

Linux Foundation Certified Kubernetes Application Developer Exam Online Questions - Outstanding Practice To your CKAD Exam



P.S. Free & New CKAD dumps are available on Google Drive shared by ExamDiscuss: <https://drive.google.com/open?id=1ZW4tMKAsmcAqr4umwDkI48EBDjvxgUo7>

As is known to us, getting the newest information is very important for all people to pass the exam and get the certification in the shortest time. In order to help all customers gain the newest information about the CKAD exam, the experts and professors from our company designed the best CKAD test guide. The experts will update the system every day. If there is new information about the exam, you will receive an email about the newest information about the CKAD Learning Materials. We can promise that you will never miss the important information about the CKAD exam.

The benefit of Obtaining the CNCF Certified Kubernetes Application Developer

Those who pass the CNCF Certified Kubernetes Application Developer Exam with the help of **CNCF CKAD Dumps** gain several benefits:

- The CNCF Certified Kubernetes Application Developer allows you to work on the projects that you create.
- You might be eligible for a higher salary.
- It will help employers assess how qualified candidates are for their job requirements.
- This certification will provide you with a good foundation for the Oracle Certified Associate Cloud + Container Administrator exam.
- CNCF Certified Kubernetes Application Developer tells employers that you have the knowledge they need to do the job.

Linux Foundation Certified Kubernetes Application Developer (CKAD) certification exam is a globally recognized certification program that validates the skills of Kubernetes application developers. The CKAD Certification Exam is designed to test the proficiency of developers in creating and deploying applications on Kubernetes clusters. It is one of the most sought-after certifications in the world of containerization and cloud-native computing.

Preparing for the CKAD certification exam requires a significant amount of time and effort. Candidates must have a deep understanding of Kubernetes concepts and be able to perform tasks quickly and accurately using the command line. However, achieving CKAD certification can open many doors for developers, including new job opportunities, higher salaries, and increased credibility in the industry.

>> CKAD Test Answers <<

New CKAD Dumps Free, Relevant CKAD Exam Dumps

With the rise of internet and the advent of knowledge age, mastering knowledge about computer is of great importance. This CKAD exam is your excellent chance to master more useful knowledge of it. Up to now, No one has questioned the quality of our CKAD training materials, for their passing rate has reached up to 98 to 100 percent. Our Kubernetes Application Developer study dumps are priced reasonably so we made a balance between delivering satisfaction to customers and doing our own jobs. So in this critical moment, our CKAD real materials will make you satisfied. Our CKAD exam materials can provide integrated functions. You can learn a great deal of knowledge and get the certificate of the exam at one order like win-win outcome at one try.

Linux Foundation Certified Kubernetes Application Developer Exam Sample Questions (Q72-Q77):

NEW QUESTION # 72

You're building a Kubernetes application that manages user profiles and requires a custom resource for storing profile information. Design a custom resource definition (CRD) and its corresponding controller, ensuring that every time a profile is created or updated, a unique user ID is assigned to the profile.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create the Custom Resource Definition (CRD)

- Define the CRD Spec:

```
apiVersion: apiextensions.k8s.io/v1
kind: CustomResourceDefinition
metadata:
  name: profiles.example.com
spec:
  group: example.com
  versions:
    - name: v1
      served: true
      storage: true
      schema:
        openAPIV3Schema:
          type: object
          properties:
            metadata:
              type: object
            spec:
              type: object
              properties:
                name:
                  type: string
                  description: User name
                email:
                  type: string
                  description: User email
                userId:
                  type: string
                  description: Unique user ID
```

