

Linux-Foundation

Exam Questions KCNA

Kubernetes and Cloud Native Associate (KCNA)



NEW QUESTION 1

kubeadm is an administrative dashboard for kubernetes

- A. False
- B. True

Answer: A

Explanation:

<https://kubernetes.io/docs/reference/setup-tools/kubeadm/>
 Graphical user interface, text, application Description automatically generated

Kubeadm

Kubeadm is a tool built to provide `kubeadm init` and `kubeadm join` as best-practice "fast paths" for creating Kubernetes clusters.

kubeadm performs the actions necessary to get a minimum viable cluster up and running. By design, it cares only about bootstrapping, not about provisioning machines. Likewise, installing various nice-to-have addons, like the Kubernetes Dashboard, monitoring solutions, and cloud-specific addons, is not in scope.



Instead, we expect higher-level and more tailored tooling to be built on top of kubeadm, and ideally, using kubeadm as the basis of all deployments will make it easier to create conformant clusters.

NEW QUESTION 2

What is the most common way to scale the application in the cloud environment?

- A. Parallel Scaling
- B. Horizontal Scaling
- C. Vertical Scaling

Answer: B

Explanation:

<https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/>

NEW QUESTION 3

What does CNCF stand for?

- A. Cloud Native Computing Foundation
- B. Cloud Native Cloud Foundation
- C. Cloud Native Container Foundation

Answer: A

Explanation:

<https://www.cncf.io/about/who-we-are/>
 Graphical user interface, text, application Description automatically generated

The Cloud Native Computing Foundation (CNCF) hosts critical components of the global technology infrastructure. CNCF brings together the world's top developers, end users, and vendors and runs the largest open source developer conferences. CNCF is part of the nonprofit Linux Foundation.

NEW QUESTION 4

What Kubernetes resource would allow you to run one Pod on some of your Nodes?

- A. DaemonSet
- B. ClusterSet
- C. Deployment
- D. ReplicaSet

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/> Graphical user interface, text, application Description automatically generated

DaemonSet

A DaemonSet ensures that all (or some) Nodes run a copy of a Pod. As nodes are added to the cluster, Pods are added to them. As nodes are removed from the cluster, those Pods are garbage collected. Deleting a DaemonSet will clean up the Pods it created.

Some typical uses of a DaemonSet are:

- running a cluster storage daemon on every node
- running a logs collection daemon on every node
- running a node monitoring daemon on every node

In a simple case, one DaemonSet, covering all nodes, would be used for each type of daemon. A more complex setup might use multiple DaemonSets for a single type of daemon, but with different flags and/or different memory and cpu requests for different hardware types.

NEW QUESTION 5

What standard does kubelet use to communicate with the container runtime?

- A. Service Mesh Interface (SMI)
- B. CRI-O
- C. ContainerD
- D. Container Runtime Interface (CRI)

Answer: D

Explanation:

kubelet can communicate with any runtime that supports the CRI standard.

NEW QUESTION 6

What is the primary interface for Kubernetes cluster?

- A. Kubernetes Api
- B. Kubelet
- C. YAML
- D. Control Plane
- E. JSON

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/overview/components/#kube-apiserver>

A screenshot of a computer Description automatically generated with medium confidence

kube-apiserver

The API server is a component of the Kubernetes control plane that exposes the Kubernetes API. The API server is the front end for the Kubernetes control plane.

The main implementation of a Kubernetes API server is [kube-apiserver](#). kube-apiserver is designed to scale horizontally—that is, it scales by deploying more instances. You can run several instances of kube-apiserver and balance traffic between those instances.

NEW QUESTION 7

Stateful set requires which service for the network identity of pods?

- A. Ingress
- B. Load Balancer Service
- C. Headless Service

Answer: C

Explanation:

<https://kubernetes.io/docs/concepts/workloads/controllers/statefulset/> Graphical user interface, text, application Description automatically generated

Using StatefulSets

StatefulSets are valuable for applications that require one or more of the following.

- Stable, unique network identifiers.
- Stable, persistent storage.
- Ordered, graceful deployment and scaling.
- Ordered, automated rolling updates.

NEW QUESTION 8

Which kubernetes resource type allows defining which pods are isolated when it comes to network-ing?

