove unused images
docker image pull	Pull an image or a repository from a registry
docker image push	Push an image or a repository to a registry
docker image rm	Remove one or more images
docker image save	Save one or more images to a tar archive (streamed to STDOUT by default)
docker image tag	Create a tag TARGET_IMAGE that refers to SOURCE_IMAGE

• Registry Service: Developers can store and manage Docker and OCI images using Harbor. Harbor is an open-source container image registry that secures images with role-based access control.
Procedure

Login to Harbor Registry with the vSphere Docker Credential Helper.

docker-credential-vsphere login <container-registry-IP> --user username@domain.com

Note:While providing --user username is acceptable for login, you should use the UserPrincipalName (UPN) syntax (--user username@domain.com) to login and use docker push commands.

Tag the image that you want to push to the project in Harbor Registry with same name as the

namespace, where you want to use it:

docker tag <image-name>[:TAG] <container-registry-IP>/<project-name>/<image-name>[:TAG]

For example:

docker tag hello-world:latest 10.179.145.77/tkgs-cluster-ns/hello-world:latest

docker images

REPOSITORY TAG IMAGE ID CREATED SIZE

10.179.145.77/tkgs-cluster-ns/hello-world latest bf756fb1ae65 10 months ago

13.3kB

hello-world latest bf756fb1ae65 10 months ago 13.3kB

To push an image to a project in Harbor, run the following command:Syntax:

docker push <container-registry-IP>/<namespace-name>/<image_name>

For example:

docker push 10.179.145.77/tkgs-cluster-ns/hello-world:latest

Expected result.

The push refers to repository [10.179.145.77/tkgs-cluster-ns/hello-world]

9c27e219663c: Pushed

latest: digest: sha256:90659bf80b44ce6be8234e6ff90a1ac34acbeb826903b02cfa0da11c82cbc042

size: 525

QUESTION 15
Which functionality does the Cloud Native Storage (CNS) component take advantage of to support the creation of container volumes?

A. First Class Disk
B. VMware Disk Encryption
C. Virtual Disk
D. Storage Based Policy Management

Correct Answer: A
Section:
Explanation:
The Cloud Native Storage server resides in vCenter Server:
• Provisions and manages life cycle operations for container volumes
• Creates First Class Disks (FCDs) to support the container volumes
• First Class Disks exist as .vmdk and -flat.vmdk files on a vSphere datastore • Integrates with storage policy based management (SPBM) for the placement of disks A First Class Disk (FCD) is also called an improved virtual disk. It is a named virtual disk that is unassociated with a VM. These disks reside on a VMFS, NFS, or vSAN datastore and support container volumes.
Storage policy based management (SPBM) is a vCenter Server service that supports provisioning of persistent volumes according to specified storage requirements. After provisioning, the service monitors compliance of the volume with the required policy characteristics.

QUESTION 16
Which step in vSphere with Tanzu enablement using the vSphere Distributed Switch process is done prior to using the Workload Management Enablement Wizard?

A. Deploy the load balancer
B. Choose the Kubernetes content library that should be used in the Supervisor Cluster
C. Define the Primary workload network P range
D. Define the Management network interfaces for the Supervisor Cluster

Correct Answer: D
Section:
Explanation:
As a vSphere administrator, you can enable the Workload Management platform on a vSphere cluster by configuring the vSphere networking stack to provide connectivity to workloads. A Supervisor Cluster that is configured

with vSphere networking supports the deployment of Tanzu Kubernetes clusters created by using the Tanzu Kubernetes Grid Service. It does not support running vSphere Pod or using the embedded Harbor Registry.
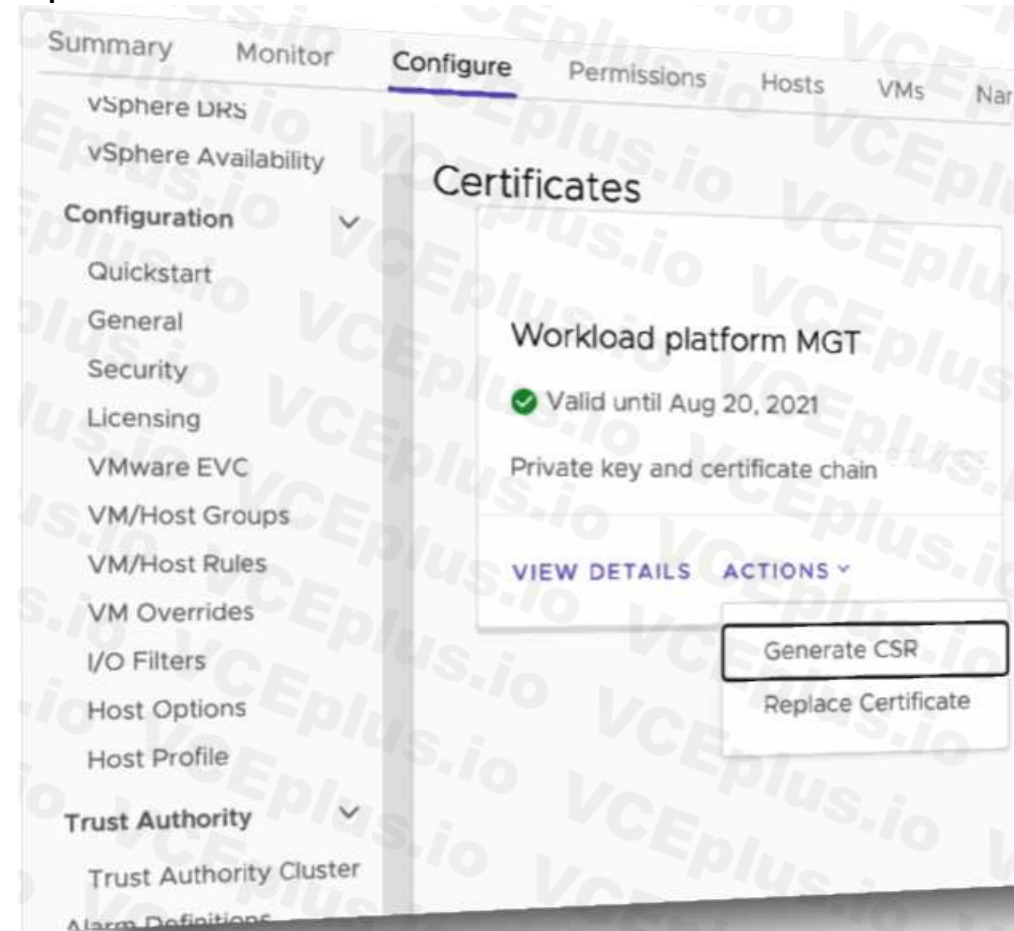
QUESTION 17
How can a vSphere administrator replace the Supervisor Cluster API endpoint certificate?

