

CKAD Reliable Test Labs & Valid Exam CKAD Book

reetasingh/
CKAD_labs



Includes labs I have done as part of preparation for CKAD exam

1 Contributor 0 Issues 1 Star 4 Forks

What's more, part of that Prep4sureGuide CKAD dumps now are free: https://drive.google.com/open?id=1DZZCK5cIYdPSIPxGOkHu9sOX_6NYNjIy

Your life will take place great changes after obtaining the CKAD certificate. Many companies like to employ versatile and comprehensive talents. What you have learnt on our CKAD preparation prep will meet their requirements. So you will finally stand out from a group of candidates and get the desirable job. At the same time, what you have learned from our CKAD Exam Questions are the latest information in the field, so that you can obtain more skills to enhance your capacity.

Linux Foundation Certified Kubernetes Application Developer (CKAD) exam is a certification program designed to evaluate and certify the skills and knowledge of developers in using Kubernetes to develop cloud-native applications. The CKAD Certification is recognized worldwide as a benchmark for Kubernetes application development expertise, and it validates the ability of developers to create and deploy cloud-native applications using Kubernetes.

>> **CKAD Reliable Test Labs** <<

Quiz 2026 Linux Foundation Newest CKAD: Linux Foundation Certified Kubernetes Application Developer Exam Reliable Test Labs

The client can try out and download our Linux Foundation CKAD Training Materials freely before their purchase so as to have an understanding of our product and then decide whether to buy them or not. The website pages of our product provide the details of our Linux Foundation Certified Kubernetes Application Developer Exam learning questions.

Linux Foundation Certified Kubernetes Application Developer Exam Sample Questions (Q149-Q154):

NEW QUESTION # 149



Given a container that writes a log file in format A and a container that converts log files from format A to format B, create a deployment that runs both containers such that the log files from the first container are converted by the second container, emitting logs in format B.

Task:

* Create a deployment named deployment-xyz in the default namespace, that:

*Includes a primary

lfcncf/busybox:1 container, named logger-dev

*includes a sidecar lfcncf/fluentd:v0.12 container, named adapter-zen

*Mounts a shared volume /tmp/log on both containers, which does not persist when the pod is deleted

*Instructs the logger-dev container to run the command

```
while true; do
  echo "i luv cncf" >> /
  tmp/log/input.log;
  sleep 10;
done
```

which should output logs to /tmp/log/input.log in plain text format, with example values:

```
i luv cncf
i luv cncf
i luv cncf
```