- A. Network policy
- B. Domain Name System 'DNS'
- C. Role Binding
- D. Service

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/services-networking/network-policies/#the-two-sorts-of-pod-isolation> Text, letter Description automatically generated

The Two Sorts of Pod Isolation

There are two sorts of isolation for a pod: isolation for egress, and isolation for ingress. They concern what connections may be established. "Isolation" here is not absolute, rather it means "some restrictions apply". The alternative, "non-isolated for \$direction", means that no restrictions apply in the stated direction. The two sorts of isolation (or not) are declared independently, and are both relevant for a connection from one pod to another.

By default, a pod is non-isolated for egress; all outbound connections are allowed. A pod is isolated for egress if there is any NetworkPolicy that both selects the pod and has "Egress" in its `policyTypes`; we say that such a policy applies to the pod for egress. When a pod is isolated for egress, the only allowed connections from the pod are those allowed by the `egress` list of some NetworkPolicy that applies to the pod for egress. The effects of those `egress` lists combine additively.

By default, a pod is non-isolated for ingress; all inbound connections are allowed. A pod is isolated for ingress if there is any NetworkPolicy that both selects the pod and has "Ingress" in its `policyTypes`; we say that such a policy applies to the pod for ingress. When a pod is isolated for ingress, the only allowed connections into the pod are those from the pod's node and those allowed by the `ingress` list of some NetworkPolicy that applies to the pod for ingress. The effects of those `ingress` lists combine additively.

NEW QUESTION 9

What kind of limitation cgroups allows?

- A. Prioritization

- B. Resource limiting
- C. Accounting
- D. None of the options
- E. Control
- F. Server cpu and memory

Answer: ABCE

NEW QUESTION 10

Which of the following is used to request storage in Kubernetes?

- A. PersistentVolume 'PV'
- B. PersistentVolumeClaim 'PVC'
- C. Container Storage Interface 'CSI'
- D. StorageClasses

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/storage/persistent-volumes/>

NEW QUESTION 10

A _____ is an application running on kubernetes.

- A. node
- B. pod
- C. workload
- D. container

Answer: C

Explanation:

<https://kubernetes.io/docs/concepts/workloads/> Text Description automatically generated

Workloads

A workload is an application running on Kubernetes. Whether your workload is a single component or several that work together, on Kubernetes you run it inside a set of *Pods*. In Kubernetes, a Pod represents a set of running containers on your cluster.

Kubernetes pods have a *defined lifecycle*. For example, once a pod is running in your cluster then a critical fault on the *node* where that pod is running means that all the pods on that node fail. Kubernetes treats that level of failure as final: you would need to create a new Pod to recover, even if the node later becomes healthy.

NEW QUESTION 14

The Kubernetes rolling update is used for _____ .

- A. Updating a service
- B. Scaling an application
- C. Updating a deployment

Answer: C

Explanation:

<https://kubernetes.io/docs/tutorials/kubernetes-basics/update/update-intro/>
Graphical user interface, text Description automatically generated with medium confidence

Performing a Rolling Update

Objectives

- Perform a rolling update using kubectl.

Updating an application

Users expect applications to be available all the time and developers are expected to deploy new versions of them several times a day. In Kubernetes this is done with rolling updates. **Rolling updates** allow Deployments' update to take place with zero downtime by incrementally updating Pods instances with new ones. The new Pods will be scheduled on Nodes with available resources.

In the previous module we scaled our application to run multiple instances. This is a requirement for performing updates without affecting application availability. By default, the maximum number of Pods that can be unavailable during the update and the maximum number of new Pods that can be created, is one. Both options can be configured to either numbers or percentages (of Pods). In Kubernetes, updates are versioned and any Deployment update can be reverted to a previous (stable) version.

Summary:

- Updating an app

Rolling updates allow Deployments' update to take place with zero downtime by incrementally updating Pods instances with new ones.

NEW QUESTION 19

What is container orchestration?

- A. Packaging code and all of its dependencies into a single executable
- B. Adding code to a container image so it can run as a container
- C. Using automation to manage containers
- D. Spinning a new containers to replace old ones

Answer: C

Explanation:

<https://www.redhat.com/en/topics/containers/what-is-container-orchestration> Text Description automatically generated

Container orchestration automates the deployment, management, scaling, and networking of containers. Enterprises that need to deploy and manage hundreds or thousands of **Linux® containers** and hosts can benefit from container orchestration.

Container orchestration can be used in any environment where you use containers. It can help you to deploy the same application across different environments without needing to redesign it. And **microservices** in containers make it easier to orchestrate services, including storage, networking, and security.