A. Use the certificate-manager CLI utility to replace the Supervisor Cluster API endpoint certificate.

B. Use the vSphere Client to replace the Workload platform MTG certificate.

C. Use the vSphere Client to replace the NSX Load Balancer certificate.

D. Use kubectl to replace the Supervisor Cluster API endpoint certificate.

Correct Answer: B
Section:
Explanation:

As a vSphere administrator, you can replace the certificate for the virtual IP address (VIP) to securely connect to the Supervisor Cluster API endpoint with a certificate signed by a CA that your hosts already trust. The certificate authenticates the Kubernetes control plane to DevOps engineers, both during login and subsequent interactions with the Supervisor Cluster.
Prerequisites
Verify that you have access to a CA that can sign CSRs. For DevOps engineers, the CA must be installed on their system as a trusted root.
Procedure
In the vSphere Client, navigate to the Supervisor Cluster.
Click Configure then under Namespaces select Certificates.
In the Workload platform MTG pane, select Actions > Generate CSR.
Provide the details for the certificate.
Once the CSR is generated, click Copy.

Sign the certificate with a CA.
From the Workload platform MTG pane, select Actions > Replace Certificate.
Upload the signed certificate file and click Replace Certificate.
Validate the certificate on the IP address of the Kubernetes control plane.

QUESTION 18
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?

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

Correct Answer: A
Section:
Explanation:

• 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 19
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?

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

Correct Answer: B
Section:
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 20
An organization is preparing to deploy vSphere with Tanzu and will be using the vSphere Networking stack.
How should the administrator allocate management network IP addresses for the Kubernetes Control Plane within the Supervisor Cluster?

A. Five IP addresses are required, one for each of the Control Plane VMs. one for the floating IP address of the Control Plane VM, and one spare for performing rolling cluster upgrades
B. Four IP addresses are required, one for each of the Control Plane VMs and one spare for performing rolling cluster upgrades
C. Three IP addresses are required, one for each of the Control Plane VMs
D. Six IP addresses are required, one for each of the Control Plane VMs, one for the floating IP address of the Control Plane VM. one for performing rolling cluster upgrades and one for the image Registry VM.

Correct Answer: A
Section:
Explanation:
Static IPs for Kubernetes control plane VMs
Block of 5A block of 5 consecutive static IP addresses to be assigned to the Kubernetes control plane VMs in the Supervisor Cluster.

QUESTION 21
Which command should be used by a developer to log in to the vSphere with Tanzu Supervisor Cluster?

A. vmwarectl login --server-<KUBERNETES-CONTROL-PLANE-IP-ADDRESS> --vsphere-username
    <VCENTER-SSO-USER>
B. kubectl vsphere login --server=<KUBERNETES-CONTROL-PLANE-IP-ADDRESS> --vsphereusername
    <vcENTER-SSO-USER>
C. vmwarectl vsphere login --server—<KUBERNETES-CONTROL-PLANE-IP-ADDRESS> --vsphereusername
    <VCENTER-SSO-USER>
D. kubectl login --server=<KUBERNETES-CONTROL-PLANE-lP-ADDRESS> --vsphere-username
    <VCENTER-SSO-USER>

Correct Answer: B
Section:
Explanation:
Authenticating Using kubectl
In vSphere with Tanzu, authentication is performed using vCenter Single Sign-On. You run the command kubectl vsphere login to authenticate a user through vCenter Single Sign-On to Kubernetes clusters. kubectl vsphere login --server <kubernetes control plane> -u <username> The --insecure-skip-tls-verify flag is required if the certificate presented by the vSphere with Tanzu control plane is not trusted by the client machine. The control plane is signed by the vCenter Server VMware CA by default. You can replace the control plane certificate with a trusted certificate if needed. Alternatively, install the vCenter Server VMware CA root certificate into your client machine to remove the need for the --insecure-skip-tls-verify flag.

QUESTION 22
Which external load balancer is supported in vSphere 7 U1 using the vSphere networking stack?

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

Correct Answer: D
Section:
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 23
A user needs to identify the namespaces that may be accessed.
Which command will provide the desired output?

A. kubect1 get storageclasses
B. kubect1 config use-context
C. kubect1 config get-contexts
D. kubect1 get contexts

Correct Answer: C
Section:
Explanation:
A user can have permissions on multiple namespaces. The kubectl commands are typically actioned against the current active namespace.
View the list of available namespaces: kubectl config get-contexts Change the current active namespace: kubectl config use-context <namespace>

QUESTION 24

What is required to enable Workload Management?

A. vSphere Distributed Switch
B. Windows Network Load Balancer
C. Github repository
D. NSX-V

Correct Answer: A
Section:
Explanation:
https://docs.vmware.com/en/VMware-vSphere/7.0/vsphere-esxi-vcenter-server-702-vsphere-withtanzu-guide.pdfConfiguring Workload Networks You configure one or more workload networks and their respective IP address ranges.
Each workload network is assigned a vSphere Distributed Switch port group and uses a defined IP range to allocate IP addresses to workloads (VMware Tanzu Kubernetes clusters).

QUESTION 25
Which role should the vSphere administrator apply for the developer?

A. Assign the developer user with the "can edit" role at the vSphere Namespace object.
B. Assign the developer user with the 'VSphere Kubernetes Manager" role at the vSphere Namespace object.
C. Assign the developer user with the 'VSphere Kubernetes Manager" role at the cluster object.
D. Assign the developer user with the "can edit" role at the cluster object.

