VMware.5V0-23.20.v2023-02-20.q20

Exam Code:	5V0-23.20
Exam Name:	VMware vSphere with Tanzu Specialist
Certification Provider:	VMware
Free Question Number:	20
Version:	v2023-02-20
# of views:	543
# of Questions views:	200
https://www.freepdfdumps.com/VMware.5V0-23.20.v2023-02-20.q20.html	

NEW QUESTION: 1
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
Answer: A (LEAVE A REPLY)
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.

NEW QUESTION: 2
A development team has deployed a Tanzu Kubernetes cluster and would like to verify the version of Kubernetes that is running. Which command will show this information?
A. kubectl describe tkc dev-cluster
B. kubectl explain tkg dev-cluster
C. kubectl get version
D. kubectl get vm dev-cluster
Answer: C (LEAVE A REPLY)
Print the client and server version information.
Synopsis
Print the client and server version information.
kubectl version

NEW QUESTION: 3
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
Answer: C (LEAVE A REPLY)
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.

NEW QUESTION: 4
Where is a storage policy applied to enable Persistent Volumes?
A. Namespace
B. Datastore
C. Virtual Machine
D. Cluster
Answer: A (LEAVE A REPLY)
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.

NEW QUESTION: 5
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
Answer: (SHOW ANSWER)