NEW QUESTION 23

Which project is not a dominant CNCF project in the storage landscape?

- A. Envoy
- B. Vitess
- C. Rook
- D. TiKV

Answer: A

Explanation:

<https://github.com/cncf/landscape#trail-map>

A picture containing timeline Description automatically generated



CLOUD NATIVE TRAIL MAP

The Cloud Native Landscape (CNCF.io) has a large number of options. This Cloud Native Trail Map is a recommended process for leveraging open source, cloud native technologies. At each step, you can choose a vendor-supported offering or do it yourself, and everything after step #3 is optional based on your circumstances.

HELP ALONG THE WAY

A. Training and Certification

Consider training offerings from CNCF and then take the exam to become a Certified Kubernetes Administrator or a Certified Kubernetes Application Developer cncf.io/training

B. Consulting Help

If you want assistance with Kubernetes and the surrounding ecosystem, consider leveraging a Kubernetes Certified Service Provider: cncf.io/kspp

C. Join CNCF's End User Community

For companies that don't offer cloud native services externally: cncf.io/enduser

WHAT IS CLOUD NATIVE?

Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds. Containers, service meshes, microservices, immutable infrastructure, and declarative APIs exemplify this approach.

These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with robust automation, they allow engineers to make high-impact changes frequently and predictably with minimal toil.

The Cloud Native Computing Foundation seeks to drive adoption of this paradigm by fostering and sustaining an ecosystem of open source, vendor-neutral projects. We democratize state-of-the-art patterns to make these innovations accessible for everyone.

cncf.io

v20200501



NEW QUESTION 28

Which tool is built on the GitOps toolkit?

- A. Jenkins-X
- B. GitHub Workflow & Actions
- C. Flux
- D. Jenkins
- E. ArgoCD
- F. Travis CI

Answer: C

Explanation:

<https://fluxcd.io/#gitops-toolkit>

Graphical user interface, text, application Description automatically generated

GitOps Toolkit

The set of APIs and controllers that make up the runtime for Flux. You can use the GitOps Toolkit to extend Flux, and to build your own systems for continuous delivery.

Note: Argo CD is a GitOps tool and not using GitOps toolkit

NEW QUESTION 30

Fluentd is the only way to export logs from Kubernetes cluster or applications running in cluster

- A. True
- B. False

Answer: B

Explanation:

<https://github.com/cncf/landscape#trail-map>

A picture containing timeline Description automatically generated

CLOUD NATIVE COMPUTING FOUNDATION

CLOUD NATIVE TRAIL MAP

The Cloud Native Landscape (2022.0) has a large number of options. This Cloud Native Trail Map is a recommended process for leveraging open source, cloud native technologies. At each step, you can choose a vendor-supported offering or do it yourself, and everything after step #3 is optional based on your circumstances.

HELP ALONG THE WAY

A. Training and Certification
 Consider training offerings from CNCF and then take the exams to become a Certified Kubernetes Administrator or a Certified Kubernetes Application Developer. cncf.io/training

B. Consulting Help
 If you want assistance with Kubernetes and the surrounding ecosystem, consider leveraging a Kubernetes Certified Service Provider. cncf.io/csp

C. Join CNCF's End User Community
 For companies that don't offer cloud native services externally. cncf.io/enduser

WHAT IS CLOUD NATIVE?

Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds. Containers, service meshes, microservices, immutable infrastructure, and declarative APIs exemplify this approach.

These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with robust automation, they allow engineers to make high-impact changes frequently and predictably with minimal toil.

The Cloud Native Computing Foundation seeks to drive adoption of this paradigm by fostering and sustaining an ecosystem of open source, vendor-neutral projects. We democratize state-of-the-art patterns to make these innovations accessible for everyone.

l.cncf.io
v20200501

1. CONTAINERIZATION

- Commonly done with Docker containers
- Any size application and dependencies (even PDF-1.1 code running on an emulator) can be containerized
- Over time, you should aspire towards splitting suitable applications and writing future functionality as microservices

2. CI/CD

- Setup Continuous Integration/Continuous Delivery (CI/CD) so that changes to your source code automatically result in a new container being built, tested, and deployed to staging and eventually, perhaps, to production
- Setup automated rollouts, roll backs and testing
- Argo is a set of Kubernetes-native tools for deploying and running jobs, applications, workflows, and events using GitOps paradigms such as continuous and progressive delivery and MLops

