

[November-2022Braindump2go 5V0-23.20 PDF Dumps for Free[Q20-Q50


November/2022 Latest Braindump2go 5V0-23.20 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go 5V0-23.20 Real Exam Questions!
QUESTION 20Which two considerations needs to be made when deciding on a virtual machine class type during the process of creating a Tanzu Kubernetes cluster? (Choose two)
A. Whether the resources provided by the virtual machine class type should be reserved on the host
B. The configuration parameters which need to be edited in the cluster
C. The amount of CPU, memory, and storage the virtual machine should have
D. Connectivity between the Tanzu Kubernetes cluster and the Subscribed Content Library
E. The storage classes which need to be made available to the cluster
Answer: AC
Explanation:

Class	CPU	Memory (GB)	Storage (GB)	Reserved CPU and Memory
guaranteed-8xlarge	32	128	16	Yes
best-effort-8xlarge	32	128	16	No
guaranteed-6xlarge	16	128	16	Yes
best-effort-6xlarge	16	128	16	No
guaranteed-2xlarge	8	64	16	Yes
best-effort-2xlarge	8	64	16	No
guaranteed-xlarge	4	32	16	Yes
best-effort-xlarge	4	32	16	No
guaranteed-large	4	16	16	No
best-effort-large	4	16	16	No
guaranteed-medium	2	8	16	Yes
best-effort-medium	2	8	16	No
guaranteed-small	2	4	16	Yes
best-effort-small	2	4	16	No
guaranteed-xsmall	2	2	16	Yes
best-effort-xsmall	2	2	16	No


Class	CPU	Memory (GB)	Storage (GB)	Reserved CPU and Memory
guaranteed-8xlarge	32	128	16	Yes
best-effort-8xlarge	32	128	16	No
guaranteed-6xlarge	16	128	16	Yes
best-effort-6xlarge	16	128	16	No
guaranteed-2xlarge	8	64	16	Yes
best-effort-2xlarge	8	64	16	No
guaranteed-xlarge	4	32	16	Yes
best-effort-xlarge	4	32	16	No
guaranteed-large	4	16	16	No
best-effort-large	4	16	16	No
guaranteed-medium	2	8	16	Yes
best-effort-medium	2	8	16	No
guaranteed-small	2	4	16	Yes
best-effort-small	2	4	16	No
guaranteed-xsmall	2	2	16	Yes
best-effort-xsmall	2	2	16	No

A virtual machine class is a request for resource reservations for processing power on the virtual machine (VM), including CPU and memory (RAM). For example, the VM class type named "guaranteed-large" reserves 4 CPU and 16 GB RAM. See Default Virtual Machine Classes for a list of default VM classes and their corresponding CPU and RAM reservations. The VM disk size is set by the OVA template, not the VM class definition. For Tanzu Kubernetes releases, the disk size is 16GB. See About Tanzu Kubernetes release Distributions. There are two reservation types for VM classes: guaranteed and best effort. The guaranteed class fully reserves its configured resources. This means that for a given cluster the spec.policies.resources.requests matches the spec.hardware settings. The best effort class allows resources to be overcommitted. For production workloads it is recommended that you use the guaranteed VM class type.
QUESTION 21An administrator working in a vSphere with Tanzu environment wants to ensure that all persistent volumes configured by developers within a namespace are placed on a defined subset of datastores. The administrator has applied tags to the required datastores in the vSphere Client. Which action should the administrator take next to meet the requirement?
A. Create a storage policy containing the tagged datastores, and apply it to the vSphere Namespace.
B. Create a storage class containing the tagged datastores, and apply it to the Supervisor Cluster.
C. Create a persistent volume claim containing the tagged