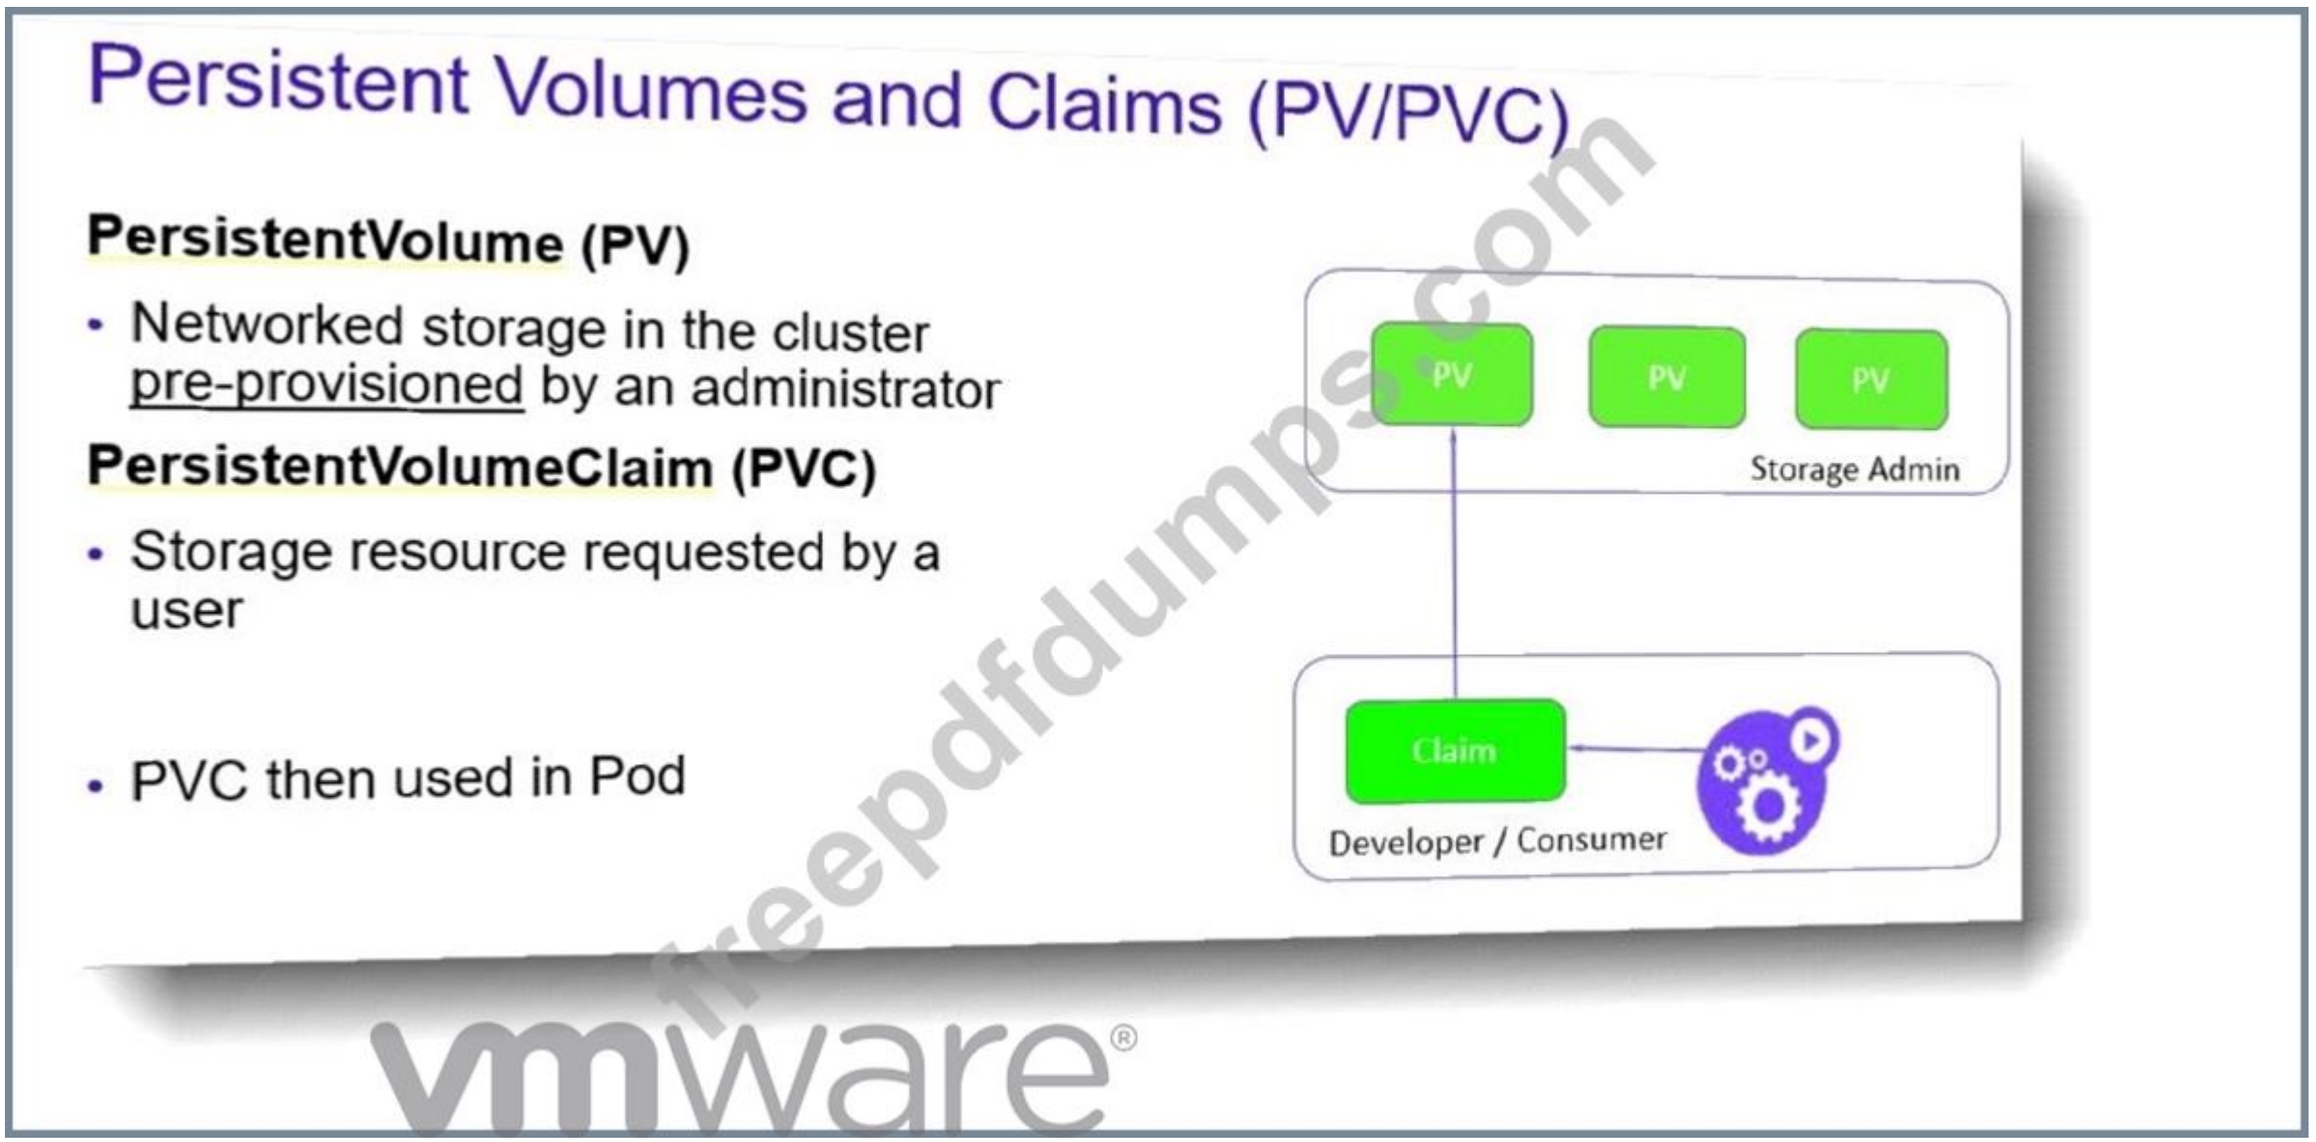

Also, kubectl delete pvc, which is much shorter.

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.

NEW QUESTION: 6

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

Answer: (SHOW ANSWER)

* 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.

NEW QUESTION: 7
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
Answer: C (LEAVE A REPLY)
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>

NEW QUESTION: 8
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
Answer: B (LEAVE A REPLY)
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.

NEW QUESTION: 9
How can you remove unreferenced container images from a project in an embedded Registry Service?
A. Delete images in Content Library.
B. Use kubectl to delete the images.
C. Delete the namespace using the vSphere Client.
D. Purge a namespace using the vSphere Client.
Answer: A (LEAVE A REPLY)
Deleting Artifact:
When an artifact is not referenced by any OCI index, you can delete the artifact freely which will delete its manifest and all associated tags.
When an artifact is referenced by an OCI index, you cannot delete it. In order to delete this artifact, you must first delete all OCI indexes referencing this artifact first, remembering that an artifact can be referenced by multiple parents artifacts pushed onto Harbor by different users. So when deleting an OCI index holding 9 children artifacts not referenced by any other index and 1 child artifact referenced by another index, only 9 out of 10 children artifacts will be deleted.
To delete any artifact in the Harbor interface, click on the artifact and select 'Delete' and confirm.
Not Purge: As a vSphere administrator, you can purge the images for a project in the private image registry by request from DevOps engineers. Purging images from the private image registry deletes all references to the images made by pods, but it does not remove the images from the image registry.

NEW QUESTION: 10
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
Answer: (SHOW ANSWER)

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

NEW QUESTION: 11
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
Answer: B (LEAVE A REPLY)
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

NEW QUESTION: 12
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 (LEAVE A REPLY)
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.