3. ORCHESTRATION & APPLICATION DEFINITION

- Kubernetes is the market-leading orchestration solution
- You should select a Certified Kubernetes Distribution, Hosted Platform, or Installer: cncf.io/ark
- Helm Charts help you define, install, and upgrade even the most complex Kubernetes application

4. OBSERVABILITY & ANALYSIS

- Pick solutions for monitoring, logging and tracing
- Consider CNCF projects Prometheus for monitoring, Fluentd for logging and Jaeger for Tracing
- For tracing, look for an OpenTracing-compatible implementation like Jaeger

5. SERVICE PROXY, DISCOVERY, & MESH

- CoreDNS is a fast and flexible tool that is useful for service discovery
- Envoy and Linkerd each enable service mesh architectures
- They offer health checking, routing, and load balancing

6. NETWORKING, POLICY, & SECURITY

To enable more flexible networking, use a CNI-compliant network project like Calico, Flannel, or Weave Net. Open Policy Agent (OPA) is a general purpose policy engine with uses ranging from authorization and admission control to data filtering. Falco is an anomaly detection engine for cloud native.

7. DISTRIBUTED DATABASE & STORAGE

When you need more resiliency and scalability than you can get from a single database, Vitess is a good option for running MySQL at scale through sharding. Rook is a storage orchestrator that integrates a diverse set of storage solutions into Kubernetes. Serving as the "brain" of Kubernetes, etcd provides a reliable way to store data across a cluster of machines. TiKV is a high performance, distributed transactional key-value store written in Rust.

8. STREAMING & MESSAGING

When you need higher performance than JSON-RPC, consider using gRPC or NATS. gRPC is a universal RPC framework. NATS is a multi-modal messaging system that includes request/reply, pub/sub and load balanced queues. CloudEvents is a specification for describing event data in common ways.

9. CONTAINER REGISTRY & RUNTIME

Harbor is a registry that stores, signs, and scans content. You can use alternative container runtimes. The most common, both of which are OCI-compliant, are containerd and CRI-O.

10. SOFTWARE DISTRIBUTION

If you need to do secure software distribution, evaluate Notary, an implementation of The Update Framework.

NEW QUESTION 34

What is Open Container Initiative 'OCI'?

- A. A protocol for communicating with the kubernetes api
- B. The governing body of the Cloud Native Computing Foundation 'CNCF'
- C. An open standard for managing service mesh in kubernetes
- D. An organization that creates open standards for containers

Answer: D

Explanation:

<https://opencontainers.org/>

Text Description automatically generated

Open Container Initiative

The **Open Container Initiative** is an open governance structure for the express purpose of creating open industry standards around container formats and runtimes.

Established in June 2015 by Docker and other leaders in the container industry, the OCI currently contains three specifications: the Runtime Specification (runtime-spec), the Image Specification (image-spec) and the Distribution Specification (distribution-spec). The Runtime Specification outlines how to run a "filesystem bundle" that is unpacked on disk. At a high-level an OCI implementation would download an OCI Image then unpack that image into an OCI Runtime filesystem bundle. At this point the OCI Runtime Bundle would be run by an OCI Runtime.

NEW QUESTION 37

What is the command to list all the available objects in your Kubernetes cluster?

- A. kubectl get all
- B. kubectl get api-resources
- C. kubectl api-resources
- D. kubectl get pods

Answer: C

Explanation:

<https://kubernetes.io/docs/reference/kubectl/cheatsheet/>

Graphical user interface, text, application, email Description automatically generated

Resource types

List all supported resource types along with their shortnames, **API group**, whether they are **namespaced**, and **Kind**:

```
kubectl api-resources
```

NEW QUESTION 41

How would you return all the pod data in the json format using kubectl command?

- A. kubectl get pods -o json
- B. kubectl get pods --all-namespaces
- C. kubectl get pods -o wide
- D. kubectl get pods -o jsonpath

Answer: A

Explanation:

<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#get>

NEW QUESTION 43

What Linux feature is used to provide isolation for containers?

- A. Processes
- B. Services
- C. NetworkPolicy
- D. Control groups

Answer: D

Explanation:

Control groups provide isolation for container processes, keeping them separate from other process-es on the host.

NEW QUESTION 48

How should folks new to the cloud native ecosystem, go about learning the different aspects of the ecosystem?