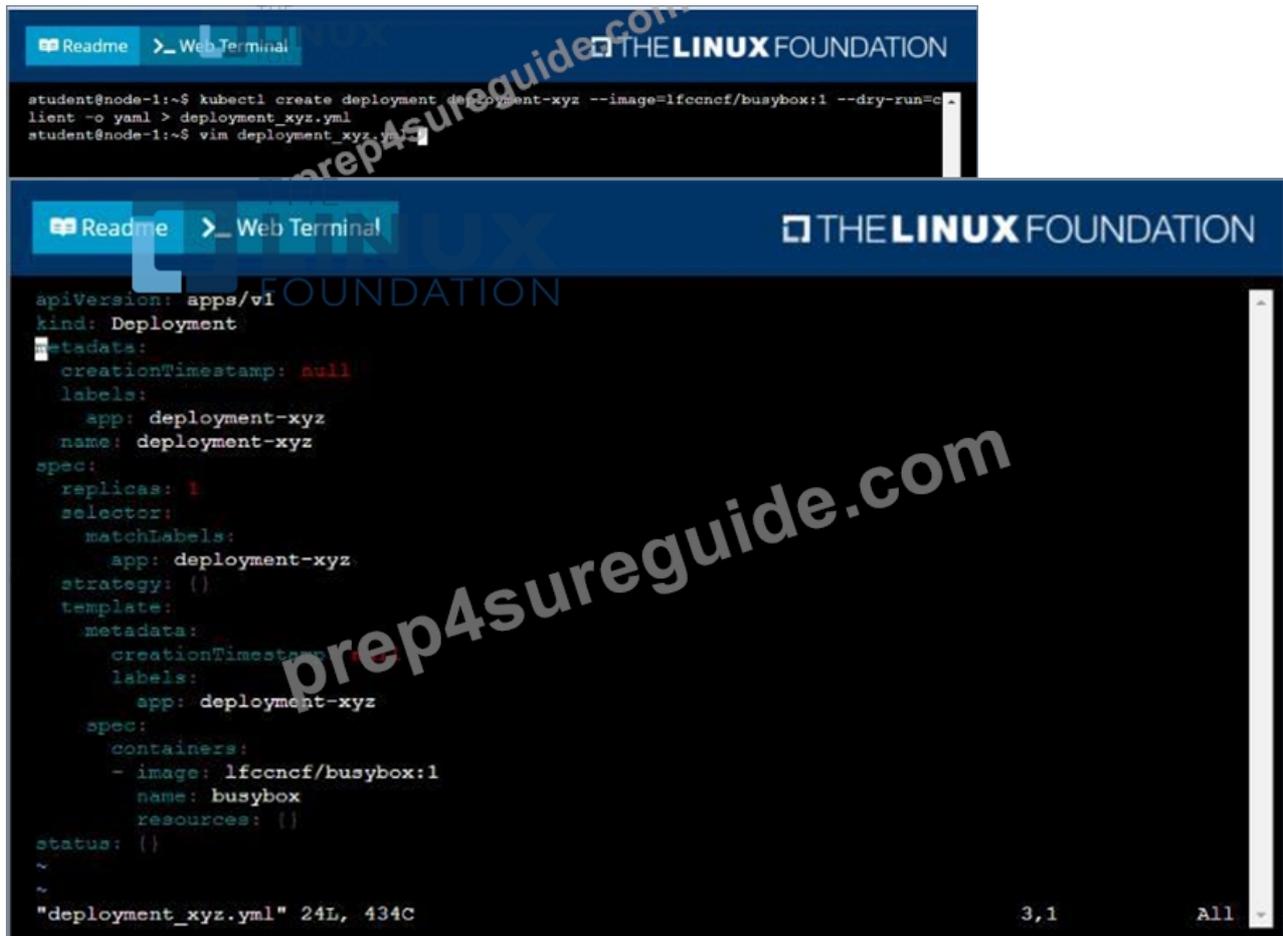
* The adapter-zen sidecar container should read /tmp/log/input.log and output the data to /tmp/log/output.* in Fluentd JSON format. Note that no knowledge of Fluentd is required to complete this task: all you will need to achieve this is to create the ConfigMap from the spec file provided at /opt/KDMC00102/fluentd-configmap.p.yaml, and mount that ConfigMap to /fluentd/etc in the adapter-zen sidecar container See the solution below.

Answer:

Explanation:

Explanation

Solution:



The screenshot shows a terminal window with a dark background and light text. The terminal prompt is 'student@node-1:~\$'. The user enters the command 'kubectl create deployment deployment-xyz --image=lfcncf/busybox:1 --dry-run=client -o yaml > deployment_xyz.yml'. The prompt changes to 'student@node-1:~\$ vim deployment_xyz.yml'. The terminal then displays the contents of the 'deployment_xyz.yml' file, which is a Kubernetes Deployment manifest. The manifest includes metadata (creationTimestamp: null, labels: app: deployment-xyz, name: deployment-xyz), spec (replicas: 1, selector: matchLabels: app: deployment-xyz, strategy: {}, template: metadata: creationTimestamp: null, labels: app: deployment-xyz, spec: containers: - image: lfcncf/busybox:1, name: busybox, resources: {}), and status: {}.

```
student@node-1:~$ kubectl create deployment deployment-xyz --image=lfcncf/busybox:1 --dry-run=client -o yaml > deployment_xyz.yml
student@node-1:~$ vim deployment_xyz.yml

apiVersion: apps/v1
kind: Deployment
metadata:
  creationTimestamp: null
  labels:
    app: deployment-xyz
  name: deployment-xyz
spec:
  replicas: 1
  selector:
    matchLabels:
      app: deployment-xyz
  strategy: {}
  template:
    metadata:
      creationTimestamp: null
      labels:
        app: deployment-xyz
    spec:
      containers:
      - image: lfcncf/busybox:1
        name: busybox
        resources: {}
status: {}
~
~
"deployment_xyz.yml" 24L, 434C                                     3,1 All
```



```

kind: Deployment
metadata:
  labels:
    app: deployment-xyz
    name: deployment-xyz
spec:
  replicas: 1
  selector:
    matchLabels:
      app: deployment-xyz
  template:
    metadata:
      labels:
        app: deployment-xyz
    spec:
      volumes:
        - name: myvoll
          emptyDir: {}
      containers:
        - image: lfcncf/busybox:1
          name: logger-dev
          volumeMounts:
            - name: myvoll
              mountPath: /tmp/log
        - image: lfcncf/fluentd:v0.12
          name: adapter-zen

```

3 lines yanked

27,22

Bot



```

replicas: 1
selector:
  matchLabels:
    app: deployment-xyz
template:
  metadata:
    labels:
      app: deployment-xyz
  spec:
    volumes:
      - name: myvoll
        emptyDir: {}
    containers:
      - image: lfcncf/busybox:1
        name: logger-dev
        command: ["/bin/sh", "-c", "while [ true ]; do echo 'i lue cnof' >> /tmp/log/input.log; sl
sleep 10; done"]
        volumeMounts:
          - name: myvoll
            mountPath: /tmp/log
      - image: lfcncf/fluentd:v0.12
        name: adapter-zen
        command: ["/bin/sh", "-c", "tail -f /tmp/log/input.log >> /tmp/log/output.log"]
        volumeMounts:
          - name: myvoll
            mountPath: /tmp/log

```

29,83

Bot

The screenshot shows a web terminal interface with a blue header containing 'Readme' and 'Web Terminal' buttons, and 'THE LINUX FOUNDATION' logo. The terminal content is as follows:

```

metadata:
  labels:
    app: deployment-xyz
spec:
  volumes:
  - name: myvol1
    emptyDir: {}
  - name: myvol2
    configMap:
      name: logconf
  containers:
  - image: lfccncf/busybox:1
    name: logger-dev
    command: ["/bin/sh", "-c", "while true; do echo 'i luv cncf' >> /tmp/log/input.log; sleep 10; done"]
    volumeMounts:
    - name: myvol1
      mountPath: /tmp/log
  - image: lfccncf/fluentd:v0.12
    name: adapter-zen
    command: ["/bin/sh", "-c", "tail -f /tmp/log/input.log >> /tmp/log/output.log"]
    volumeMounts:
    - name: myvol1
      mountPath: /tmp/log
    - name: myvol2
      mountPath: /fluentd/etc

```

Below the configuration, the terminal shows the execution of the following commands and their outputs:

```

student@node-1:~$ kubectl create -f deployment_xyz.yml
deployment.apps/deployment-xyz created
student@node-1:~$ kubectl get deployment
NAME          READY   UP-TO-DATE   AVAILABLE   AGE
deployment-xyz 0/1     1             0           3s
student@node-1:~$ kubectl get deployment
NAME          READY   UP-TO-DATE   AVAILABLE   AGE
deployment-xyz 0/1     1             0           9s
student@node-1:~$ kubectl get deployment
NAME          READY   UP-TO-DATE   AVAILABLE   AGE
deployment-xyz 1/1     1             1          12s
student@node-1:~$

```

NEW QUESTION # 150

You have a custom resource definition (CRD) named that represents a database resource in your Kubernetes cluster. You want to create a custom operator that automates the creation and management of these database instances. The operator should handle the following:

- Creation: When a new 'database.example.com' resource is created, the operator should provision a new PostgreSQL database instance on the cluster-
- Deletion: When a 'database.example_com' resource is deleted, the operator should clean up the corresponding PostgreSQL database instance.
- Scaling: If the 'spec-replicas' field of the 'database-example.com' resource is updated, the operator should scale the number of database instances accordingly.