Correct Answer: A
Section:
Explanation:
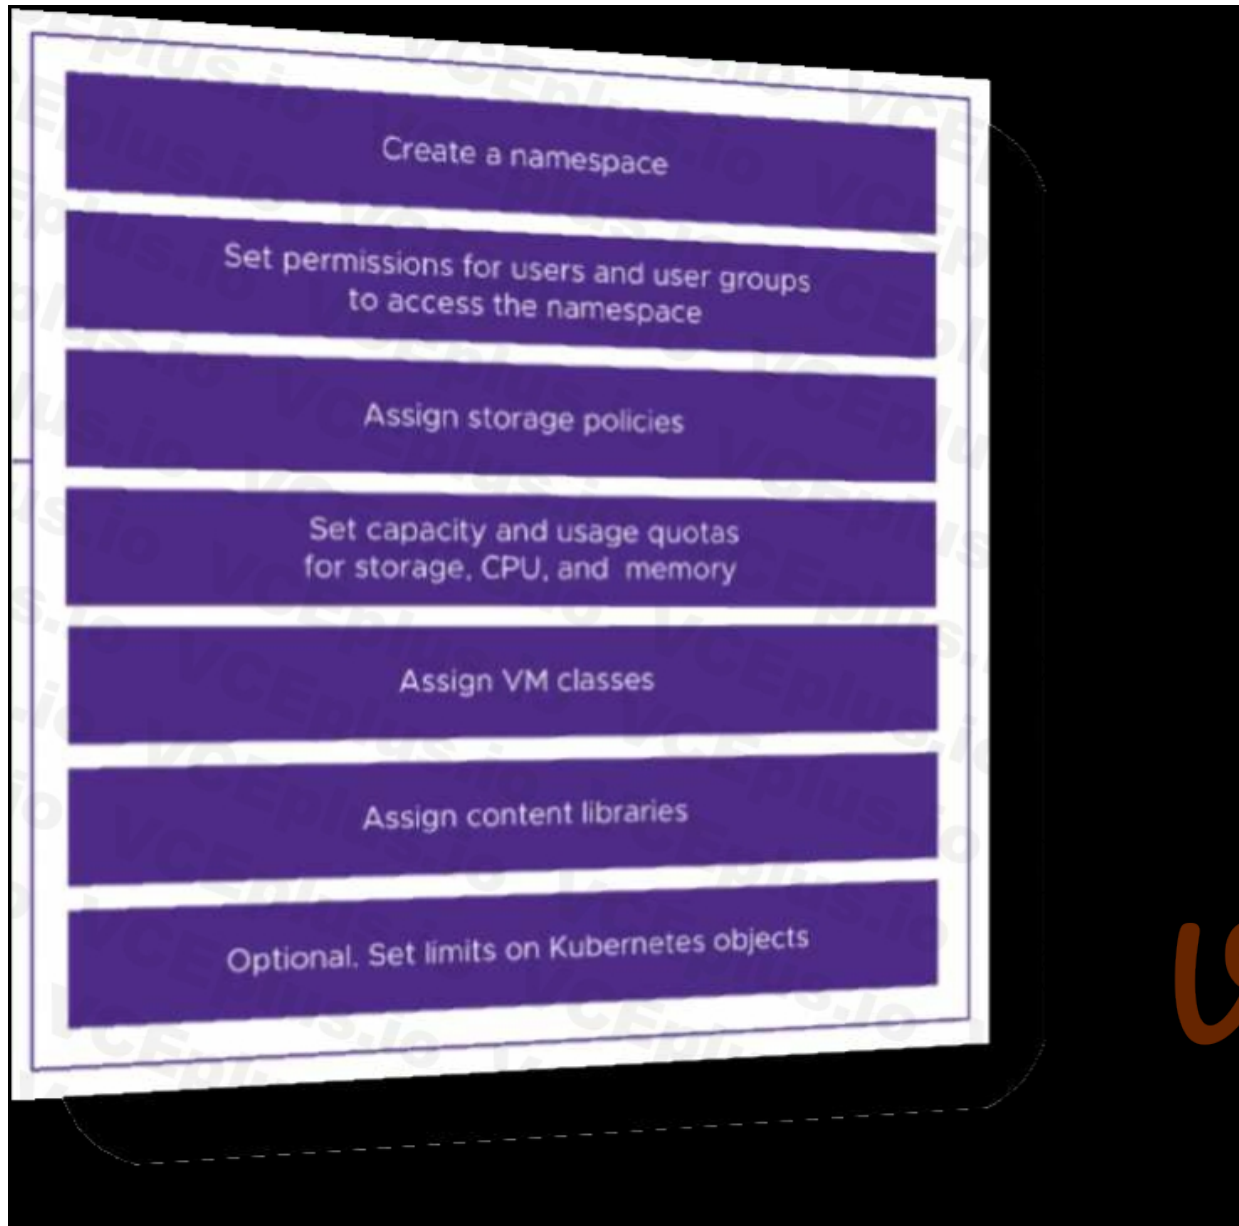
Permissions for programmers should be assign at the Namespace level, typically using groups and roles.
You assign roles for the Namespace to Active Directory groups. You can later assign access to users by adding them to these groups. You assign access to separate Active Directory groups for the edit and view roles in the Namespace.

QUESTION 26
A vSphere administrator is required to provide a developer with read/write permissions on a vSphere Namespace.

A. It is a Kubernetes cluster deployed by developers using a YAML specification file.
B. It aggregates hosts across vSphere clusters to form a Kubernetes cluster on vSphere.
C. It enables vSphere High Availability and Distributed Resource Scheduler.
D. It exposes a Kubernetes API for use by DevOps teams.

Correct Answer: B
Section:

QUESTION 27
On which cluster can a Supervisor Namespace be created?

A. A Tanzu Kubernetes Grid Integrated cluster

B.   A vSphere 7 cluster enabled with Workload Management

C.   A Tanzu Kubernetes cluster

D.   A vSphere 6.7 cluster enabled with Workload Management

Correct Answer: B
Section:
Explanation:
A Supervisor Cluster namespace is a combination of resource allocations and permissions set within the Supervisor Cluster. When you create a Supervisor Namespace, you'll assign who has access to use it, and how many of the ESXi cluster's resources you can use (much like a resource pool).
When you enabled the Workload Management components, you created a special Kubernetes cluster called the "Supervisor Cluster". You can continue to deploy virtual machines in this cluster, and you can also deploy kubernetes pods as a "pod vm" which is basically a container with some special wrapping so they are better isolated, like a virtual machine is.

QUESTION 28
Which command provides valid syntax to deploy a vSphere Pod?

A.   tkg apply -c containerName

B.   docker run containerName

C.   kubectl apply -f deployment.yaml

D.   kubectl apply -t deployment.yaml

Correct Answer: C
Section:
Explanation:
You can deploy an application on a namespace on a Supervisor Cluster. Once the application is deployed, the respective number of vSphere Pods are created on the Supervisor Cluster within the namespace.
Common kubectl commands include the apply, get, describe, and delete commands:
• The kubectl apply command applies the contents of a YAML file. Typically, this command is used to create a pod or deployment: — kubectl apply -f /path/to/my.yaml
• The kubectl get command returns basic information about an object: — kubectl get pod
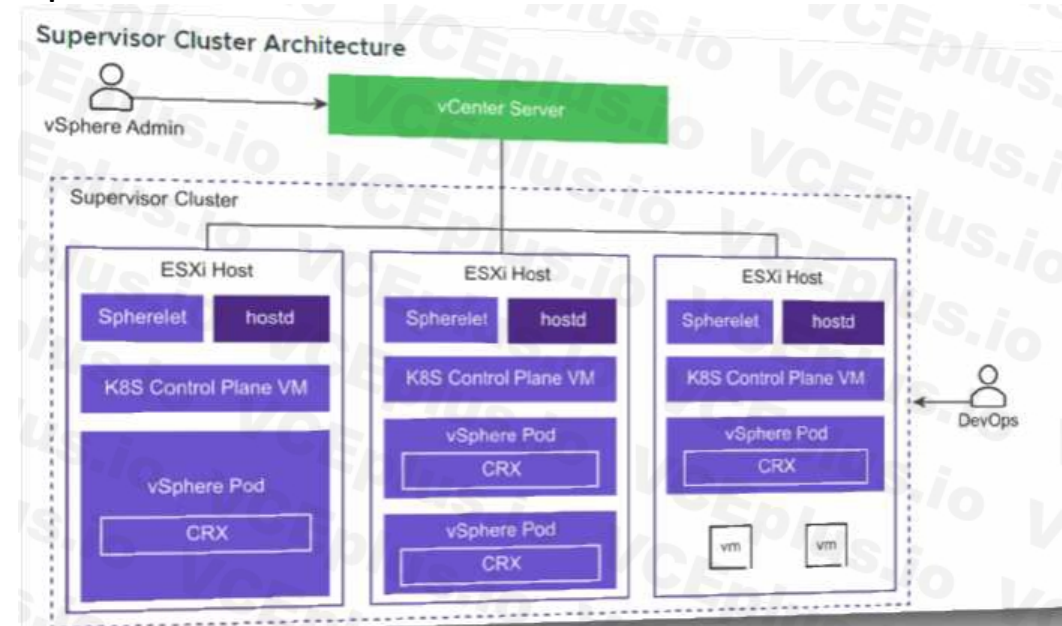<pod_name_name>