- A. by signing up the CNCF slack
- B. by reading the Kubernetes documentation
- C. by looking at the cloud native landscape
- D. by looking at the cloud native trail-map

Answer: D

Explanation:

<https://github.com/cncf/landscape#trail-map>

NEW QUESTION 53

What does the 'kops' acronym means?

- A. Kubernetes Open Platform Specification
- B. Kubernetes Operations
- C. Kubernetes Operators
- D. Kubernetes Operation Policy Specification

Answer: B

Explanation:

<https://github.com/kubernetes/kops>

Graphical user interface, text, application, email Description automatically generated

kOps - Kubernetes Operations

go report A+ reference

The easiest way to get a production grade Kubernetes cluster up and running.

What is kOps?

We like to think of it as `kubect1` for clusters.

`kops` will not only help you create, destroy, upgrade and maintain production-grade, highly available, Kubernetes cluster, but it will also provision the necessary cloud infrastructure.

AWS (Amazon Web Services) and GCE (Google Cloud Platform) are currently officially supported, with DigitalOcean, Hetzner and OpenStack in beta support, and Azure in alpha.

NEW QUESTION 58

Which of the following is not a stop on the cloud native trailmap?

- A. Microservices
- B. CI/CD
- C. Containerization
- D. Software distribution

Answer: A

Explanation:

<https://github.com/cncf/landscape#trail-map>

NEW QUESTION 63

Which of the following is not the required field to describe Kubernetes objects?

- A. metadata
- B. apiVersion
- C. Kind
- D. Container
- E. spec

Answer: D

Explanation:

<https://kubernetes.io/docs/concepts/overview/working-with-objects/kubernetes-objects/> Graphical user interface, text, application Description automatically generated

Required Fields [↔](#)

In the `.yaml` file for the Kubernetes object you want to create, you'll need to set values for the following fields:

- `apiVersion` - Which version of the Kubernetes API you're using to create this object
- `kind` - What kind of object you want to create
- `metadata` - Data that helps uniquely identify the object, including a `name` string, `UID`, and optional `namespace`
- `spec` - What state you desire for the object

The precise format of the object `spec` is different for every Kubernetes object, and contains nested fields specific to that object. The [Kubernetes API Reference](#) can help you find the `spec` format for all of the objects you can create using Kubernetes.

NEW QUESTION 65

You might need to run a stateless application in Kubernetes, and you want to be able to scale easily and perform rolling updates. What Kubernetes resource type can you use to do this?

- A. Daemon set
- B. Replica set
- C. Deployment
- D. pod
- E. service
- F. Stateful set

Answer: C

Explanation:

<https://kubernetes.io/docs/concepts/workloads/controllers/deployment/> Graphical user interface, text, application Description automatically generated

Deployments

A *Deployment* provides declarative updates for Pods and ReplicaSets.

You describe a *desired state* in a Deployment, and the Deployment Controller changes the actual state to the desired state at a controlled rate. You can define Deployments to create new ReplicaSets, or to remove existing Deployments and adopt all their resources with new Deployments.

Note: Do not manage ReplicaSets owned by a Deployment. Consider opening an issue in the main Kubernetes repository if your use case is not covered below.

NEW QUESTION 69

Have a pod 'hello' and a container in that pod 'green'. Which of the following commands would get the logs for that container?

- A. `alias k='kubectl' k logs -p hello -c green`
- B. `alias k='kubectl' k logs hello -c green`
- C. `alias k='kubectl' k get logs -p hello -c green`
- D. `alias k='kubectl' k logs -p hello green`

Answer: B

Explanation:

<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#logs> Graphical user interface, text, application, email Description automatically generated

logs

Print the logs for a container in a pod or specified resource. If the pod has only one container, the container name is optional.

Usage

```
$ kubectl logs [-f] [-p] (POD | TYPE/NAME) [-c CONTAINER]
```

Flags

Name	Shorthand	Default	Usage
all-containers		false	Get all containers' logs in the pod(s).
container	c		Print the logs of this container
follow	f	false	Specify if the logs should be streamed.

Return snapshot logs from pod nginx with only one container

```
kubectl logs nginx
```

Return snapshot logs from pod nginx with multi containers

```
kubectl logs nginx --all-containers=true
```

Return snapshot logs from all containers in pods defined by label app=nginx

```
kubectl logs -l app=nginx --all-containers=true
```

Return snapshot of previous terminated ruby container logs from pod web-1

```
kubectl logs -p -c ruby web-1
```

Begin streaming the logs of the ruby container in pod web-1

```
kubectl logs -f -c ruby web-1
```

NEW QUESTION 72

Open Container Initiative set container standards for

- A. Code, Build, Distribute, Deploy containers
- B. Run, build, and image
- C. Code, Build, Distribute containers
- D. Run, Build, Distribute containers

Answer: D

NEW QUESTION 73

How can you achieve cost optimization in the cloud environment?