Provide the necessary Kubernetes resources, custom operator code, and steps to implement this operator. You should use the 'Operator Framework' to build and deploy this operator

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create the CRD:

```

apiVersion: apiextensions.k8s.io/v1
kind: CustomResourceDefinition
metadata:
  name: databases.example.com
spec:
  group: example.com
  names:
    kind: Database
    plural: databases
    singular: database
  scope: Namespaced
  versions:
  - name: v1
    served: true
    storage: true
  subresources:
    status: {}

```

- Apply this YAML file to your cluster using 'kubectl apply -f database-crd.yaml'. 2. Create the Operator Project: - Use the Operator Framework' to initialize a new operator project `bash operator-sdk init -domain example.com -repo example.com/database-operator --version VO.O. I -license apache2` - Replace 'example.com' with your desired domain name. 3. Define the Custom Resource: - Create a 'database_types.go' file in the 'api/v1' directory of your project. - Define the 'Database' resource as a custom resource struct `Go package v1 import (metav1 "k8s.io/api/machinery/pkg/apis/meta/v1" // DatabaseSpec defines the desired state of Database type DatabaseSpec struct { If Replicas specifies the number of database instances to run.`

// Password is the password for the database users.

```

Password string `json:"password"`

```

} // DatabaseStatus defines the observed state of Database type DatabaseStatus struct { // Replicas is the actual number of database instances running.

```

Replicas int32 `json:"replicas"`

```

// Ready indicates if the database is ready to accept connections.

```

Ready bool `json:"ready"`

```

```

}
// +kubebuilder:object:root=true
// +kubebuilder:subresource:status

```

```

// Database is the Schema for the databases API
type Database struct {
  metav1.TypeMeta `json:"inline"`
  metav1.ObjectMeta `json:"metadata,omitempty"`

```

```

  Spec DatabaseSpec `json:"spec,omitempty"`
  Status DatabaseStatus `json:"status,omitempty"`
}

```

```

// +kubebuilder:object:root=true

```

```

// DatabaseList contains a list of Database
type DatabaseList struct {
  metav1.TypeMeta `json:"inline"`
  metav1.ListMeta `json:"metadata,omitempty"`
  Items []Database `json:"items"`
}

```

```

func init() {
  SchemeBuilder.Register(&Database{}, &DatabaseList{})
}

```

4. Implement the Controller Logic: - Create a 'database_controller.go' file in the 'controllers' directory- - Implement the logic for creating, deleting, and scaling database instances.

```

go
package controllers

import (
  "context"
  "fmt"
  appsv1 "k8s.io/api/apps/v1"
  corev1 "k8s.io/api/core/v1"
  "k8s.io/apimachinery/pkg/api/errors"

