

Using HPE Containers

HP HPE2-N68

Version Demo

Total Demo Questions: 7

Total Premium Questions: 72

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

You need to make a web application running in a Kubernetes cluster available to external HTTP clients by using an Ingress. Which three components are required?

- A. An Ingress controller
- B. An Ingress resource
- C. A Service that selects the application Pods
- D. A ConfigMap containing the application URL
- E. A PersistentVolume for the Ingress controller

ANSWER: A B C

Explanation:

An Ingress resource defines HTTP or HTTPS routing rules, such as the host name and path that should be routed to an application. The backend normally references a Kubernetes Service, which provides a stable endpoint for the application Pods. An Ingress controller is also required because it implements the routing behavior described by the Ingress resource.

Creating an Ingress object alone does not expose an application if no controller is installed to process it. A ConfigMap can provide configuration data, but it is not required as part of the basic Ingress path. See the [Kubernetes Ingress documentation](#).

QUESTION NO: 2

You are preparing a Kubernetes manifest for an application that must obtain its configuration from environment-specific values without changing the container image.

Which two methods can provide values from a ConfigMap to a container?

- A. Reference ConfigMap keys as container environment variables.
- B. Mount the ConfigMap as a volume in the Pod.
- C. Use the ConfigMap as a PersistentVolumeClaim.
- D. Assign the ConfigMap as the Pod's ServiceAccount.
- E. Use the ConfigMap to create a Kubernetes node.

ANSWER: A B

Explanation:

A ConfigMap can provide configuration values to a container as environment variables or as files mounted through a volume. Environment variables are useful for application settings read when the process starts. A volume mount is useful when an application expects a configuration file or can read updated configuration data from a mounted directory.

A ConfigMap is intended for nonconfidential configuration data. Sensitive values, such as passwords and tokens, should be stored in a Secret instead. See the Kubernetes documentation for [ConfigMaps](#) and [configuring Pods with ConfigMaps](#).

QUESTION NO: 3

You need to enable the AI/ML project option for a Kubernetes tenant on HPE Ezmeral

Container Platform. What Is a requirement?

- A. That the Kubernetes cluster for the tenant is running Kubernetes version 1.18.6 or above
- B. That the Ezmeral ML Ops Installation package was run on the platform and cluster before the tenant was added
- C. That the platform has enough available Ezmeral ML Ops licenses to match the core quota on the tenant
- D. That the Kubernetes cluster for the tenant has both Spark operator and Istio enabled on it

ANSWER: C

Explanation:

That the platform has enough available Ezmeral ML Ops licenses to match the core quota on the tenant is required because enabling the AI/ML project capability consumes ML Ops licensing according to the tenant's allocated compute capacity. The platform validates that sufficient unassigned ML Ops license capacity exists for the number of cores defined in that tenant's core quota before the feature can be enabled. This ensures that the tenant's AI/ML workloads are properly entitled and that the platform does not allocate more ML Ops capacity than the organization has licensed.

In practical terms, an administrator should review the tenant's core quota and the available ML Ops license inventory before selecting the AI/ML project option. If the quota is increased later, the corresponding licensing capacity must also be available to preserve the entitlement for that tenant. This licensing check is a platform-level prerequisite tied to resource allocation, rather than merely a Kubernetes-version or optional component setting. HPE describes Ezmeral as its platform for deploying and managing containerized and data-intensive workloads, including AI/ML use cases; see the [HPE Ezmeral software overview](#) and the [HPE Ezmeral Runtime Enterprise documentation](#).

QUESTION NO: 4

A company IT department has strict requirements around data protection and wants you to explain how to use the Max Target and Mm Target options for an HPE Ezmeral Data Fabric volume.

What should you explain?

The Max Target sets the desired RAID level for data in the primary volume. The Mm Target sets the desired RAID level for data in associated mirror and snapshot volumes.

- A. The Max Target sets the desired number of replicas for data; the Min Target sets a replica number at which the fabric still tries to reach the Max Target, but less aggressively.
- B. The Max Target and Min Target work together to set the erasure coding level for data in the volume. C. The Max Target is the total number of blocks and the Min Target is the total minus redundant blocks.
- C. The Max Target sets the desired number of replicas for "hot" (frequently used) data; the Min Target sets the desired number of replicas for "warm" (less frequently used) data.