- A. Use On Demand instances
- B. Use Spot Instances
- C. Use Reserved Instances
- D. Use Bare Metal

Answer: C

NEW QUESTION 78

How can persistent volume be provisioned?

- A. Automatically
- B. Bootstrap
- C. Dynamically

Answer: C

Explanation:

<https://kubernetes.io/docs/concepts/storage/persistent-volumes/> Text Description automatically generated

A *PersistentVolume* (PV) is a piece of storage in the cluster that has been provisioned by an administrator or dynamically provisioned using [Storage Classes](#). It is a resource in the cluster just like a node is a cluster resource. PVs are volume plugins like Volumes, but have a lifecycle independent of any individual Pod that uses the PV. This API object captures the details of the implementation of the storage, be that NFS, iSCSI, or a cloud-provider-specific storage system.

NEW QUESTION 82

What is not semantic versioning?

- A. 1.0.0
- B. 2022-05-04
- C. 1.0.0-alpha
- D. 1.0.0-beta.2

Answer: B

Explanation:

<https://semver.org/>

RegEx SemVer at <https://regex101.com/r/vkijKf/1/>

NEW QUESTION 84

Observability and monitoring are not the same?

- A. True
- B. False

Answer: A

NEW QUESTION 88

What command to view the kube config?

- A. kubectl view config
- B. kubectl config view
- C. kubectl get kubeconfig

Answer: B

Explanation:

<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#-em-view-em-> Graphical user interface, text, application Description automatically generated

view

Display merged kubeconfig settings or a specified kubeconfig file.

You can use `--output jsonpath={...}` to extract specific values using a jsonpath expression.

Usage

```
$ kubectl config view
```



NEW QUESTION 89

The Kubernetes API provides an interface for storing objects. Which of the following describes the type of objects stored by the Kubernetes API?

- A. Containers
- B. REST
- C. YAML
- D. ETCD

Answer: B

Explanation:

Kubernetes objects are RESTful objects.

NEW QUESTION 90

What is the name for a service that has no clusterIp address?

- A. Headless
- B. NodePort
- C. ClusterIP
- D. LoadBalancer

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/services-networking/service/#headless-services>
 Text, letter Description automatically generated

Headless Services

Sometimes you don't need load-balancing and a single Service IP. In this case, you can create what are termed "headless" Services, by explicitly specifying "None" for the cluster IP (`.spec.clusterIP`).

You can use a headless Service to interface with other service discovery mechanisms, without being tied to Kubernetes' implementation.

For headless Services , a cluster IP is not allocated, kube-proxy does not handle these Services, and there is no load balancing or proxying done by the platform for them. How DNS is automatically configured depends on whether the Service has selectors defined:

NEW QUESTION 94

What kubectl command is used to edit a resource on the server?

- A. kubectl resource modify
- B. kubectl update resource
- C. kubectl edit
- D. kubectl resource edit

Answer: C

Explanation:

<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#edit> Graphical user interface, text, application, email Description automatically generated

edit

Edit a resource from the default editor.

The edit command allows you to directly edit any API resource you can retrieve via the command-line tools. It will open the editor defined by your KUBE_EDITOR, or EDITOR environment variables, or fall back to 'vi' for Linux or 'notepad' for Windows. You can edit multiple objects, although changes are applied one at a time. The command accepts file names as well as command-line arguments, although the files you point to must be previously saved versions of resources.

Editing is done with the API version used to fetch the resource. To edit using a specific API version, fully-qualify the resource, version, and group.

The default format is YAML. To edit in JSON, specify "-o json".

The flag --windows-line-endings can be used to force Windows line endings, otherwise the default for your operating system will be used.

In the event an error occurs while updating, a temporary file will be created on disk that contains your unapplied changes. The most common error when updating a resource is another editor changing the resource on the server. When this occurs, you will have to apply your changes to the newer version of the resource, or update your temporary saved copy to include the latest resource version.