scope: Namespaced

- Apply the CRD: `bash kubectl apply -f profile-crd.yaml` 2. Create a Controller for the Custom Resource: - Define the Controller Logic:

```

package main

import (
    "context"
    "fmt"
    "time"

    metav1 "k8s.io/apimachinery/pkg/apis/meta/v1"
    "k8s.io/apimachinery/pkg/runtime"
    "k8s.io/apimachinery/pkg/types"
    "k8s.io/client-go/kubernetes/scheme"
    ctrl "sigs.k8s.io/controller-runtime"
    "sigs.k8s.io/controller-runtime/pkg/client"
    "sigs.k8s.io/controller-runtime/pkg/log"

    examplev1 "github.com/your-org/example-operator/api/v1"
)

// ProfileReconciler reconciles a Profile object
type ProfileReconciler struct {
    client.Client
    Scheme runtime.Scheme
}

//+kubebuilder:rbac:groups=example.com,resources=profiles,verbs=get;list;watch;create;update;patch;delete
//+kubebuilder:rbac:groups=example.com,resources=profiles/status,verbs=get;update;patch
//+kubebuilder:rbac:groups=example.com,resources=profiles/finalizers,verbs=update

func (r ProfileReconciler) Reconcile(ctx context.Context, req ctrl.Request) (ctrl.Result, error) {
    _ = log.FromContext(ctx)

    // Fetch the Profile instance
    profile := &examplev1.Profile{}
    err := r.Get(ctx, req.NamespacedName, profile)
    if err != nil {
        return ctrl.Result{}, client.IgnoreNotFound(err)
    }

    // Generate a unique user ID (replace with your logic)
    userID := generateUniqueUserID(profile)

    // Update the Profile with the user ID
    profile.Spec.UserID = userID

    // Update the Profile resource
    err = r.Update(ctx, profile)
    if err != nil {
        return ctrl.Result{}, err
    }

    return ctrl.Result{RequeueAfter: time.Second * 5}, nil
}

// Generate a unique user ID (replace with your logic)
func generateUniqueUserID(profile examplev1.Profile) string {
    // You can use a combination of namespace, name, and a random ID
    return fmt.Sprintf("%s-%s-%d", profile.Namespace, profile.Name, 12345)
}

func main() {
    // Register the scheme for the custom resource
    scheme.AddToScheme(scheme.Scheme)
    examplev1.AddToScheme(scheme.Scheme)

    // Create a new controller manager
    mgr, err := ctrl.NewManager(ctrl.GetConfigOrDie(), ctrl.Options{Scheme: scheme.Scheme})
    if err != nil {
        setupLog.Fatal(err, "unable to start manager")
    }

    // Register the ProfileReconciler
    if err := (&ProfileReconciler{
        Client: mgr.GetClient(),
        Scheme: mgr.GetScheme(),
    }).SetupWithManager(mgr); err != nil {
        setupLog.Fatal(err, "unable to create controller", "controller", "Profile")
    }

    // Start the controller
    if err := mgr.Start(ctrl.SetupSignalHandler()); err != nil {
        setupLog.Fatal(err, "problem running manager")
    }
}

// Define the Profile Resource:
go
package v1
import (
    metav1 "k8s.io/apimachinery/pkg/apis/meta/v1"
)

```

```
// ProfileSpec defines the desired state of Profile
type ProfileSpec struct {
    // Name of the profile
    Name string `json:"name"`
    // Email address
    Email string `json:"email"`
    // Unique user ID
    UserID string `json:"userId"`
}

// ProfileStatus defines the observed state of Profile
type ProfileStatus struct {
    // Insert your status fields here
}

//+kubebuilder:object:root=true
//+kubebuilder:subresource:status
//+kubebuilder:resource:path=profiles,scope=Namespaced
type Profile struct {
    metav1.TypeMeta `json:"",inline"`
    metav1.ObjectMeta `json:"metadata,omitempty"`

    Spec ProfileSpec `json:"spec,omitempty"`
    Status ProfileStatus `json:"status,omitempty"`
}

//+kubebuilder:object:root=true
type ProfileList struct {
    metav1.TypeMeta `json:"",inline"`
    metav1.ListMeta `json:"metadata,omitempty"`
    Items []Profile `json:"items"`
}

func init() {
    SchemeBuilder.Register(&Profile{}, &ProfileList{})
}
```

3. Build and Deploy the Controller:

- Build the controller binary:

```
bash
go build
```

- Deploy the controller to your Kubernetes cluster:

- You can use 'kubectl apply -f deployment.yaml' to deploy the controller as a Deployment, making sure to include the necessary resources in the YAML file.

4. Create a Profile Resource:

- Create a profile YAML:

```
apiVersion: example.com/v1
kind: Profile
metadata:
  name: my-profile
spec:
  name: John Doe
  email: john.doe@example.com
```

- Apply the YAML:

```
bash
kubectl apply -f profile.yaml
```

5. Verify the User ID Generation:

- Get the profile resource:

```
bash
kubectl get profile my-profile -o yaml
```

- Check the 'userId' field in the output. You should see the automatically generated unique user ID. This comprehensive solution demonstrates how to implement a custom resource definition and a controller for managing user profiles in Kubernetes. You can adjust the code and logic according to your specific requirements. ,

NEW QUESTION # 73

You have a Deployment named 'api-deployment' that runs an API server. The API server handles sensitive data and must have strong security measures. You want to ensure that all pods within the Deployment are running with a specific security context that limits their capabilities. Describe the steps to configure a Securitycontext in the Deployment to enforce these security restrictions.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Define the SecurityContext:

- Add a 'securityContext' section to the container definition Within the Deployment's template.
- Define the desired security restrictions Within the 'securityContext' section
- 'runAsUser': Specifies the user ID under which the container should run.
- 'runAsGroup': Defines the group ID for the container.
- 'fsGroup': Sets the supplemental group ID for the container, giving access to specific files and directories.
- 'readOnlyRootFilesystem': Specifies whether the container should have read-only access to the root filesystem.
- 'capabilities': Configures the allowed capabilities for the container, limiting its privileges.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: api-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: api-server
  template:
    metadata:
      labels:
        app: api-server
    spec:
      containers:
        - name: api-server
          image: example/api-server:latest
          securityContext:
            runAsUser: 1000
            runAsGroup: 1000
            readOnlyRootFilesystem: true
            capabilities:
              drop: ["ALL"]
```

2. Apply the Deployment: - Use 'kubectl apply -f api-deployment.yaml' to update the Deployment with the security context configuration. 3. Verify the Security Context: - Examine the pod details using 'kubectl describe pod -l app=api-server' to confirm that the SecurityContext is applied to the containers. 4. Test Security Measures: - Run tests to ensure the security context is effectively limiting the capabilities of the API server pods.

NEW QUESTION # 74



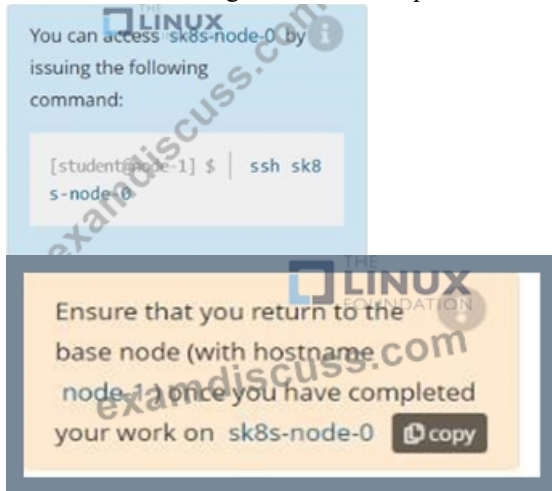
Context

A project that you are working on has a requirement for persistent data to be available.

Task

To facilitate this, perform the following tasks:

- * Create a file on node sk8s-node-0 at /opt/KDSP00101/data/index.html with the content Acct=Finance
- * Create a PersistentVolume named task-pv-volume using hostPath and allocate 1Gi to it, specifying that the volume is at /opt/KDSP00101/data on the cluster's node. The configuration should specify the access mode of ReadWriteOnce . It should define the StorageClass name exam for the PersistentVolume , which will be used to bind PersistentVolumeClaim requests to this PersistentVolume.
- * Create a PersistentVolumeClaim named task-pv-claim that requests a volume of at least 100Mi and specifies an access mode of ReadWriteOnce
- * Create a pod that uses the PersistentVolumeClaim as a volume with a label app: my-storage-app mounting the resulting volume to a mountPath /usr/share/nginx/html inside the pod




Answer:

Explanation:

See the solution below.

Explanation

Solution:



Readme

Web Terminal

THE LINUX FOUNDATION

```
student@sk8s-node-0:~$ echo 'Acct=Finance' > /opt/KDSP00101/data/index.html
student@sk8s-node-0:~$ vim pv.yml
```

Readme

Web Terminal

THE LINUX FOUNDATION

```
-- INSERT --
```

Readme

Web Terminal

THE LINUX FOUNDATION

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: task-pv-volume
spec:
  capacity:
    storage: 1Gi
  accessModes:
    - ReadWriteOnce
  storageClassName: storage
  hostPath:
    path: /opt/KDSP00101/data
    type: Directory
```

Readme

Web Terminal

THE LINUX FOUNDATION

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: task-pv-claim
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 100Mi
  storageClassName: storage
```

```

student@sk8s-node-0:~$ kubectl create -f pv.yml
persistentvolume/task-pv-volume created
student@sk8s-node-0:~$ kubectl create -f pvc.yml
persistentvolumeclaim/task-pv-claim created
student@sk8s-node-0:~$ kubectl get pv
NAME                CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS   CLAIM                STORAGECLASS   AGE
task-pv-volume      1Gi        RWO            Retain           Bound    default/task-pv-claim storage          11s
student@sk8s-node-0:~$ kubectl get pvc
NAME                STATUS   VOLUME             CAPACITY   ACCESS MODES   STORAGECLASS   AGE
task-pv-claim       Bound    task-pv-volume      1Gi        RWO            storage          9s
student@sk8s-node-0:~$ vim pod.yml

```

Readme
Web Terminal
THE LINUX FOUNDATION

