


VMware 5V0-23.20 Dumps Questions [2023 Pass for 5V0-23.20 Exam [Q31-Q53]



VMware 5V0-23.20 Dumps Questions [2023] Pass for 5V0-23.20 Exam
Updated VMware Study Guide 5V0-23.20 Dumps Questions

What are the prerequisites for the VMware 5V0-23.20 Exam

At least 12 months of experience with vSphere is required for the minimally qualified candidate.

The MQC has experience with NSX-T, network, virtualization, and OS. Candidates need to have completed all recommended training courses in order to hold the VCP-DCV 2020.

The successful candidate will hold certifications that are recognized by the industry.

NO.31 Where is a storage policy applied to enable Persistent Volumes?

- * Namespace
- * Datastore
- * Virtual Machine
- * Cluster

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.

NO.32 To which set of networks are the Supervisor Cluster nodes attached when deploying with an NSX-T network topology?

- * Frontend and Workload
- * Frontend and Management
- * Workload and NSX Overlay
- * Management and NSX Overlay

NO.33 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?

- * `kubectl describe tkc dev-cluster`
- * `kubectl explain tkg dev-cluster`
- * `kubectl get version`
- * `kubectl get vm dev-cluster`

NO.34 The application development team is pushing a Kubernetes application into production. It consists of an application server and a database. The team wants to ensure that only the production application server can access the production database.

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

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

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.

NO.35 How can you remove unreferenced container images from a project in an embedded Registry Service?

- * Delete images in Content Library.
- * Use `kubectl` to delete the images.
- * Delete the namespace using the vSphere Client.
- * Purge a namespace using the vSphere Client.

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 parent 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.

NO.36 Which functionality does the Cloud Native Storage (CNS) component take advantage of to support the creation of container volumes?

- * First Class Disk
- * VMware Disk Encryption
- * Virtual Disk
- * Storage Based Policy Management

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.

NO.37 How is the storage selected for the Harbor pods when the embedded Harbor image registry is enabled?

- * vCenter Server automatically selects a local ESXi host datastore.
- * An administrator selects a VM storage policy as part of enablement.
- * vCenter Server automatically chooses a VM storage policy.
- * An administrator selects a specific datastore as part of enablement.

Enabling Harbor Image Registry

The vSphere administrator uses the vSphere Client to enable Harbor. To enable this component, select a cluster, select `Configure > Namespaces > Image Registry`, and click `ENABLE HARBOR`:

- * A VM Storage Policy is required to allocate storage for the Harbor pods.
- * An IP Address, based on the ingress CIDR range, is allocated for the Harbor management interface.
- * After a few minutes, Harbor is deployed and running. 184

NO.38 Why would developers choose to deploy an application as a vSphere Pod instead of a Tanzu Kubernetes cluster?

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

NO.39 Which step in vSphere with Tanzu enablement using the vSphere Distributed Switch process is done prior to using the Workload Management Enablement Wizard?

- * Deploy the load balancer
- * Choose the Kubernetes content library that should be used in the Supervisor Cluster
- * Define the Primary workload network P range
- * Define the Management network interfaces for the Supervisor Cluster

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.

NO.40 What is true of a Tanzu Kubernetes Cluster?

- * It is a Kubernetes cluster deployed by developers using a YAML specification file.
- * It use resources from hosts across vSphere clusters to form a Kubernetes cluster on vSphere.
- * It enables vSphere High Availability and Distributed Resource Scheduler.
- * It removes the Kubernetes API for use by DevOps teams.

NO.41 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?

- * 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
- * Four IP addresses are required, one for each of the Control Plane VMs and one spare for performing rolling cluster upgrades
- * Three IP addresses are required, one for each of the Control Plane VMs
- * 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.

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.

NO.42 Which 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)

- * Whether the resources provided by the virtual machine class type should be reserved on the host
- * The configuration parameters which need to be edited in the cluster
- * The amount of CPU. memory, and storage the virtual machine should have
- * Connectivity between the Tanzu Kubernetes cluster and the Subscribed Content Library
- * The storage classes which need to be made available to the cluster

NO.43 A user needs to identify the namespaces that may be accessed.

Which command will provide the desired output?

- * `kubect1 get storageclasses`
- * `kubect1 config use-context`
- * `kubect1 config get-contexts`
- * `kubect1 get contexts`

A user can have permissions on multiple namespaces. The kubect1 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>

NO.44 Which object helps maintain copies of a vSphere pod?

- * ReplicaSets
- * Network Policies
- * Namespaces
- * Persistent Volume

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.

The diagram illustrates a Replica Set. At the top, a box labeled 'replica set controller pod' contains a purple square with a white 'C'. Three arrows point downwards from this controller pod to three separate boxes, each labeled 'Pod-1', 'Pod-2', and 'Pod-3'. Each pod box contains a purple square with a white 'C'. To the right of the controller pod, text indicates 'desired replicas = 3', 'current replicas = 3', and 'current == desired'.

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/>

NO.45 Why would an organization set up private image registries?

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

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

NO.46 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?

- * `image: /<namespace>/<image name>:latest`
- * `template: <image registry url>/<namespace name>/<image name> : latest`
- * `image: <image registry url>/<namespace name>/<image name>:latest`
- * `template: /<namespace name>/<image name>:latest`

NO.47 How can a vSphere administrator replace the Supervisor Cluster API endpoint certificate?

- * Use the certificate-manager CLI utility to replace the Supervisor Cluster API endpoint certificate.
- * Use the vSphere Client to replace the Workload platform MTG certificate.
- * Use the vSphere Client to replace the NSX Load Balancer certificate.
- * Use kubectl to replace the Supervisor Cluster API endpoint certificate.

NO.48 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?

- * 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
- * Four f3 addresses are required, one for each of the Control Plane VMs and one spare for performing rolling cluster upgrades
- * Three P addresses are required, one for each of the Control Plane VMs

* 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.

NO.49 What is the proper way to delete a Persistent Volume Claim?


- * By using the `kubectl delete persistentvolumeclaim` command
- * By using the `kubectl remove pvc` command
- * Through the SPBM policy engine using the vSphere Client
- * By unmounting the volume from the VM and deleting it from the vSphere datastore

NO.50 On which cluster can a Supervisor Namespace be created?

- * A Tanzu Kubernetes Grid Integrated cluster
- * A vSphere 7 cluster enabled with Workload Management
- * A Tanzu Kubernetes cluster
- * A vSphere 6.7 cluster enabled with Workload Management

A Supervisor Cluster namespace is a combination of resource allocations and permissions set within the Supervisor Cluster. When you create a Supervisor Namespace, you must 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.



NO.51 Which is a valid version change for a Tanzu Kubernetes cluster running Kubernetes version 1.16.7?

- * Upgrade one major version (e.g., 2.0.1)
- * Upgrade two minor versions (e.g., 1.18.0)
- * Downgrade one patch version (e.g., 1.16.5)
- * Upgrade one minor version (e.g., 1.17.0)

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.

NO.52 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?

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

NO.53 An 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?

- * Create a storage policy containing the tagged datastores, and apply it to the vSphere Namespace.
- * Create a storage class containing the tagged datastores, and apply it to the Supervisor Cluster
- * Create a persistent volume claim containing the tagged datastores, and apply it to the vSphere Namespace.
- * Create a storage Policy containing the tagged datastores, and apply it to the Supervisor Cluster.

Achieve Success in Actual 5V0-23.20 Exam 5V0-23.20 Exam Dumps:

<https://www.examcollectionpass.com/VMware/5V0-23.20-practice-exam-dumps.html>