QUESTION 29

Which statement describes a characteristic of Supervisor Cluster control plane VMs?

A. Manage the lifecycle of ESXi hosts

B. Are manually created by a vSphere administrator

C. Host developer workloads

D. Run system and infrastructure pods

Correct Answer: D
Section:
Explanation:

The Supervisor Cluster provides the management layer on which Tanzu Kubernetes clusters are built.
The Tanzu Kubernetes Grid Service is a custom controller manager with a set of controllers that is part of the Supervisor Cluster. The purpose of the Tanzu Kubernetes Grid Service is to provision Tanzu Kubernetes clusters.
While there is a one-to-one relationship between the Supervisor Cluster and the vSphere cluster, there is a one-to-many relationship between the Supervisor Cluster and Tanzu Kubernetes clusters.
You can provision multiple Tanzu Kubernetes clusters within a single Supervisor Cluster. The workload management functionality provided by the Supervisor Cluster gives you control over the cluster configuration and lifecycle, while allowing you to maintain concurrency with upstream Kubernetes.

QUESTION 30
The creation of which object by an administrator in the vSphere client automatically results in the creation of a new segment within NSX -T?

A. Service

B. Pod

C. Network policy

D. Namespace

Correct Answer: D
Section:
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 31
How does Kubernetes implement the vSphere storage policy in vSphere with Tanzu?

A. Storage class

B. Paravirtual CSI

C. Static Persistent Volume

D. Persistent Volume

Correct Answer: A
Section:
Explanation:
When vSphere with Tanzu converts storage policies that you assign to namespaces into Kubernetes storage classes, it changes all upper case letters into lower case and replaces spaces with dashes (-).
To avoid confusion, use lower case and no spaces in the VM storage policy names.
Storage Policy Based Management is a vCenter Server service that supports provisioning of persistent volumes and their backing virtual disks according to storage requirements described in a storage policy.

QUESTION 32
Which two functions are provided by the NSX Container Rug-in (NCP)? (Choose two.)

A. Implements Kubernetes Ingress with an NSX-T layer 7 load balancer

B. Integrates with container-based PaaS such as Docker

C. Creates an NSX-T logical topology for a Kubernetes cluster and a separate logical network for each Kubernetes namespace

D. Configures Overlay Transport Zones

E. Implements Kubernetes Ingress with an NSX-T layer 4 load balancer

Correct Answer: A, C
Section:
Explanation:
NCP provides the following functionalities:
Automatically creates an NSX-T Data Center logical topology for a Kubernetes cluster, and creates a separate logical network for each Kubernetes namespace.
Implements Kubernetes Ingress with NSX-T layer 7 load balancer
Connects Kubernetes pods to the logical network, and allocates IP and MAC addresses.
Supports network address translation (NAT) and allocates a separate SNAT IP for each Kubernetes namespace.Note:When configuring NAT, the total number of translated IPs cannot exceed 1000.
Implements Kubernetes network policies with NSX-T Data Center distributed firewall.
Implements Kubernetes service of type ClusterIP and service of type LoadBalancer.

QUESTION 33
How do Tanzu Kubemetes clusters communicate with Storage Policy Based Management to request PersistentVolumes?

A. Through a proxy VM

B. Directly with vCenter Server and the underlying ESXi hosts

C. Through the Supervisor Cluster

D. Directly with the vCenter Server
    Answer: D
    Explanation:
    The Cloud Native Storage for vSphere with Tanzu workflow is as follows:

E.   A developer deploys a pod using the kubectl CLI.

F.   The vSphere with Tanzu Cloud Native Storage-Container Storage Interface (CNS-CSI) reads this request from the control plane API server.

G.   CNS-CSI informs the vCenter Server CNS of the need for a disk with storage class Gold.

H.   CNS interfaces with SPBM for a suitable datastore that satisfies the Gold storage class (storage policy).

I.   SPBM decides on a suitable datastore and interfaces with DRS for a suitable ESXi host.

J.   Hostd on the ESXi host creates a First Class Disk (VMDK) on the datastore.

K.   Spherelet on the ESXi host takes the created VMDK.

L.   Spherelet mounts the VMDK to the vSphere Pod.

M.

Correct Answer: D
Section:
Explanation:
The Cloud Native Storage for vSphere with Tanzu workflow is as follows:
1. A developer deploys a pod using the kubectl CLI.
2. The vSphere with Tanzu Cloud Native Storage-Container Storage Interface (CNS-CSI) reads this request from the control plane API server.
3. CNS-CSI informs the vCenter Server CNS of the need for a disk with storage class Gold.
4. CNS interfaces with SPBM for a suitable datastore that satisfies the Gold storage class (storage policy).
5. SPBM decides on a suitable datastore and interfaces with DRS for a suitable ESXi host.
6. Hostd on the ESXi host creates a First Class Disk (VMDK) on the datastore.
7. Spherelet on the ESXi host takes the created VMDK.
8. Spherelet mounts the VMDK to the vSphere Pod.
9. Spherelet reports the mount as a successful event to the control plane API server.

QUESTION 34
Which kubectl command is used to list al pods in the current active namespace?

A.   kubectl get nodes

B.   kubectl get pods

C.   kubectl get services

D.   kubectl list pods

Correct Answer: B
Section:
Explanation:
Fetch all Pods in all namespaces using kubectl get pods --all-namespaces

Pods

Shortcode = po

List one or more pods

kubectl get pod

Delete a pod

kubectl delete pod <pod_name>

Display the detailed state of a pods

kubectl describe pod <pod_name>

Create a pod

kubectl create pod <pod_name>

Execute a command against a container in a pod

kubectl exec <pod_name> -c <container_name> <command>

Get interactive shell on a a single-container pod

kubectl exec -it <pod_name> /bin/sh

Display Resource usage (CPU/Memory/Storage) for pods

kubectl top pod

Add or update the annotations of a pod

kubectl annotate pod <pod_name> <annotation>

Add or update the label of a pod

kubectl label pod <pod_name>

QUESTION 35
Which vSphere with Tanzu Workload Network topology provides Layer 2 isolation between Tanzu Kubernetes clusters across namespaces when vSphere Distributed Switches are used?