```

```

    metav1 "k8s.io/apimachinery/pkg/apis/meta/v1"
    "k8s.io/apimachinery/pkg/runtime"
    "k8s.io/apimachinery/pkg/types"
    "sigs.k8s.io/controller-runtime/pkg/client"
    "sigs.k8s.io/controller-runtime/pkg/controller"
    "sigs.k8s.io/controller-runtime/pkg/handler"
    "sigs.k8s.io/controller-runtime/pkg/manager"
    "sigs.k8s.io/controller-runtime/pkg/reconcile"
    "sigs.k8s.io/controller-runtime/pkg/source"
    "sigs.k8s.io/controller-runtime/pkg/webhook/admission"
)

// DatabaseReconciler reconciles a Database object
type DatabaseReconciler struct {
    client.Client
    Scheme runtime.Scheme
}

// +kubebuilder:rbac:groups=example.com,resources=databases,verbs=get;list;watch;create;update;patch;delete
// +kubebuilder:rbac:groups=example.com,resources=databases/status,verbs=get;update;patch
// +kubebuilder:rbac:groups=apps,resources=deployments,verbs=get;list;watch;create;update;patch;delete
// +kubebuilder:rbac:groups=core,resources=services,verbs=get;list;watch;create;update;patch;delete

func (r DatabaseReconciler) Reconcile(ctx context.Context, req reconcile.Request) (reconcile.Result, error) {
    log := r.Log.WithValues("database", req.NamespacedName)

    // Fetch the Database instance
    instance := &v1.Database{}
    err := r.Get(ctx, req.NamespacedName, instance)
    if err != nil {
        if errors.IsNotFound(err) {
            // Request object not found, could have been deleted after reconcile request.
            // Owned objects are automatically garbage collected. For additional cleanup logic use finalizers.
            // Return and don't requeue
            return reconcile.Result{}, nil
        }
        // Error reading the object - requeue the request
        return reconcile.Result{}, err
    }

    // Check if the number of replicas needs to be updated
    if instance.Spec.Replicas != instance.Status.Replicas {
        // Scale the deployment to match the desired number of replicas
        err = r.scaleDeployment(ctx, instance)
        if err != nil {
            return reconcile.Result{}, err
        }
        instance.Status.Replicas = instance.Spec.Replicas
    }

    // Set the status of the database instance
    instance.Status.Ready = true

    // Update the database instance status
    err = r.Status().Update(ctx, instance)
    if err != nil {
        return reconcile.Result{}, err
    }

    return reconcile.Result{}, nil
}

func (r DatabaseReconciler) scaleDeployment(ctx context.Context, instance v1.Database) error {
    // Create or update the Deployment
    deployment := &apps1.Deployment{
        ObjectMeta: metav1.ObjectMeta{
            Name: fmt.Sprintf("database-%s", instance.Name),
            Namespace: instance.Namespace,
        },
        Spec: apps1.DeploymentSpec{
            Replicas: &instance.Spec.Replicas,
            Selector: &metav1.LabelSelector{
                MatchLabels: map[string]string{
                    "app": "database",
                },
            },
            Template: corev1.PodTemplateSpec{
                ObjectMeta: metav1.ObjectMeta{
                    Labels: map[string]string{
                        "app": "database",
                    },
                },
                Spec: corev1.PodSpec{
                    Containers: []corev1.Container{
                        {
                            Name: "database",
                            Image: "postgres:latest", // Use your desired PostgreSQL image
                            Env: []corev1.EnvVar{
                                {

```

```

Name: "POSTGRES_USER",
Value: "postgres",
},
{
Name: "POSTGRES_PASSWORD",
Value: instance.Spec.Password,
},
{
Name: "POSTGRES_DB",
Value: instance.Name,
},
},
},
Ports: []corev1.ContainerPort{
{
ContainerPort: 5432,
Name: "database",
},
},
},
},
},
},
},
},
},
},
},
}

// Create or update the deployment
if err := r.CreateOrUpdate(ctx, deployment, func() error {
return r.Client.Get(ctx, types.NamespacedName{Name: deployment.Name, Namespace: deployment.Namespace}, deployment)
}); err != nil {
return err
}

// Create or update the service
service := &corev1.Service{
ObjectMeta: metav1.ObjectMeta{
Name: fmt.Sprintf("database-%s", instance.Name),
Namespace: instance.Namespace,
},
Spec: corev1.ServiceSpec{
Selector: map[string]string{
"app": "database",
},
Ports: []corev1.ServicePort{
{
Port: 5432,
TargetPort: intstr.FromInt(5432),
Name: "database",
},
},
},
},
},
}

// Create or update the service
if err := r.CreateOrUpdate(ctx, service, func() error {
return r.Client.Get(ctx, types.NamespacedName{Name: service.Name, Namespace: service.Namespace}, service)
}); err != nil {
return err
}

return nil
}

func (r DatabaseReconciler) SetupWithManager(mgr manager.Manager) error {
return ctrl.NewControllerManagedBy(mgr).
For(&v1.Database{}).
Owns(&apps1.Deployment{}).
Owns(&corev1.Service{}).
Complete(r)
}