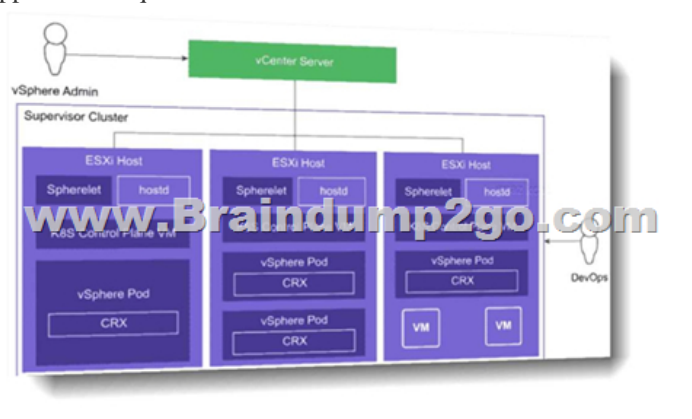
datastores, and apply it to the vSphere Namespace.D. Create a storage Policy containing the tagged datastores, and apply it to the Supervisor Cluster.Answer: AExplanation:



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.Storage class names are created in the following way:- Spaces in VM Storage Policy names are replaced with hyphens (-). - Special characters are replaced with a digit. A VM Storage Policy called My Gold Policy \$ is called my-gold-policy-0 as a storage class.QUESTION 22Which three roles does the Spherelet perform? (Choose three)A. Determines placement of vSphere podsB. Manages node configurationC. Starts vSphere podsD. Provides a key-value store for pod configurationE. Communicates with Kubernetes APIF. Provisions Tanzu Kubernetes clustersAnswer: BCEExplanation:



Spherelet is a kubelet that is ported natively to ESXi. It allows the ESXi host to become part of a Kubernetes cluster. Spherelet performs the following functions:- Communicates with the control plane VMs- Manages node configuration- Starts vSphere Pods- Monitors vSphere PodsQUESTION 23Why 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.Answer: BExplanation:



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 HostvSphere 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 24A company needs to provide global visibility and consistent policy management across multiple Tanzu Kubernetes Clusters, namespaces, and clouds.Which VMware solution will meet these requirements?A. vSphere with Tanzu Supervisor ClusterB. vCenter ServerC. Tanzu Mission ControlD. Tanzu Kubernetes Grid ServiceAnswer: CExplanation:



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 25A 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 Plane IP?A. The path to the existing kubeconfig file and the SSO UsernameB. The path to the existing kubeconfig file and the Token D for the SSO credentialsC. The name of the Supervisor Namespace and the Token ID for the SSO credentialsD. The name of the Supervisor Namespace and the SSO UsernameAnswer: DExplanation: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-IS-DEPLOYED--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 26Which value must be increased or decreased to horizontally scale a Tanzu Kubernetes cluster?A. NamespacesB. etcd instanceC. Worker node countD. ReplicaSetsAnswer: CExplanation: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 27The 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?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 databaseC. Yes, by using kubectl to create a policy that disables pod to pod communication in the NamespaceD. No, Kubernetes Network Policy does not support this action.Answer: AExplanation: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 28Which type of service is created by default when publishing a Kubernetes service?A. Cluster IPB. Node PortC. LoadBalancerD. ExternalNameAnswer: AExplanation:For some parts of your application (for example, frontends)

you may want to expose a Service onto an external IP address, that's outside of your cluster. Kubernetes ServiceTypes allow you to specify what kind of Service you want. The default is ClusterIP.

QUESTION 29 Which kubectl command should be used to change the active vSphere namespace to namespace-01?

A. `kubectl config use-context namespace-01`
B. `kubectl describe ns namespace-01`
C. `kubectl get ns namespace-01`
D. `kubectl config change-context namespace-01`

Answer: A
Explanation: A context element in a kubeconfig file is used to group access parameters under a convenient name. Each context has three parameters: cluster, namespace, and user. By default, the kubectl command-line tool uses parameters from the current context to communicate with the cluster. To choose the current context: `kubectl config use-context ctx001`

QUESTION 30 Which two items must be provided before a vSphere with Tanzu Supervisor Namespace can be created? (Choose two.)