A. A dedicated Primary Workload Network for the Supervisor Cluster control plane VMs and separate Workload Networks for each namespace A dedicated Primary
B. Workload Network for the Supervisor Cluster control plane VMs and a single Workload Network for namespaces
C. A single Workload Network for the Supervisor Cluster control plane VMs and Tanzu Kubernetes clusters
D. Distributed firewall rules to isolate namespaces

Correct Answer: A
Section:
Explanation:
For a Supervisor Cluster that is configured with the vSphere networking stack, you can provide Layer 2 isolation for your Kubernetes workloads by creating Workload Networks and assigning them to namespaces. Workload Networks provide connectivity to Tanzu Kubernetes clusters in the namespace and are backed by distributed port groups on the switch that is connected to the hosts in the Supervisor Cluster.

QUESTION 36
To which network are HA Proxy virtual server IP addresses issued when using the vSphere networking stack default configuration?

A. vMotion

B. Overlay

C. Primary workload

D. Management

Correct Answer: C
Section:
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 37
How is information found about all Kubernetes Persistent Volumes in a vSphere environment?

A. Navigating to the Cloud Native Storage view in vCenter Server

B. Using: kubectl get persistentvolumes

C. Accessing the FCD folder on a Datastore

D. Using: esxcli storage cloud native get

Correct Answer: A
Section:
Explanation:

This is the textbook answer, I know kubectl does give you some information.

QUESTION 38
To which set of networks are the Supervisor Cluster nodes attached when deploying with an NSX-T network topology?

A. Frontend and Workload

B. Frontend and Management

C. Workload and Management

D. Management and NSX Overlay

Correct Answer: C
Section:
Explanation:
The Network Service has been extended to support the vSphere Distributed Switch (vDS). Start by configuring the switch with appropriate portgroups. Management will carry traffic between vCenter and the Kubernetes

Control Plane (Supervisor Cluster control plane). As we will see in a moment, not having the built in Load Balancing capability of NSX means you will need to deploy your own load balancer externally from the cluster. We will give you a choice of integrated load balancers. The first one we support is HAProxy.

The Management network will also carry traffic between the supervisor cluster nodes and HAProxy.

The Frontend network will carry traffic to the Load Balancer virtual interfaces. It must be routable from any device that will be a client for your cluster. Developers will use this to issue kubectl commands to the Supervisor cluster or their TKG clusters. You can have one or more Workload networks.

The primary Workload network will connect the cluster interfaces of the Supervisor cluster.

Namespaces can be defined with their own Workload network allowing for isolation between development teams assigned different Namespaces. The Namespace assigned Workload Networks will connect the TKG cluster nodes in that Namespace.

QUESTION 39
Kubernetes object types are going to be limited by an administrator within a vSphere with Tanzu namespace. Which three Kubernetes object types may be limited? (Choose three.)

A. Number of Persistent Volume Claims
B. Number of Pods
C. Number of Operators
D. Number of DaemonSets
E. Number of Ingress frontends
F. Number of Load Balancer Services

Correct Answer: A, B, F
Section:
Explanation:
https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-1238AFD8-

232C-4EFC-BD54-796CB9F8C45F.html

Resource Name Description

configmaps The total number of ConfigMaps that can exist in the namespace.

persistentvolumeclaims The total number of PersistentVolumeClaims that can exist in the namespace.

pods The total number of Pods in a non-terminal state that can exist in the namespace. A pod is in a terminal state if .status.phase in (Failed, Succeeded) is true.

replicationcontrollers The total number of ReplicationControllers that can exist in the namespace.

resourcequotas The total number of ResourceQuotas that can exist in the namespace.

services The total number of Services that can exist in the namespace.

services.loadbalancers The total number of Services of type LoadBalancer that can exist in the namespace.

services.nodeports The total number of Services of type NodePort that can exist in the namespace.

secrets The total number of Secrets that can exist in the namespace.

QUESTION 40
Which two capabilities are associated with vSphere Pod? (Choose two.)

A. Compatibility with vSphere vMotion
B. Compatibility with vSphere performance charts
C. Compatibility with NSX-V Datacenter
D. Compatibility with vSphere HA and DRS
E. Compatibility with Windows and Linux kernels

Correct Answer: C, D
Section:
Explanation:
Resource Management. vSphere DRS handles the placement of vSphere Pods on the Supervisor Cluster.

QUESTION 41
Why would an organization set up private image registries?

A. Role-based access control can be assigned by integrating the image registry with user identity management.
B. DevOps engineers are able to store virtual machine images in a central location.
C. Open source registry server projects enable organizations to modify them as necessary.
D. Public image registries lack enterprise support.

Correct Answer: A
Section:
Explanation:

VMware created Harbor in 2014. Harbor was shared with the community through an open-source license in 2016 and donated to the Cloud Native Computing Foundation (CNCF) in 2018.
Harbor is integrated into VMware products: vSphere Integrated Containers, Tanzu Kubernetes Grid Integrated Edition, and vSphere with Tanzu. The embedded Harbor for vSphere with Tanzu includes the following features: • Identity integration and role-based access control
• Graphical user interface
• Auditing of operations
• Management with labels

QUESTION 42
What needs to be deployed to access applications on vSphere with Tanzu?

A. vSphere Distributed Switch
B. Virtual router
C. Internal load balancer

D. External load balancer

Correct Answer: D
Section:

QUESTION 43
The application development team plans to test a few CPU intensive applications. The virtualization team is concerned about these applications impacting other teams.
How should the team manage this problem?

A. Set a CPU reservation on the Namespace Resource Pool in vCenter.
B. Add a CPU limit to the Namespace Resource Pool in vCenter.
C. Set a network policy to limit network bandwidth within the vSphere Namespace.
D. Add a CPU limit to the vSphere Namespace for the development team.

Correct Answer: B
Section:

QUESTION 44
An administrator needs to limit resource consumption within a vSphere with Tanzu namespace.
Which three actions should be taken to meet this goal? (Choose three.)

A. Limit the capacity of the storage able to be consumed.
B. Change the size of the Supervisor Control Plane nodes.
C. Limit the number of load balancers able to be deployed.
D. Limit the CPU and Memory able to be consumed.
E. Limit the number of Operators able to be deployed.
F. Limit the number of vSphere Pods able to be provisioned.

Correct Answer: A, D, F
Section:

QUESTION 45
Which statement describes the characteristics of vSphere with Tanzu using vSphere Distributed Switch network topology?