```

apiVersion: v1
kind: Pod
metadata:
  name: mypod
  labels:
    app: my-storage-app
spec:
  containers:
    - name: myfrontend
      image: nginx
      volumeMounts:
        - mountPath: "/usr/share/nginx/html"
          name: mypod
  volumes:
    - name: mypod
      persistentVolumeClaim:
        claimName: task-pv-claim

```

17,32 All

Readme
Web Terminal
THE LINUX FOUNDATION

```

student@sk8s-node-0:~$ kubectl create -f pod.yml
pod/mypod created
student@sk8s-node-0:~$ kubectl get pods
NAME    READY   STATUS              RESTARTS   AGE
mypod   0/1     ContainerCreating   0           4s
student@sk8s-node-0:~$ kubectl get pods
NAME    READY   STATUS              RESTARTS   AGE
mypod   0/1     ContainerCreating   0           8s
student@sk8s-node-0:~$ kubectl get pods
NAME    READY   STATUS              RESTARTS   AGE
mypod   1/1     Running             0           10s
student@sk8s-node-0:~$ logout
Connection to 10.250.3.115 closed.
student@node-1:~$

```

NEW QUESTION # 75

You are managing a Kubernetes cluster that runs several microservices. One of these services, called "web-app", needs to have its pods scaled based on the current load. You have created a custom resource called "autoscaling.k8s.io/v1 alpha1" which represents the desired number of replicas for the "web-app" service- Implement a Kubernetes Controller that watches for changes in this custom resource and automatically scales the "web-app" Deployment accordingly.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create the Custom Resource Definition (CRD):

- Define the custom resource "autoscaling.k8s.io/v1alpha1" using the following YAML file:

```
apiVersion: apiextensions.k8s.io/v1
kind: CustomResourceDefinition
metadata:
  name: autoscalings.autoscaling.k8s.io
spec:
  group: autoscaling.k8s.io
  versions:
  - name: v1alpha1
    served: true
    storage: true
  scope: Namespaced
  names:
    plural: autoscalings
    singular: autoscaling
```

- Apply the CRD using 'kubectl apply -f crd.yaml' 2. Create the Custom Resource: - Create an instance of the custom resource 'Autoscaling' with the desired number of replicas for the "web-app" Deployment:

```
apiVersion: autoscaling.k8s.io/v1alpha1
kind: Autoscaling
metadata:
  name: web-app-autoscaler
spec:
  replicas: 3
```

- Apply the custom resource using 'kubectl apply -f autoscaling.yaml' 3. Create the Kubernetes Controller: - Create a Kubernetes Controller that watches for changes in the "Autoscaling" custom resource and updates the "web-app" Deployment.

```
go
package main

import (
    "context"
    appsv1 "k8s.io/api/apps/v1"
    corev1 "k8s.io/api/core/v1"
    metav1 "k8s.io/apimachinery/pkg/apis/meta/v1"
    "k8s.io/apimachinery/pkg/runtime"
    "k8s.io/apimachinery/pkg/runtime/schema"
    "k8s.io/apimachinery/pkg/watch"
    "k8s.io/client-go/kubernetes"
    "k8s.io/client-go/rest"
    "k8s.io/client-go/tools/cache"
    "k8s.io/client-go/tools/record"
    "k8s.io/client-go/util/workqueue"
    "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"
)

// Define the custom resource schema
var (
    scheme = runtime.NewScheme()
    gvk = schema.GroupVersionKind{
        Group: "autoscaling.k8s.io",
        Version: "v1alpha1",
        Kind: "Autoscaling",
    }
)

func init() {
    _ = corev1.AddToScheme(scheme)
    _ = appsv1.AddToScheme(scheme)
    _ = metav1.AddToGroupVersion(scheme, schema.GroupVersion{Group: "autoscaling.k8s.io", Version: "v1alpha1"})
}
```



```

// Define the controller logic
func reconcileAutoscaling(request reconcile.Request, c client.Client) (reconcile.Result, error) {
    // Fetch the Autoscaling resource
    autoscaling := &autoscalingv1alpha1.Autoscaling{}
    err := c.Get(context.Background(), request.NamespacedName, autoscaling)
    if err != nil {
        // Handle error
        return reconcile.Result{}, err
    }
    // Fetch the web-app Deployment
    deployment := &appsv1.Deployment{}
    err = c.Get(context.Background(), request.NamespacedName, deployment)
    if err != nil {
        // Handle error
        return reconcile.Result{}, err
    }