```



5. Build and Deploy the Operator: - Build the operator using the 'operator-sdk build' command: `bash operator-sdk build example.com/database-operator:v0.0.1 --local` - Deploy the operator to your Kubernetes cluster: `bash kubectl apply -f deploy/operator.yaml` 6. Test the Operator: - Create a new 'database-example-conf' resource:

```

apiVersion: example.com/v1
kind: Database
metadata:
  name: my-database
  namespace: default
spec:
  replicas: 1
  password: "mypassword"

```

- Apply the YAML file to your cluster: `bash kubectl apply -f my-database.yaml` - Verify that the operator creates a PostgreSQL database instance. - Test scaling the database by updating the 'spec.replicas' field of the 'database.example.com' resource. - Delete the 'database.example.com' resource and verify that the operator cleans up the database instance. This step-by-step guide demonstrates a basic example of a custom operator using the Operator Framework. You can customize this operator further to handle more complex operations and integrate with other Kubernetes resources. ,

NEW QUESTION # 151

Context

Anytime a team needs to run a container on Kubernetes they will need to define a pod within which to run the container.

Task

Please complete the following:

* Create a YAML formatted pod manifest

/opt/KDPD00101/pod1.yml to create a pod named app1 that runs a container named app1cont using image lfcncf/arg-output with these command line arguments: `-lines 56 -F`

* Create the pod with the kubectl command using the YAML file created in the previous step

* When the pod is running display summary data about the pod in JSON format using the kubectl command and redirect the output to a file named /opt/KDPD00101/out1.json

* All of the files you need to work with have been created, empty, for your convenience

When creating your pod, you do not need to specify a container command,

only args

Answer:

Explanation:

See the solution below.

Explanation

Solution:

```

student@node-1:~$ kubectl run app1 --image=lfcncf/arg-output --dry-run=client -o yaml > /opt/KD
PD00101/pod1.yml
student@node-1:~$ vim /opt/KDPD00101/pod1.yml

```

```
apiVersion: v1
kind: Pod
metadata:
  creationTimestamp: null
  labels:
    run: app1
  name: app1
spec:
  containers:
  - image: lfcncf/arg-output
    name: app1
    resources: {}
  dnsPolicy: ClusterFirst
  restartPolicy: Always
status: {}
```

"/opt/KDPD00101/pod1.yml" 15L, 242C

3,1

All

```
apiVersion: v1
kind: Pod
metadata:
  labels:
    run: app1
  name: app1
spec:
  containers:
  - image: lfcncf/arg-output
    name: app1
    args: ["--lines", "56", "--s"]
```

11,30

All

```

pod/app1 created
student@node-1:~$ kubectl get pods
NAME          READY   STATUS             RESTARTS   AGE
app1          0/1     ContainerCreating  0           5s
counter       1/1     Running            0           4m44s
liveness-http 1/1     Running            0           6h50m
nginx-101     1/1     Running            0           6h51m
nginx-configmap 1/1     Running            0           6m21s
nginx-secret  1/1     Running            0           11m
poller        1/1     Running            0           6h51m
student@node-1:~$ kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
app1          1/1     Running  0           26s
counter       1/1     Running  0           5m5s
liveness-http 1/1     Running  0           6h50m
nginx-101     1/1     Running  0           6h51m
nginx-configmap 1/1     Running  0           6m42s
nginx-secret  1/1     Running  0           12m
poller        1/1     Running  0           6h51m
student@node-1:~$ kubectl delete pod app1
pod "app1" deleted
student@node-1:~$ vim /opt/KDPD00101/pod1.yml