NEW QUESTION: 13
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
Answer: D (LEAVE A REPLY)

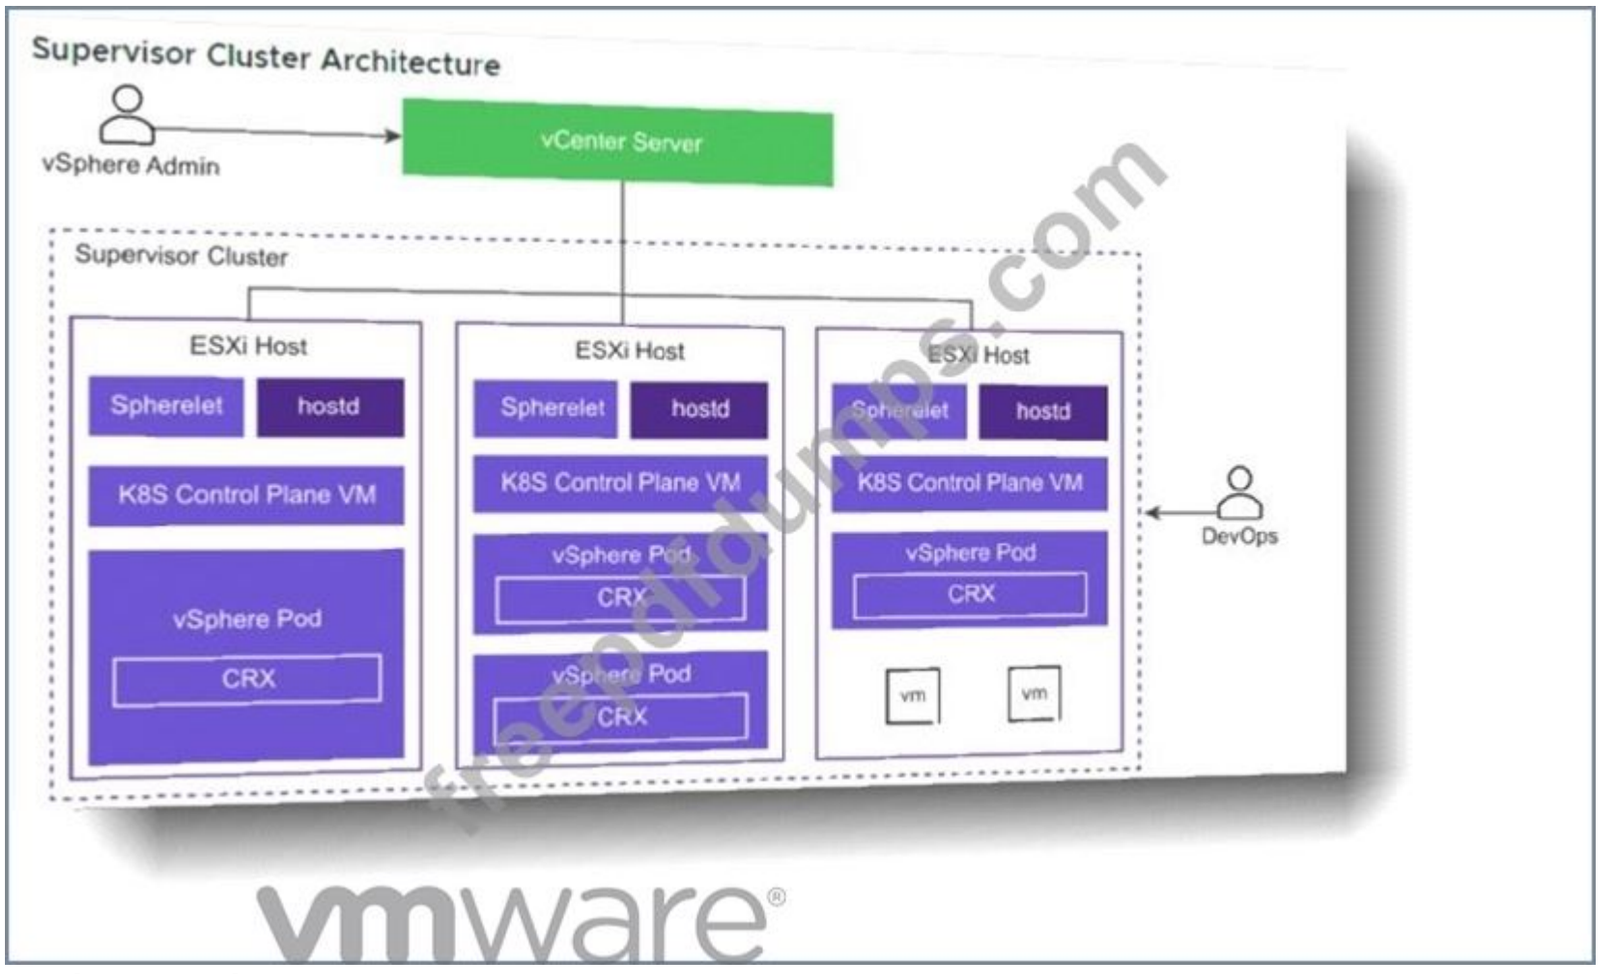

The Supervisor Cluster provides the management layer on which Tanzu Kubernetes clusters are built. The Tanzu Kubernetes Grid Service is a custom controller manager with a set of controllers that is part of the Supervisor Cluster. The purpose of the Tanzu Kubernetes Grid Service is to provision Tanzu 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.

NEW QUESTION: 14
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
Answer: C (LEAVE A REPLY)
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>

NEW QUESTION: 15

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.

Answer: B (LEAVE A REPLY)

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.

NEW QUESTION: 16
Which open-source project extends the Docker registry source code to provide an enterprise-class registry server?
A. Namespace
B. Manifest
C. Harbor
D. Github
Answer: C (LEAVE A REPLY)
VMware Harbor Registry
VMware Harbor Registry is an enterprise-class registry server that stores and distributes container images. Harbor allows you to store and manage images for use with VMware Tanzu Kubernetes Grid Integrated Edition (TKGI).
Overview
Harbor extends the open source Docker Distribution by adding the functionalities usually required by an enterprise, such as security, identity, and management. As an enterprise private registry, Harbor offers enhanced performance and security. Deploying a registry alongside the TKGI environment improves image management efficiency.

NEW QUESTION: 17
Which external load balancer is supported in vSphere 7 U1 using the vSphere networking stack?
A. Nginx
B. Seesaw
C. Loadmaster
D. HAProxy
Answer: D (LEAVE A REPLY)
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.

NEW QUESTION: 18
Which object helps maintain copies of a vSphere pod?
A. ReplicaSets
B. Network Policies
C. Namespaces
D. Persistent Volume
Answer: A (LEAVE A REPLY)