A. Supervisor Cluster control plane VMs are attached to the primary workload network.
B. Supervisor Cluster control plane VMs are attached to primary and non-primary workload networks.
C. vCenter Server is attached to all workload networks.
D. Load balancer appliance is attached to all workload networks.

Correct Answer: B
Section:

QUESTION 46
Which two actions are orchestrated by the NSX Container Plugin (NCP)? (Choose two.)

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

Correct Answer: A, D
Section:

QUESTION 47
Which two networks are used to attach Supervisor Cluster control plane VMs when using the vSphere networking stack? (Choose two.)

A. vMotion

B. Frontend

C. Primary workload

D. Management

E. Non-primary workload

Correct Answer: D, E
Section:

QUESTION 48
What should be increased or reduced in order to scale a Tanzu Kubernetes cluster up or down?

A. The number of Pods that it contains

B. The number of Namespaces that it contains

C. The number of Deployments that it contains

D. The number of worker nodes that it contains

Correct Answer: D
Section:

QUESTION 49
Which command will display the container image(s) used in a vSphere pod deployment name nginxdeployment?

A. kubectl get deployment nginx-deployment

B. kubectl get pod nginx

C. kubectl describe deployment nginx-deployment

D. docker ps

Correct Answer: C
Section:

QUESTION 50
What does vSphere with Tanzu integrate with to provision persistent storage for Kubernetes workloads?

A. Cloud Native Storage (CNS)

B. DaemonSet

C. S3 external volume

D. vSAN Direct for vSphere

Correct Answer: D
Section:

QUESTION 51
An administrator is tasked with horizontally scaling an existing Tanzu Kubernetes cluster from 6 to 4 nodes.
Which action does the administrator need to complete to ensure the cluster scales successfully when updating the YAML definition?

A. Decrease the number of control plane nodes to 4.
B. Decrease the number of worker nodes to 4.
C. Update the Virtual Machine Class Type.
D. Update the Kubernetes version to the latest supported version.

Correct Answer: B
Section:

QUESTION 52
Which statement accurately describes the upgrade of a vSphere with Tanzu Supervisor Cluster?

A. vCenter Server performs an in-place upgrade of the Supervisor Cluster control plane VMs.
B. vCenter Server orchestrates the rolling upgrade of Supervisory control plane VMs and upgrades the ESXi host spherelet component.
C. An administrator manually deploys new Supervisor Cluster control plane VMs and uses vSphere Lifecycle Manager to update the ESXi host spherelet component.
D. An administrator downloads and installs new RPMs to the Supervisor Cluster control plane VMs.

Correct Answer: C
Section:

QUESTION 53
Which description reflects a use case for an embedded image registry?

A. Manage source code changes
B. Configure network settings of the image
C. Maintain desired state configuration for containers
D. Secure artifacts with policies and role-based access control

Correct Answer: D
Section:

QUESTION 54
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?

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

Correct Answer: B
Section:

QUESTION 55
Which three elements are part of the Container Runtime for ESXi (CRX)? (Choose three.)

A. Image Service
B. Kube Proxy
C. Linux kernel
D. Customized VMX
E. Spherelet
F. Container Engine

Correct Answer: A, C, E
Section:

QUESTION 56
Which statement accurately describes the Primary Workload Network?

A. It carries traffic between load balancer and vCenter.
B. It carries traffic between load balancer and Supervisory control plane.
C. It is created by developers as part of TKG cluster deployment.
D. It carries traffic between Supervisory control plane and vCenter.

Correct Answer: D
Section:

QUESTION 57
When creating a Supervisor Namespace, which item must be added to ensure that the end user can consume the namespace?

A. Content Library
B. Permissions
C. Storage
D. Limits

Correct Answer: B
Section:

QUESTION 58
Which command would be used to deploy a Tanzu Kubernetes cluster?

A. kubectl create cluster -f CLUSTER-NAME.yaml
B. kubectl apply cluster CLUSTER-NAME.yaml
C. kubectl apply -f CLUSTER-NAME.yaml
D. kubectl create -f CLUSTER-NAME.yaml

Correct Answer: C

Section:

QUESTION 59
An administrator is configuring a vSphere with Tanzu Supervisor Cluster with the vSphere networking stack.
Which two minimum requirements must be met for the compute and networking components?
(Choose two.)

A. The cluster configured with vSphere High Availability enabled
B. A DHCP IP address range for the Kubernetes control plane VMs
C. A DHCP IP address range for the HA Proxy virtual IPs
D. A minimum of three distinct subnets
E. The cluster configured with vSphere DRS enabled and automation level set to Fully Automated

Correct Answer: D, E
Section:

QUESTION 60
Which element defines the type of storage backing the Persistent Volume?

A. Datastore Cluster
B. Storage Class
C. vVols
D. Persistent Volume Claim

Correct Answer: B
Section:

QUESTION 61
At which three levels of the vSphere Client is information about Kubernetes Persistent Volumes found? (Choose three.)

A. Resource Pool
B. Cluster
C. vCenter
D. Datastore
E. Virtual Machine
F. Network

Correct Answer: B, C, D
Section:

QUESTION 62
Which summary page should an administrator use in the vSphere Client to identify the CPU resources consumed in a namespace?

A. ESXi host
B. Namespace
C. vSphere cluster

D. Supervisor Cluster

Correct Answer: B
Section:

QUESTION 63
Which statement accurately describes a characteristic of load balancers in vSphere with Tanzu using the vSphere networking stack?

A. A load balancer balances pods between Tanzu Kubernetes cluster nodes.

B. A load balancer distributes Tanzu Kubernetes cluster nodes between ESXi hosts.

C. A load balancer provides external access to Supervisor Clusters and Tanzu Kubernetes clusters.

D. A load balancer performs authentication and authorization for Supervisor Clusters and Tanzu Kubernetes clusters.

Correct Answer: C
Section:

QUESTION 64
A Namespace contains multiple Tanzu Kubernetes clusters.
How is access granted to a single Tanzu Kubernetes cluster?

A. Create a custom Role and RoleBinding, and then apply to the Namespace using kubectl commands.

B. Use the vSphere Client to grant access to the Namespace.

C. Create a custom Role and RoleBinding, and then apply to the Tanzu Kubernetes cluster using kubectl commands.

D. Use the vSphere Client to grant access to the Tanzu Kubernetes cluster.

Correct Answer: C
Section:

QUESTION 65
What is automatically created within NSX-T when an administrator creates a new namespace in the vSphere client?

A. A new segment connected to a Tier-1 gateway and a new overlay transport zone

B. A new segment connected to a Tier-0 gateway and a new overlay transport zone

C. A new segment connected to a Tier-0 gateway and an existing overlay transport zone

D. A new segment connected to a Tier-1 gateway and an existing overlay transport zone

Correct Answer: D
Section:

QUESTION 66
Which kubectl command is used to list the Kubernetes services in the current active namespace?