```

Readme > Web Terminal

```

nginx-configmap 1/1     Running  0           6m21s
nginx-secret    1/1     Running  0           11m
poller          1/1     Running  0           6h51m
student@node-1:~$ kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
app1          1/1     Running  0           26s
counter       1/1     Running  0           5m5s
liveness-http 1/1     Running  0           6h50m
nginx-101     1/1     Running  0           6h51m
nginx-configmap 1/1     Running  0           6m42s
nginx-secret  1/1     Running  0           12m
poller        1/1     Running  0           6h51m
student@node-1:~$ kubectl delete pod app1
pod "app1" deleted
student@node-1:~$ vim /opt/KDPD00101/pod1.yml
student@node-1:~$ kubectl create -f /opt/KDPD00101/pod1.yml
pod/app1 created
student@node-1:~$ kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
app1          1/1     Running  0           20s
counter       1/1     Running  0           6m57s
liveness-http 1/1     Running  0           6h52m
nginx-101     1/1     Running  0           6h53m
nginx-configmap 1/1     Running  0           8m34s
nginx-secret  1/1     Running  0           14m
poller        1/1     Running  0           6h53m
student@node-1:~$ kubectl get pod app1 -o json >

```

```

Readme  Web Terminal  THE LINUX FOUNDATION

poller          1/1      Running          0          6h51m
student@node-1:~$ kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
appl          1/1     Running    0           26s
counter       1/1     Running    0           5m5s
liveness-http 1/1     Running    0           6h50m
nginx-101     1/1     Running    0           6h51m
nginx-configmap 1/1     Running    0           6m42s
nginx-secret  1/1     Running    0           12m
poller        1/1     Running    0           6h51m
student@node-1:~$ kubectl delete pod appl
pod "appl" deleted
student@node-1:~$ vim /opt/KDPD00101/pod1.yml
student@node-1:~$ kubectl create -f /opt/KDPD00101/pod1.yml
pod/appl created
student@node-1:~$ kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
appl          1/1     Running    0           20s
counter       1/1     Running    0           6m57s
liveness-http 1/1     Running    0           6h52m
nginx-101     1/1     Running    0           6h53m
nginx-configmap 1/1     Running    0           8m34s
nginx-secret  1/1     Running    0           14m
poller        1/1     Running    0           6h53m
student@node-1:~$ kubectl get pod appl -o json > /opt/KDPD00101/out1.json
student@node-1:~$
student@node-1:~$

```

NEW QUESTION # 152

You are managing a Kubernetes cluster running a highly-available application that uses a custom resource called 'Orders'. The 'orders' resource is created and managed by a custom controller that ensures the order processing workflow runs flawlessly. However, the 'order' resource's validation rules have changed, requiring a new schema to be applied. How can you ensure that the existing 'Order' resources conform to the new schema without disrupting the application's functionality?

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1). Define the New Schema:

- Create a new CustomResourceDefinition (CRD) file with the updated schema for the 'Order' resource.
- Ensure that the CRD's 'spec-validation.openAPIV3Schema' field includes all the new validation rules.

```

apiVersion: apiextensions.k8s.io/v1
kind: CustomResourceDefinition
metadata:
  name: orders.example.com
spec:
  group: example.com
  version: v1
  names:
    kind: Order
    plural: orders
  scope: Namespaced
  validation:
    openAPIV3Schema:
      type: object
      properties:
        # Updated validation rules for the 'Order' resource
      ...

```

2. Update the CRD: - Apply the new CRD definition using 'kubectl apply -f order-crd.yaml'.
3. Create a Webhook for Validation: - Define a webhook in your Kubernetes cluster that will be responsible for validating the 'order' resources against the new schema. - Configure the webhook to be invoked during resource creation and update operations.

```

apiVersion: admissionregistration.k8s.io/v1
kind: ValidatingWebhookConfiguration
metadata:
  name: order-validation-webhook
webhooks:
- name: order-validation
  rules:
  - apiGroups: ["example.com"]
    apiVersions: ["v1"]
    resources: ["orders"]
    operations: ["CREATE", "UPDATE"]
  failurePolicy: Fail
  admissionReviewVersions: ["v1", "v1beta1"]
  clientConfig:
    service:
      name: order-validation-service
      namespace: default
      path: /validate

```

4. Deploy the Validation Service: - Create a deployment for the validation service that implements the logic for validating the 'Order' resources against the new schema. - The service should expose an endpoint that the webhook can communicate with.