The screenshot shows a list of kubectl edit command examples:

- Edit the service named 'registry'**
`kubectl edit svc/registry`
- Use an alternative editor**
`KUBE_EDITOR="nano" kubectl edit svc/registry`
- Edit the job 'myjob' in JSON using the v1 API format**
`kubectl edit job.v1.batch/myjob -o json`
- Edit the deployment 'mydeployment' in YAML and save the modified config in its annotation**
`kubectl edit deployment/mydeployment -o yaml --save-config`
- Edit the deployment/mydeployment's status subresource**
`kubectl edit deployment mydeployment --subresource='status'`

NEW QUESTION 99

Which is NOT a use case for the Kubernetes dashboard?

- A. Troubleshooting any issues with applications
- B. Managing running applications
- C. Installing new Kubernetes cluster
- D. Managing the entire Kubernetes cluster

Answer: C

NEW QUESTION 104

The three typical opentelemetry data is?

- A. Metrics
- B. Traces

- C. Logs
- D. All of the options

Answer: D

Explanation:

<https://opentelemetry.io/docs/concepts/data-sources/> Text Description automatically generated

What is OpenTelemetry?

OpenTelemetry is a set of APIs, SDKs, tooling and integrations that are designed for the creation and management of *telemetry data* such as traces, metrics, and logs. The project provides a vendor-agnostic implementation that can be configured to send telemetry data to the backend(s) of your choice. It supports a variety of popular open-source projects including Jaeger and Prometheus.

NEW QUESTION 105

Which of the following best describes the way K8S Role-based access control (RBAC) works?

- A. K8S does not do RBAC or Cluster role
- B. RBAC lists which operations are denied to users
- C. States which users can perform which actions against the resources.

Answer: C

Explanation:

<https://kubernetes.io/docs/reference/access-authn-authz/rbac/> Text Description automatically generated

When the kube-apiserver is run with a log level of 5 or higher for the RBAC component (`--vmodule=rbac*=5` or `--v=5`), you can see RBAC denials in the API server log (prefixed with `RBAC`). You can use that information to determine which roles need to be granted to which users, groups, or service accounts.

Once you have [granted roles to service accounts](#) and workloads are running with no RBAC denial messages in the server logs, you can remove the ABAC authorizer.

NEW QUESTION 110

Which organizational persona creates Service Level Agreements 'SLA', Service Level Objectives 'SLO', and Service Level Indicator 'SLI'?

- A. Developer
- B. DevSecOps
- C. Site Reliability Engineer (SRE)
- D. Security and Compliance Engineer
- E. DevOps

Answer: C

Explanation:

SREs create SLAs, SLOs, and SLIs to define and implement standards for application and infra-structure reliability.

NEW QUESTION 113

Which of the following best describes the way kubernetes Role-based access control (RBAC) works?

- A. Kubernetes does not do RBAC
- B. Kubernetes RBAC states which users can perform which actions against which re-source
- C. Kubernetes RBAC lists which operations on which resources are denied to users
- D. Kubernetes RBAC is responsible for authenticating subjects such as users and groups

Answer: B

Explanation:

<https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

Graphical user interface, text, application, email Description automatically generated

Using RBAC Authorization

Role-based access control (RBAC) is a method of regulating access to computer or network resources based on the roles of individual users within your organization.

RBAC authorization uses the `rbac.authorization.k8s.io` API group to drive authorization decisions, allowing you to dynamically configure policies through the Kubernetes API.

To enable RBAC, start the `API server` with the `--authorization-mode` flag set to a comma-separated list that includes `RBAC`; for example:

```
kube-apiserver --authorization-mode=Example,RBAC --other-options --more-options
```

NEW QUESTION 114

Which kubernetes object do deployments use behind the scenes when they need to scale pods?

- A. Horizontal pod autoscaler
- B. ReplicaSets
- C. kubectl
- D. Replication controller

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/workloads/controllers/replicaset/>

Graphical user interface, text, application, email Description automatically generated

ReplicaSet

A ReplicaSet's purpose is to maintain a stable set of replica Pods running at any given time. As such, it is often used to guarantee the availability of a specified number of identical Pods.

NEW QUESTION 116

Which statement is true about Pod Networking?

- A. All pod requires an external DNS server to get the hostname
- B. All containers in a pod get a unique IP address
- C. All containers in a pod share a single IP address
- D. All pod requires NAT to get a unique IP address.