    // Update the Deployment's replicas
    deployment.Spec.Replicas = &autoscaling.Spec.Replicas
    err = c.Update(context.Background(), deployment)
    if err != nil {
        // Handle error
        return reconcile.Result{}, err
    }

    return reconcile.Result{}, nil
}

func main() {
    // Create a Kubernetes client
    cfg, err := rest.InClusterConfig()
    if err != nil {
        panic(err)
    }
    clientset, err := kubernetes.NewForConfig(cfg)
    if err != nil {
        panic(err)
    }

    // Create a manager
    mgr, err := manager.New(cfg, manager.Options{Scheme: scheme})
    if err != nil {
        panic(err)
    }

    // Create a controller for the Autoscaling custom resource
    err = controller.New("autoscaling-controller", mgr, controller.Options{Reconciler: &ReconcileAutoscaling{client: mgr.GetClient(), scheme: mgr.GetScheme(), recorder: mgr.GetEventRecorderFor("autoscaling-controller")}})
    if err != nil {
        panic(err)
    }

    // Watch for changes in the Autoscaling custom resource
    err = mgr.GetFieldIndexer().IndexField(context.Background(), &autoscalingv1alpha1.Autoscaling{}, "metadata.name", func(rawObj runtime.Object) []string {
        autoscaling := rawObj.(*autoscalingv1alpha1.Autoscaling)
        return []string{autoscaling.GetName()}
    })
    if err != nil {
        panic(err)
    }

    // Create a new source for the controller
    source := &source.Kind{Type: &autoscalingv1alpha1.Autoscaling{}}

    // Create a new event handler for the controller
    handler := &handler.EnqueueRequestForObject{}

    // Create a new controller for the Autoscaling custom resource
    err = mgr.Add(controller.New("autoscaling-controller", mgr, controller.Options{Reconciler: &ReconcileAutoscaling{client: mgr.GetClient(), scheme: mgr.GetScheme(), recorder: mgr.GetEventRecorderFor("autoscaling-controller")}}).
        Watch(source, handler))
    if err != nil {
        panic(err)
    }

    // Start the manager
    if err := mgr.Start(ctrl.SetupSignalHandler()); err != nil {
        panic(err)
    }
}

```

4. Deploy the Controller: - Build and deploy the controller to your Kubernetes cluster. 5. Verify the Controllers Functionality: - Make a change to the "Autoscaling" custom resource (e.g., increase the desired replicas). - Observe that the "web-app" Deployment is automatically scaled based on the new desired replica count. This code implements a basic Kubernetes Controller that monitors the "Autoscaling" custom resource. When the desired replicas are changed, the controller updates the "web-app" Deployment, ensuring the desired number of replicas is maintained. Note: This example assumes that the "web-app" Deployment exists in the same namespace as the "Autoscaling" custom resource. You might need to adapt the code for different deployments and

namespaces.,

NEW QUESTION # 76

You must switch to the correct cluster/configuration context. Failure to do so may result in a zero score.

[candidate@node-1] \$ kubectl config use-context sk8s



Task:

Update the Deployment app-1 in the frontend namespace to use the existing ServiceAccount app.

Answer:

Explanation:

See the solution below.

Explanation

Solution:

Text Description automatically generated

```
File Edit View Terminal Tabs Help
The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

candidate@node-1:~$ vi ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context sk8s
Switched to context "sk8s".
candidate@node-1:~$ vim .vimrc
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl apply -f ~/spicy-pikachu/backend-deployment.yaml
deployment.apps/backend-deployment configured
candidate@node-1:~$ kubectl get pods -n staging
NAME                                READY   STATUS    RESTARTS   AGE
backend-deployment-59d449b99d-cxct6 1/1     Running   0           20s
backend-deployment-59d449b99d-h2zjq 0/1     Running   0           9s
backend-deployment-78976f74f5-b8c85 1/1     Running   0           6h40m
backend-deployment-78976f74f5-flfsj 1/1     Running   0           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment 3/3      3             3           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME                READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment 3/3      3             3           6h41m
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl set serviceaccount deploy app-1 app -n frontend
deployment.apps/app-1 serviceaccount updated
candidate@node-1:~$
```

NEW QUESTION # 77

.....

The clients can download our CKAD exam questions and use our them immediately after they pay successfully. Our system will send our CKAD learning prep in the form of mails to the client in 5-10 minutes after their successful payment. The mails provide the links and if only the clients click on the links they can log in our software immediately to learn our CKAD Guide materials. It is fast and convenient!

New CKAD Dumps Free: <https://www.examdumps.com/Linux-Foundation/exam/CKAD/>

- [illegible]

BONUS!!! Download part of ExamDiscuss CKAD dumps for free: <https://drive.google.com/open?id=1ZW4tMKAsmcAqr4unwDkI48EBDjvxgUo7>