A. vSphere with Tanzu Enabled Cluster
B. DNS-compliant Name
C. Permissions
D. Storage Policy
E. Resource Limits

Answer: AD
Explanation: Prerequisites

Configure a cluster with vSphere with Tanzu. Create users or groups for all DevOps engineers who will access the namespace. Create storage policies for persistent storage. Storage policies can define different types and classes of storage, for example, gold, silver, and bronze. Create VM classes and content libraries for stand-alone VMs. Create a content library for Tanzu Kubernetes releases for use with Tanzu Kubernetes clusters. See [Creating and Managing Content Libraries for Tanzu Kubernetes releases](#). Required privileges: Namespaces. Modify cluster-wide configuration. Namespaces. Modify namespace configuration

QUESTION 31 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.

Answer: A
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 32 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 IP range, how many non-overlapping IP address ranges are needed?

A. 3
B. 1
C. 4
D. 2

Answer: B
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 33 Where is a storage policy applied to enable Persistent Volumes?

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

Answer: A
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 34 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 datastoreD. Disconnect non-NFS datastores from the ESXi hosts that make up the Supervisor cluster.

Answer: C

Explanation: The storage policy would become a storage class in Kubernetes.

QUESTION 35 Which command displays the storage limits that have been set together with the amount of resources consumed?

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

Answer: A

Explanation: Create the ResourceQuota: `kubectl apply -f https://k8s.io/examples/admin/resource/quota-mem-cpu.yaml --namespace=quota-mem-cpu-example` View detailed information about the ResourceQuota: `kubectl get resourcequota mem-cpu-demo --namespace=quota-mem-cpu-example --output=yaml`

QUESTION 36 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)

Answer: D

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 37 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

Answer: A

QUESTION 38 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`

Answer: C

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 39 How does Kubernetes implement the vSphere storage policy in vSphere with Tanzu?



A. Storage class B. Paravirtual CSIC. Static Persistent Volume D. Persistent Volume



Answer: A



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 40 Which two functions are provided by the NSX Container Plug-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



Answer: A



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 41 How do Tanzu Kubernetes clusters communicate with Storage Policy Based Management to request Persistent Volumes?



A. Through a proxy VMB. 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:



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 42 Which kubectl command is used to list all pods in the current active namespace? A. kubectl get nodes B. kubectl get pods C. kubectl get services D. kubectl list pods

Answer: B

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 pod 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 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 43 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 B. A dedicated Primary

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

Answer: A

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 44 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

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 45 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: A

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

QUESTION 46 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 NSX Overlay D. Management and NSX Overlay

Answer: C

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 47 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

Answer: ABF

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 48 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 DRSE.
E. Compatibility with Windows and Linux kernels

Answer: CDE
Explanation: vSphere Pods are only supported on Supervisor Clusters that use NSX-T Data Center as their networking stack.

Resource Management. vSphere DRS handles the placement of vSphere Pods on the Supervisor Cluster.

QUESTION 49 Which requirement is valid for vSphere with Tanzu on vSphere Distributed Switch Network?

A. Workload networks that are routable to the primary workload network
B. HAProxy Virtual Server IP range that is allocated to NSX-T edge router external interface
C. Network Interface Cards with Single Root IO Visualization Support (SR-IOV)
D. HAProxy Frontend interface that has a common subnet and bridged interface to workload networks

Answer: B
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.

Resources From: 1. 2022 Latest Braindump2go 5V0-23.20 Exam Dumps (PDF & VCE) Free Share:
<https://www.braindump2go.com/5v0-23-20.html> 2. 2022 Latest Braindump2go 5V0-23.20 PDF and 5V0-23.20 VCE Dumps Free Share:
<https://drive.google.com/drive/folders/1tYEvMOleEzNwIyKGl63OyULAy4M3vcOQ?usp=sharing> 3. 2021 Free Braindump2go 5V0-23.20 Exam Questions Download:
[https://www.braindump2go.com/free-online-pdf/5V0-23.20-PDF-Dumps\(1-20\).pdf](https://www.braindump2go.com/free-online-pdf/5V0-23.20-PDF-Dumps(1-20).pdf)
[https://www.braindump2go.com/free-online-pdf/5V0-23.20-VCE-Dumps\(21-41\).pdf](https://www.braindump2go.com/free-online-pdf/5V0-23.20-VCE-Dumps(21-41).pdf) Free Resources from Braindump2go, We Devoted to Helping You 100% Pass All Exams!