Answer: C

Explanation:

<https://kubernetes.io/docs/concepts/workloads/pods/#pod-networking> Text Description automatically generated

Pod networking

Each Pod is assigned a unique IP address for each address family. Every container in a Pod shares the network namespace, including the IP address and network ports. Inside a Pod (and **only** then), the containers that belong to the Pod can communicate with one another using `localhost`. When containers in a Pod communicate with entities *outside the Pod*, they must coordinate how they use the shared network resources (such as ports). Within a Pod, containers share an IP address and port space, and can find each other via `localhost`. The containers in a Pod can also communicate with each other using standard inter-process communications like SystemV semaphores or POSIX shared memory. Containers in different Pods have distinct IP addresses and can not communicate by OS-level IPC without special configuration. Containers that want to interact with a container running in a different Pod can use IP networking to communicate.

Containers within the Pod see the system hostname as being the same as the configured `name` for the Pod. There's more about this in the [networking](#) section.

NEW QUESTION 118

What is the functionality of the daemon set?

- A. To run a copy of the pod in all the nodes of the cluster
- B. To initialize the pod before starting the main pod
- C. To run a copy of the pod in a single node of the cluster

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/>

Graphical user interface, text, application Description automatically generated with medium confidence

DaemonSet

A *DaemonSet* ensures that all (or some) Nodes run a copy of a Pod. As nodes are added to the cluster, Pods are added to them. As nodes are removed from the cluster, those Pods are garbage collected. Deleting a DaemonSet will clean up the Pods it created.

Some typical uses of a DaemonSet are:

- running a cluster storage daemon on every node
- running a logs collection daemon on every node
- running a node monitoring daemon on every node

NEW QUESTION 121

How to get the logs of the previously terminated nginx container from the web pod?

- A. `kubectl logs -p -c nginx web`
- B. `kubectl logs nginx`
- C. `kubectl logs -p -c web nginx`
- D. `kubectl logs -f -c nginx web`

Answer: A

Explanation:

<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#logs> Text Description automatically generated with medium confidence

Return snapshot of previous terminated ruby container logs from pod web-1

```
kubectl logs -p -c ruby web-1
```

NEW QUESTION 126

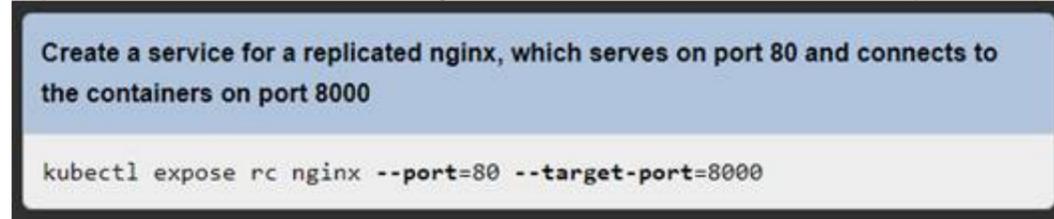
Which command is used to expose Kubernetes service

- A. kubectl expose
- B. kubectl create
- C. kubectl run

Answer: A

Explanation:

<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#expose> Text Description automatically generated



NEW QUESTION 131

Which command-line tool is used to interact with the Kubernetes cluster?

- A. kube-api
- B. kubectl
- C. kube-scheduler

Answer: B

Explanation:

<https://kubernetes.io/docs/reference/kubectl/>

Graphical user interface, text, application, email Description automatically generated

Command line tool (kubectl)

Kubernetes provides a command line tool for communicating with a Kubernetes cluster's control plane, using the Kubernetes API.

This tool is named `kubectl`.

For configuration, `kubectl` looks for a file named `config` in the `$HOME/.kube` directory. You can specify other `kubeconfig` files by setting the `KUBECONFIG` environment variable or by setting the `--kubeconfig` flag.

This overview covers `kubectl` syntax, describes the command operations, and provides common examples. For details about each command, including all the supported flags and subcommands, see the [kubectl](#) reference documentation.

For installation instructions, see [Installing kubectl](#); for a quick guide, see the [cheat sheet](#). If you're used to using the `docker` command-line tool, [kubectl for Docker Users](#) explains some equivalent commands for Kubernetes.

NEW QUESTION 133

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

KCNA Practice Exam Features:

- * KCNA Questions and Answers Updated Frequently
- * KCNA Practice Questions Verified by Expert Senior Certified Staff
- * KCNA Most Realistic Questions that Guarantee you a Pass on Your First Try
- * KCNA Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The KCNA Practice Test Here](#)