ANSWER: A

Explanation:

The Max Target sets the desired number of replicas for data; the Min Target sets a replica number at which the fabric still tries to reach the Max Target, but less aggressively. These settings control replication targets for the volume's data, rather than RAID levels or separate classes of data. Max Target is the normal, preferred replica count that HPE Ezmeral Data Fabric seeks to maintain. Min Target establishes the lower replication threshold used to prioritize recovery when replicas are unavailable, such as after a node or disk failure. When the number of available replicas is below the desired maximum but remains at or above the minimum target, the fabric can rebuild toward the maximum at a lower priority. This allows the platform to balance ongoing protection with cluster workload and available capacity. If replication falls below the minimum target, restoring protection receives higher urgency because the data has fewer remaining copies. Administrators can therefore use the two values to define both the intended steady-state resilience and the threshold at which replication recovery becomes more critical. HPE's documentation describes volume-level replication and its administrative controls in the [HPE Ezmeral Data Fabric documentation](#); additional product information is available from [HPE Ezmeral Data Fabric](#).

QUESTION NO: 5

CORRECT TEXT

Your company has developed a JSON file to define a custom KubeDirector application.

What step should you take to add the application to the catalog available to a Kubernetes Tenant?

- A. As a Tenant Member or Admin, use the Web Terminal or HPE Kubectl plug-in to issue the kubectl apply -f command.
- B. As a Site Admin, import the JSON file using the Web UI. Then enable the application within the Tenant settings.
- C. As a Tenant Admin, use Helm to download the JSON file as a chart. Then deploy the chart as a release within the tenant.
- D. As a Site Admin, place the file in the controller /srv/bluedata/catalog directory; modify permissions on the file to make it accessible in the UI.

ANSWER: B

Explanation:

As a Site Admin, import the JSON file using the Web UI. Then enable the application within the Tenant settings. This is the required workflow because a KubeDirector application definition is first added to the platform-level application catalog by a site administrator. The import process registers the custom JSON application template with HPE Ezmeral Runtime Enterprise, making it available for controlled distribution. The site administrator can then explicitly enable that catalog application for the relevant Kubernetes tenant. This two-stage process separates global catalog administration from tenant-level consumption and allows the platform team to decide which approved application templates each tenant may deploy. Once enabled, authorized users in that tenant can select the custom application from its catalog and create deployments using the template's defined services, container images, resource settings, and other application parameters. HPE's Runtime Enterprise documentation describes site-level administration and tenant access controls for Kubernetes resources and application catalogs. See the [HPE Ezmeral Runtime Enterprise documentation](#) and the [HPE KubeDirector project documentation](#).

QUESTION NO: 6 - (SIMULATION)

Your company has developed a JSON file to define a custom KubeDirector application.

What step should you take to add the application to the catalog available to a Kubernetes Tenant?

- B.As a Tenant Member or Admin, use the Web Terminal or HPE Kubectl plug-in to issue the kubectl apply -f command.
- C.As a Site Admin, import the JSON file using the Web UI. Then enable the application within the Tenant settings.
- C.As a Tenant Admin, use Helm to download the JSON file as a chart. Then deploy the chart as a release within the tenant.
- D.As a Site Admin, place the file in the controller /srv/bluedata/catalog directory; modify permissions on the file to make it accessible in the UI.

ANSWER: Check the explanation for the answer

Explanation:

Correct answer: As a **Site Admin**, import the custom KubeDirector application JSON file through the Web UI, then enable the imported application in the appropriate Tenant settings.

KubeDirector application definitions are managed in the platform application catalog by a Site Admin. Importing the JSON registers the application; enabling it for a tenant makes it available to that tenant's users for deployment.

Log in as a Site Admin.

Use the Web UI catalog/application-definition import function to upload the JSON file.

Open the target tenant configuration/settings.

Enable the newly imported application for that tenant.

Commands such as `kubectl apply -f`, Helm chart deployment, or manually copying the JSON to a controller directory are not the catalog-registration procedure.

QUESTION NO: 7

You need to update a Kubernetes Deployment with a new container image while maintaining application availability.

Which two statements correctly describe a rolling update?

- A. `maxUnavailable` controls how many desired Pods can be unavailable during the update.
- B. `maxSurge` controls how many additional Pods can be created above the desired replica count.
- C. A rolling update requires all existing Pods to be deleted before replacement Pods are created.
- D. A Deployment must be converted to a StatefulSet before it can use a rolling update.

ANSWER: A B

Explanation:

A Deployment rolling update replaces Pods gradually rather than removing all existing replicas at the same time. The `maxUnavailable` setting controls how many desired Pods can be unavailable during the update, while `maxSurge` controls how many additional Pods can be created above the desired replica count.

This behavior supports controlled application updates and allows Kubernetes to monitor readiness as replacement Pods are created. A rolling update does not require deleting the Deployment or changing it to a StatefulSet. See the [Kubernetes Deployment documentation](#).