A. kubectl get services

B. kubectl get loadbalancer

C. kubectl services get ip

D. kubectl list services

Correct Answer: A
Section:

QUESTION 67
A development team has submitted a YAML specification to create a Tanzu Kubernetes cluster called dev-cluster and wants to verify that the cluster is running.
Which command should be run to obtain this verification?

A. kubectl show tkc dev-cluster

B. kubectl get cluster dev-cluster

C. kubectl get tkc dev-cluster

D. kubectl list cluster dev-cluster

Correct Answer: B
Section:

QUESTION 68
How are updates applied to Tanzu Kubernetes clusters?

A. The cluster is shutdown, Update Manager applies patches to VMs, and the cluster is restarted.

B. Worker nodes are shutdown while Controllers are patched and restarted after the update.

C. New nodes are incrementally added with new Kubernetes versions so as a new node is added, an older node is removed.

D. A new cluster is completely configured and runs in parallel to the primary cluster prior to switchover.

Correct Answer: C
Section:

QUESTION 69
Which kubectl command is used to deploy the application when using a Kubernetes deployment specification file, my-app.yaml?

A. kubectl run my-app.yaml

B. kubectl apply -f my-app.yaml

C. kubectl create my-app.yaml

D. kubectl apply sp?c my-app.yaml

Correct Answer: B
Section:

QUESTION 70
Which step is necessary in order to allocate storage to a newly created namespace?

A. Assign a VM Storage Policy

B. Assign permissions to users

C. Define resource limits

D. Add a content library

Correct Answer: A
Section:

QUESTION 71
An administrator needs to label and push an image into Harbor.
Which command should the administrator run first?

A. docker tag demoapp01:14.04 /demo/demoapp01:14.04
B. kubectl tag demoapp01:14.04 /demo/demoapp01:14.04
C. kubectl label demoapp01:14.04 /demo/demoapp01:14.04
D. docker label demoapp01:14.04 /demo/demoapp01:14.04

Correct Answer: A
Section:

QUESTION 72
Why would developers choose to deploy an application as a vSphere Pod instead of a Tanzu Kubernetes cluster?

A. They need the application to run as privileged pods.
B. The application works with sensitive customer data, and they want strong resource and security isolation.
C. They want to have root level access to the control plane and worker nodes in the Kubernetes cluster.
D. The application requires a version of Kubernetes that is above the version running on the supervisor cluster.

Correct Answer: B
Section:
Explanation:
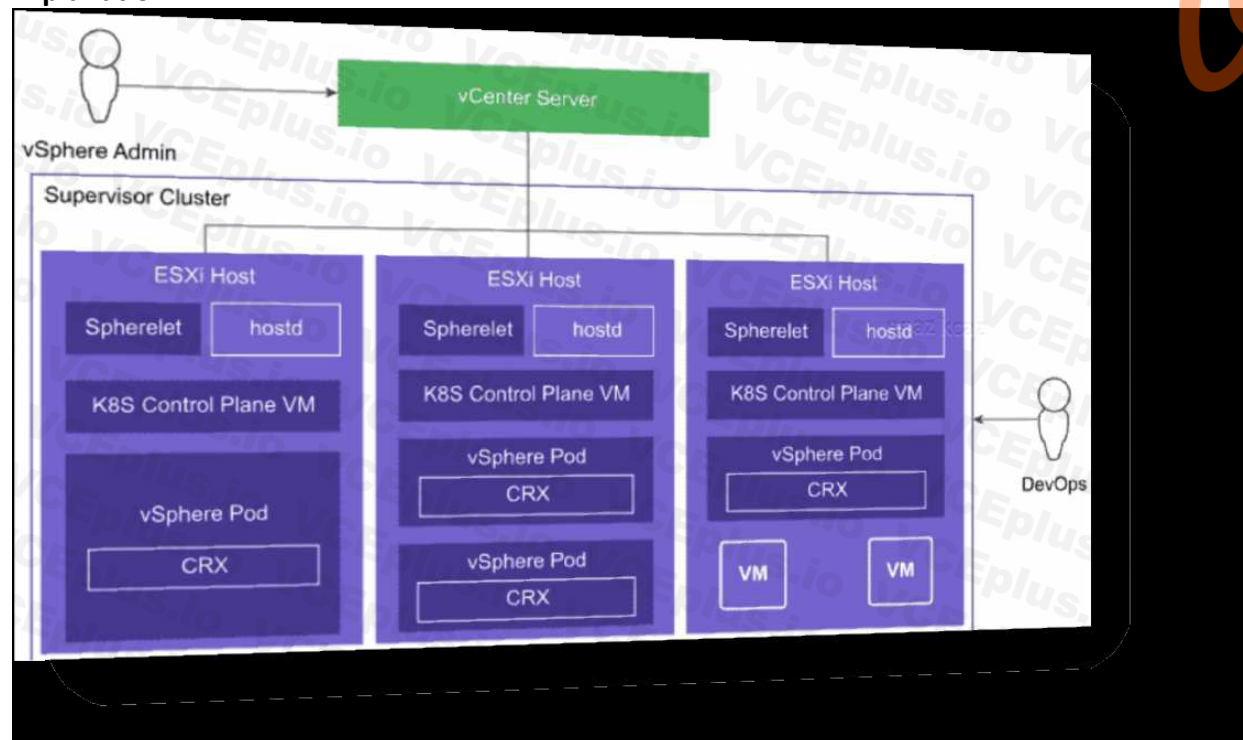
A vSphere Pod is a VM with a small footprint that runs one or more Linux containers. With vSphere Pods, workloads have the following capabilities:
• Strong isolation from a Linux kernel based on Photon OS
• Resource management using DRS
• Same level of resource isolation as VMs
• Open Container Initiative (OCI) compatible
• Equivalent to a Kubernetes Container Host vSphere Pods are not compatible with vSphere vMotion. When an ESXi host is placed into maintenance mode, running vSphere Pods are drained and redeployed on another ESXi host, but only if the vSphere Pod is part of a ReplicaSet.
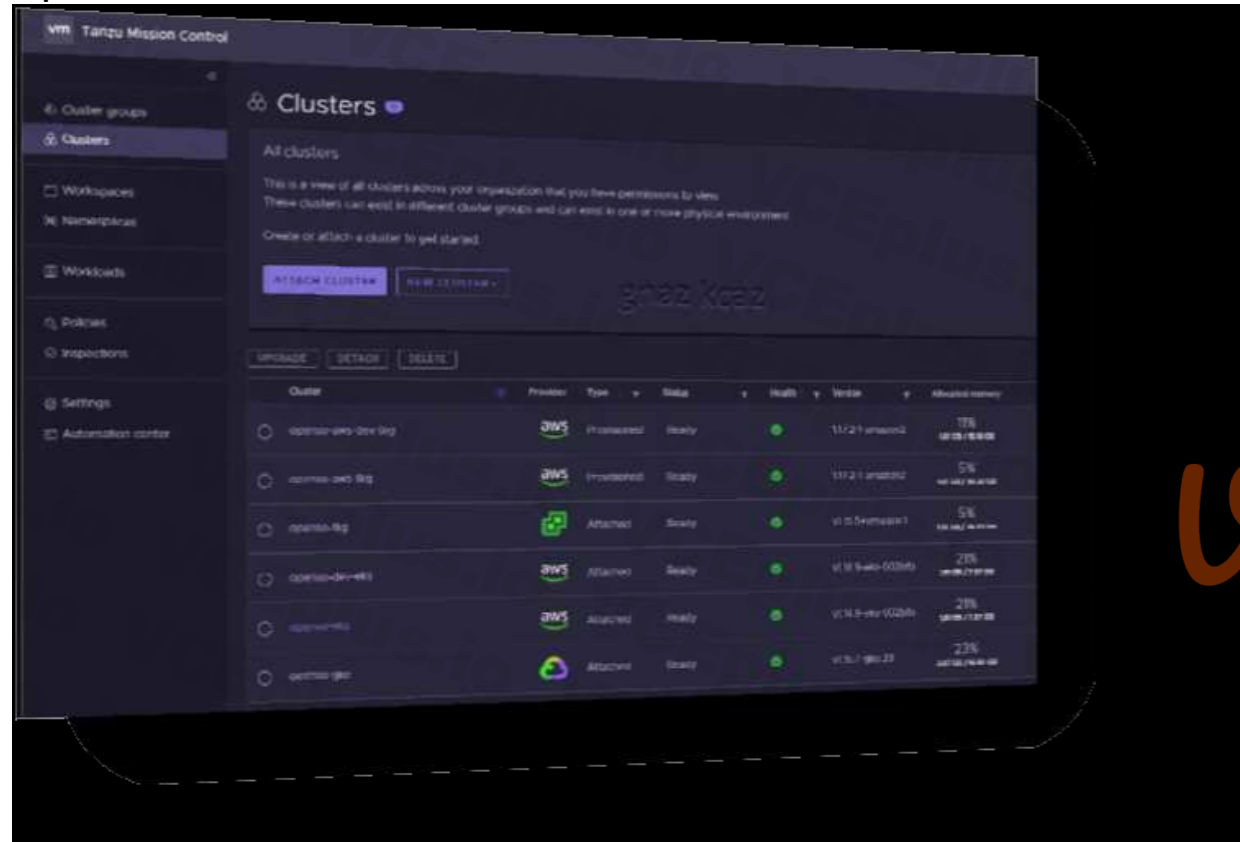
QUESTION 73