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: order-validation-service
spec:
  replicas: 1
  selector:
    matchLabels:
      app: order-validation
  template:
    metadata:
      labels:
        app: order-validation
    spec:
      containers:
      - name: order-validation
        image: your-validation-image:latest
        ports:
        - containerPort: 8443

```

5. Reconcile Existing Resources: - Once the validation webhook and service are deployed, create a temporary job that iterates through all existing 'Order' resources. - The job should validate each resource against the new schema and automatically update any resources that do not comply.

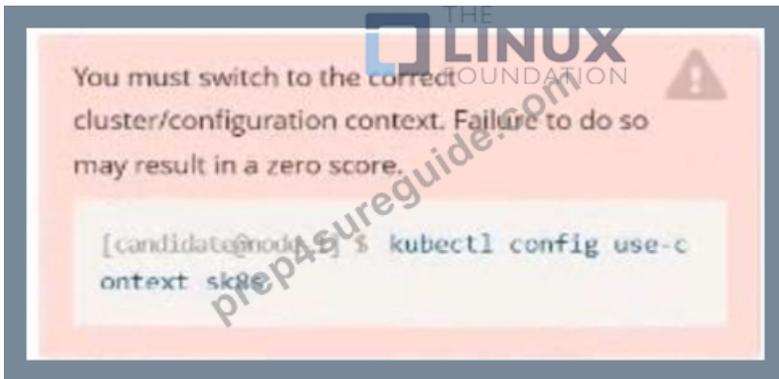
```

apiVersion: batch/v1
kind: Job
metadata:
  name: order-reconciliation-job
spec:
  template:
    spec:
      containers:
      - name: order-reconciliation
        image: your-reconciliation-image:latest
        command: ["bin/sh", "-c", "kubectl get orders --all-namespaces -o jsonpath='{.items[*].metadata.name}' | xargs -n1 kubectl get -f - -o jsonpath='{.spec}' | jq -r '. | . + { \"metadata\": { \"annotations\": { \"order.example.com/schema-version\": \"2\" } } }' | kubectl apply -f -"]

```

By following these steps, you can ensure that your 'order' resources conform to the new schema without disrupting the application's functionality. The validation webhook prevents new invalid resources from being created, and the reconciliation job ensures that existing resources are updated to meet the new schema requirements. This approach allows for smooth schema evolution and maintains the consistency of your data.,

NEW QUESTION # 153



Task:

- 1) Create a secret named app-secret in the default namespace containing the following single key-value pair:
Key3: value1
- 2) Create a Pod named nginx secret in the default namespace. Specify a single container using the nginx:stable image. Add an environment variable named BEST_VARIABLE consuming the value of the secret key3.

Answer:

Explanation:

See the solution below.

Explanation

Solution:

```
candidate@node-1:~$ kubectl config use-context k8s  
Switched to context "k8s".  
candidate@node-1:~$ kubectl create secret generic app-secret -n default --from-literal=key3=value1  
secret/app-secret created  
candidate@node-1:~$ kubectl get secrets  
NAME          TYPE          DATA   AGE  
app-secret    Opaque        1       4s  
candidate@node-1:~$ kubectl run nginx-secret -n default --image=nginx:stable --dry-run=client -o yaml > sec.yaml  
candidate@node-1:~$ vim sec.yaml
```

Text Description automatically generated

```
File Edit View Terminal Tabs Help  
apiVersion: v1  
kind: Pod  
metadata:  
  creationTimestamp: null  
  labels:  
    run: nginx-secret  
  name: nginx-secret  
  namespace: default  
spec:  
  containers:  
  - image: nginx:stable  
    name: nginx-secret  
    env:  
    - name: BEST_VARIABLE  
      valueFrom:  
        secretKeyRef:  
          name: app-secret  
          key: key3
```

Text Description automatically generated