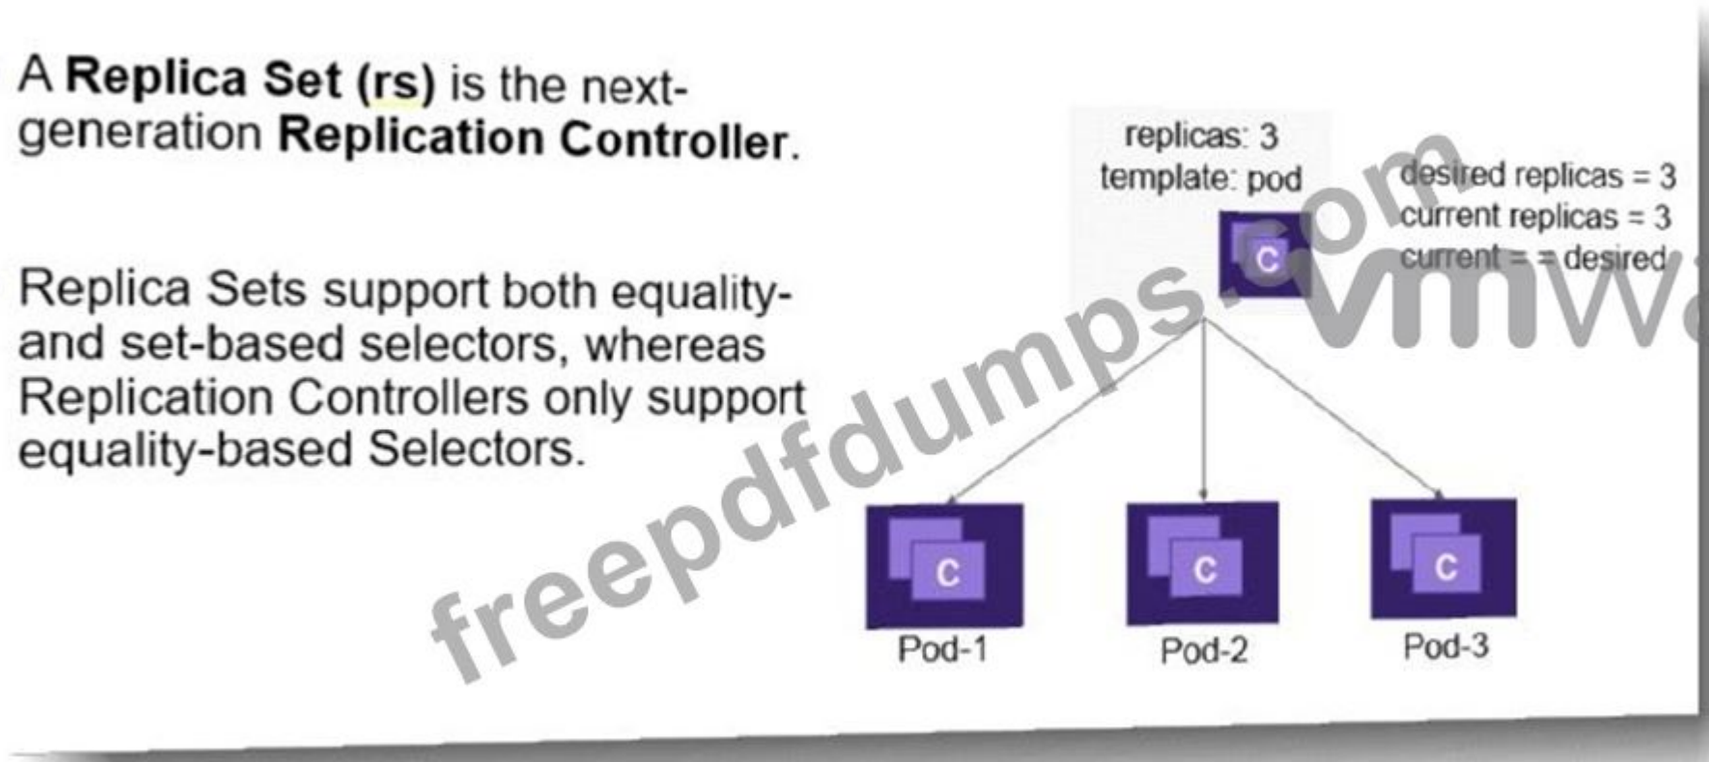

A ReplicaSet declares how the functionality of a pod is made scalable and resilient through redundancy. The ReplicaSet ensures that a specified number of pods is kept running. Example: Deploy a ReplicaSet. * The ReplicaSet name is nginx-replica-demo. * Two replicas are expected to be running. * The ReplicaSet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

NEW QUESTION: 19
Which process should be used to upgrade the vSphere with Tanzu Supervisor Cluster?
A. Use the vSphere Client, navigate to Workload Management, and apply updates.
B. Use kubectl, and apply an update manifest specification to the Supervisor Cluster.
C. Use the vSphere Client, navigate to vSphere lifecycle Manager, and apply updates
D. Allow vSphere with Tanzu Supervisor Cluster to upgrade automatically when new versions are available.
Answer: (SHOW ANSWER)
From the vSphere Client menu, select Workload Management.
Select the Updates tab.
Select the Available Version that you want to update to.
For example, select the version v1.17.4-vsc0.0.2-16293900.
Select the Supervisor Cluster to apply the update to.
To initiate the update, click Apply Updates.
Use the Recent Tasks pane to monitor the status of the update.

NEW QUESTION: 20
Which functionality does the Cloud Native Storage (CNS) component take advantage of to support the creation of container volumes?
A. Virtual Disk
B. VMware Disk Encryption
C. First Class Disk
D. Storage Based Policy Management
Answer: C (LEAVE A REPLY)