A company needs to provide global visibility and consistent policy management across multiple Tanzu Kubernetes Clusters, namespaces, and clouds Which VMvare solution will meet these requirements'?

A. vSphere with Tanzu Supervisor Cluster
B. vCenter Server
C. Tanzu Mission Control
D. Tanzu Kubernetes Grid Service

Correct Answer: C
Section:
Explanation:

VMware Tanzu Mission Control™ is a centralized management platform for consistently operating and securing your Kubernetes infrastructure and modern applications across multiple teams and clouds.

QUESTION 74

A developer is connecting to a Tanzu Kubernetes Cluster using the kubectl vsphere login command Which information must be specified, in addition to both the name of the cluster and the Supervisor Cluster Control Wane IP?

A. The path to the existing kubeconfig file and the SSO Username
B. The path to the existing kubeconfig file and the Token D for the SSO credentials
C. The name of the Supervisor Namespace and the Token ID for the SSO credentials
D. The name of the Supervisor Namespace and the SSO Username

Correct Answer: D
Section:
Explanation:
To connect to the Supervisor Cluster, run the following command. kubectl vsphere login --server=SUPERVISOR-CLUSTER-CONTROL-PLANE-IP --tanzu-kubernetes-cluster-name TANZU-KUBERNETES-CLUSTER-NAME --tanzu-kubernetes-cluster-namespace SUPERVISOR-NAMESPACE-WHERE-THE-CLUSTER-ISDEPLOYED --vsphere-username VCENTER-SSO-USER-NAME For example: kubectl vsphere login --server=10.92.42.137 --tanzu-kubernetes-

cluster-name tanzu-kubernetes-cluster-01 --tanzu-kubernetes-cluster-namespace tanzu-ns-1 --vsphere-username administrator@example.com

QUESTION 75
Which value must be increased or decreased to horizontally scale a Tanzu Kubernetes cluster?

A. Namespaces

B. etcd instance

C. Worker node count

D. ReplicaSets

Correct Answer: C
Section:
Explanation:
Scale a Cluster Horizontally With the Tanzu CLI
To horizontally scale a Tanzu Kubernetes cluster, use the tanzu cluster scale command. You change the number of control plane nodes by specifying the --controlplane-machine-count option. You change the number of worker nodes by specifying the --worker-machine-count option.

QUESTION 76
Which object is deployed on the workload network when vSphere with Tanzu is configured with vSphere network stack?

A. Windows Virtual Machines

B. Registry Service

C. Tanzu Kubernetes clusters

D. NSX Edge

Correct Answer: C
Section:

QUESTION 77
The network topology for a Supervisor Cluster deployed using the vSphere networking stack, and a HAProxy load balancer is being planned. In addition to the control plane management IP range and services P range, how many non-overlapping P address ranges are needed?

A. 3

B. 1

C. 4

D. 2

Correct Answer: B
Section:
Explanation:
A dedicated IP range for virtual IPs. The HAProxy VM must be the only owner of this virtual IP range.
The range must not overlap with any IP range assigned to any Workload Network owned by any Supervisor Cluster.

QUESTION 78
How do vSphere Namespaces allow different teams to share a Kubernetes cluster?

A. Attaches authorization and policy to a subsection of the cluster

B. Installs Calico CNI
C. Installs a network policy provider
D. Advertises extended resources for a node

Correct Answer: A
Section:

QUESTION 79
What can be used to define whether a vSphere pod can be scaled?

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

Correct Answer: A
Section:

QUESTION 80
How can the vSphere administrator purge unused images from the image registry?

A. Enable the garbage collector service to purge all unreferenced images from the corresponding namespace.
B. From the Harbor UI, navigate to the project associated with namespace, and select the option to purge unreferenced images.
C. Download the vSphere Docker Credential Helper CLI Tool, and use it to purge all unreferenced images from the corresponding namespace.
D. From vSphere Client, navigate to the namespace, and select the option to purge unreferenced images.

Correct Answer: B
Section:

QUESTION 81
Which corresponding object is automatically created in the embedded Registry Service when a new Namespace is created in the Supervisor Cluster?

A. Pod
B. Project
C. Container
D. Image

Correct Answer: B
Section: