


KCNA Dumps

Kubernetes and Cloud Native Associate (KCNA)

<https://www.certleader.com/KCNA-dumps.html>



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

How to create deployment name app-dep, image=nginx, and replicas 5 using imperative command?

- A. `kubectrl create app-dep deployment --image=nginx --replicas=5`
- B. `kubectrl create deployment app-dep --image=nginx --replicas=5`
- C. `kubectrl create app-dep deployment --replicas=5 --image=nginx`

Answer: B

Explanation:

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

Create a deployment named my-dep that runs the nginx image with 3 replicas

```
kubectrl create deployment my-dep --image=nginx --replicas=3
```

NEW QUESTION 3

What is the name of the Kubernetes agent that runs on each worker nodes?


- A. kubelet
- B. systemd
- C. kube-proxy
- D. pod

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/overview/components/>

Graphical user interface, diagram, application Description automatically generated



NEW QUESTION 4

To specify a Kubernetes object which language is used?

- A. JSON
- B. Go
- C. YAML
- D. Node
- E. Python

Answer: C

Explanation:

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

Understanding Kubernetes Objects

This page explains how Kubernetes objects are represented in the Kubernetes API, and how you can express them in `.yaml` format.

NEW QUESTION 5

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 6


There are three Nodes in a cluster, and want to run exactly one replica of a Pod on each Node. Prefer to automatically create a replica on any new Nodes when they are added. Which Kubernetes resource should you use?

- A. DaemonSet
- B. ReplicaSet
- C. NodeSet
- D. StatefulSet
- E. Deployment

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/>
A DaemonSet runs replicas on all (or just some) Nodes in the cluster.
Table Description automatically generated with medium confidence



NEW QUESTION 7

What tool allows us to build useful visual representations of prometheus data?

- A. Grafana
- B. kubectl
- C. Distributed system tracing
- D. Rook
- E. Kibana

Answer: A

Explanation:

<https://prometheus.io/>
Graphical user interface, text, application Description automatically generated



Great visualization

Prometheus has multiple modes for visualizing data: a built-in expression browser, Grafana integration, and a console template language.

NEW QUESTION 8

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 9

What can you use to add new resource types to your cluster?

- A. start container
- B. CustomResourceDefinitions
- C. init container
- D. Flux
- E. CRI-O

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/> Graphical user interface, text, application Description automatically generated

CustomResourceDefinitions

The [CustomResourceDefinition](#) API resource allows you to define custom resources.

Defining a CRD object creates a new custom resource with a name and schema that you specify. The Kubernetes API serves and handles the storage of your custom resource. The name of a CRD object must be a valid [DNS subdomain name](#).

This frees you from writing your own API server to handle the custom resource, but the generic nature of the implementation means you have less flexibility than with [API server aggregation](#).

Refer to the [custom controller example](#) for an example of how to register a new custom resource, work with instances of your new resource type, and use a controller to handle events.

NEW QUESTION 10

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 14

A new Pod is created. Then, the Pod is assigned to a Node. Which Kubernetes component was re-sponsible for determining which Node to assign the Pod to?

- A. kubelet
- B. Scheduler
- C. API Server
- D. Controller manager

Answer: B

Explanation:

<https://kubernetes.io/docs/reference/command-line-tools-reference/kube-scheduler/> Graphical user interface, text, application Description automatically generated

The Kubernetes scheduler is a control plane process which assigns Pods to Nodes. The scheduler determines which Nodes are valid placements for each Pod in the scheduling queue according to constraints and available resources. The scheduler then ranks each valid Node and binds the Pod to a suitable Node. Multiple different schedulers may be used within a cluster; kube-scheduler is the reference implementation. See [scheduling](#) for more information about scheduling and the kube-scheduler component.

```
kube-scheduler [flags]
```

NEW QUESTION 16


Which of the following components is part of the Kubernetes control panel

- A. kubectl
- B. kube-proxy
- C. Service Mesh
- D. kubelet
- E. Cloud control manager

Answer: E

Explanation:

<https://kubernetes.io/docs/concepts/overview/components/> Diagram Description automatically generated



NEW QUESTION 21

What tool allows you to create self-managing, self-scaling, self-healing storage?

- A. Persistent Volume
- B. Persistent Volume Claim
- C. Storage Class
- D. Rook
- E. Volume

Answer: D

Explanation:

<https://rook.io/>

Text Description automatically generated

Storage Operators for Kubernetes

Rook turns distributed storage systems into self-managing, self-scaling, self-healing storage services. It automates the tasks of a storage administrator: deployment, bootstrapping, configuration, provisioning, scaling, upgrading, migration, disaster recovery, monitoring, and resource management.

Rook uses the power of the Kubernetes platform to deliver its services via a Kubernetes Operator for each storage provider.

NEW QUESTION 23

What is the command used to scale the application?

- A. kubectl run
- B. kubectl explain
- C. kubectl scale

Answer: C

Explanation:

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

scale

Set a new size for a deployment, replica set, replication controller, or stateful set.

Scale also allows users to specify one or more preconditions for the scale action.

If `--current-replicas` or `--resource-version` is specified, it is validated before the scale is attempted, and it is guaranteed that the precondition holds true when the scale is sent to the server.

Usage

```
$ kubectl scale [--resource-version=version] [--current-replicas=count] --replicas=COUNT (-f FILENAME | TYPE NAME)
```



NEW QUESTION 28

What are the two major components of service mesh?

- A. Control plane and Data plane
- B. Master plane and Data plane
- C. None of the options
- D. Controller plane and User plane
- E. Master plane and User plane

Answer: A

Explanation:

<https://istio.io/latest/about/service-mesh/>

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

How it Works

Istio has two components: the data plane and the control plane.

The data plane is the communication between services. Without a service mesh, the network doesn't understand the traffic being sent over, and can't make any decisions based on what type of traffic it is, or who it is from or to.

NEW QUESTION 30

What is OPA?


- A. Open Permission Agent
- B. Online Policy Audit
- C. Open Policy Agent
- D. Offline Policy Accessor

Answer: C

Explanation:

<https://www.cncf.io/projects/open-policy-agent-opa/>

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



NEW QUESTION 33

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 38

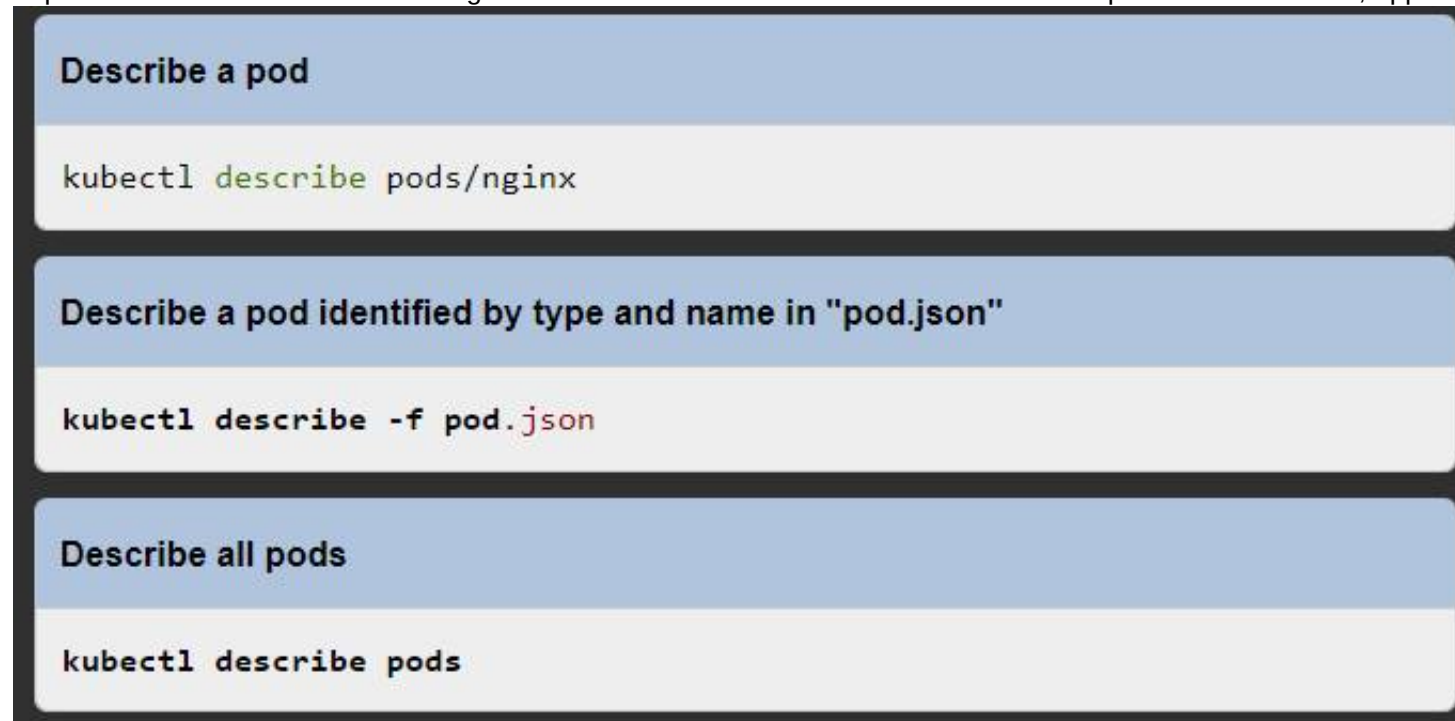
Which of the following command is used to get detailed information about the pod?

- A. kubectl info
- B. kubectl get
- C. kubectl describe
- D. kubectl explain

Answer: C

Explanation:

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



NEW QUESTION 40

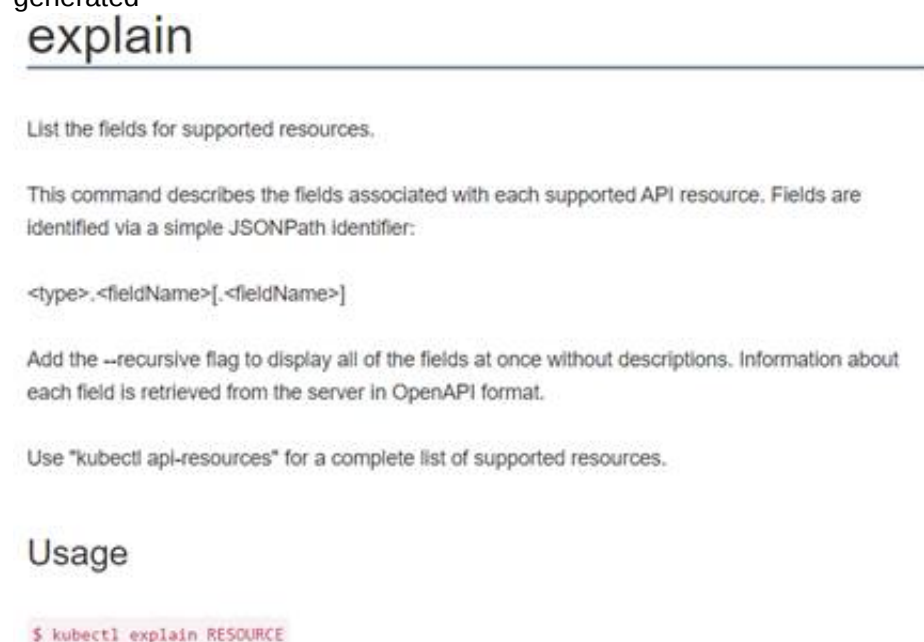
What command can you use to get documentation about a resource type from the command line?

- A. kubectl api-resources
- B. kubectl explain
- C. kubectl get
- D. kubeadm get-resource

Answer: B

Explanation:

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



NEW QUESTION 43

Which of the following is not the Kubernetes AutoScaling Strategy?

- A. Horizontal Pod Autoscaler
- B. Cluster Autoscaler
- C. Vertical Pod Autoscaler
- D. Load Balancing AutoScaler

Answer: D

Explanation:

<https://learnk8s.io/kubernetes-autoscaling-strategies>

Graphical user interface, text Description automatically generated with medium confidence

In Kubernetes, several things are referred to as "autoscaling", including:

- Horizontal Pod Autoscaler.
- Vertical Pod Autoscaler.
- Cluster Autoscaler.

NEW QUESTION 46

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 49

Flux is built using which toolkit?

- A. CI/CD
- B. DevSecOps
- C. GitOps
- D. DevOps

Answer: C

Explanation:

<https://fluxcd.io/>

Graphical user interface, text, application Description automatically generated

**Flux provides
GitOps for both
apps and
infrastructure**

Flux and Flagger deploy apps with canaries, feature flags, and A/B rollouts. Flux can also manage any Kubernetes resource. Infrastructure and workload dependency management is built in.

**Just push to Git
and Flux does the
rest**

Flux enables application deployment (CD) and (with the help of Flagger) progressive delivery (PD) through automatic reconciliation. Flux can even push back to Git for you with automated container image updates to Git (image scanning and patching).

NEW QUESTION 52


How does service logical group set of pods?

- A. Using hostname
- B. Using label and selectors
- C. Using IP address

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/services-networking/service/> Diagram Description automatically generated



NEW QUESTION 54

What is the use of labels in Kubernetes?

- A. All of the options
- B. It is used to assign annotation to an object
- C. It is used to assign key-value pair to an object
- D. It is used to assign a name to an object.

Answer: C

Explanation:

<https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/> Text Description automatically generated

Labels and Selectors

Labels are key/value pairs that are attached to objects, such as pods.

Labels are intended to be used to specify identifying attributes of objects that are meaningful and relevant to users, but do not directly imply semantics to the core system. Labels can be used to organize and to select subsets of objects. Labels can be attached to objects at creation time and subsequently added and modified at any time.

Each object can have a set of key/value labels defined. Each Key must be unique for a given object.

NEW QUESTION 59

What makes cloud native technology so important?

- A. It makes data centric
- B. It strengthens team
- C. It removes roadblocks to innovation
- D. It helps gather software requirements
- E. It makes operational centric

Answer: C

Explanation:

<https://github.com/cncf/foundation/blob/main/charter.md>

Graphical user interface, text, application Description automatically generated

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.

NEW QUESTION 63

Which of the following is an example of vertical scaling?

- A. Using cluster autoscaler
- B. Adding more resources (memory and/or cpu) to a kubernetes node
- C. Adding more nodes to kubernetes cluster
- D. Adding more replica pods to a deployment

Answer: B

Explanation:

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

Horizontal scaling means that the response to increased load is to deploy more Pods. This is different from *vertical* scaling, which for Kubernetes would mean assigning more resources (for example: memory or CPU) to the Pods that are already running for the workload.

NEW QUESTION 64

Which is not a service type in Kubernetes?

- A. ClusterIP
- B. NodePort
- C. Ingress
- D. LoadBalancer
- E. ExternalName

Answer: C

Explanation:

<https://kubernetes.io/docs/tutorials/kubernetes-basics/expose/expose-intro/>

without a Service. Services allow your applications to receive traffic. Services can be exposed in different ways by specifying a type in the ServiceSpec:

- *ClusterIP* (default) - Exposes the Service on an internal IP in the cluster. This type makes the Service only reachable from within the cluster.
- *NodePort* - Exposes the Service on the same port of each selected Node in the cluster using NAT. Makes a Service accessible from outside the cluster using `<NodeIP>: <NodePort>` . Superset of ClusterIP.
- *LoadBalancer* - Creates an external load balancer in the current cloud (if supported) and assigns a fixed, external IP to the Service. Superset of NodePort.
- *ExternalName* - Maps the Service to the contents of the `externalName` field (e.g. `foo.bar.example.com`), by returning a CNAME record with its value. No proxying of any kind is set up. This type requires v1.7 or higher of `kube-dns` , or CoreDNS version 0.0.8 or higher.

More information about the different types of Services can be found in the [Using Source IP](#) tutorial. Also see [Connecting Applications with Services](#).

Text Description automatically generated

NEW QUESTION 69

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 71

What is the default service type in Kubernetes?

- A. CusterIP
- B. NodePort
- C. serviceType
- D. loadBalancer

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types> Graphical user interface, text, application, email Description automatically generated

Kubernetes `ServiceTypes` allow you to specify what kind of Service you want. The default is `ClusterIP`.

`Type` values and their behaviors are:

- `ClusterIP` : Exposes the Service on a cluster-internal IP. Choosing this value makes the Service only reachable from within the cluster. This is the default `ServiceType`.
- `NodePort`: Exposes the Service on each Node's IP at a static port (the `NodePort`). A `ClusterIP` Service, to which the `NodePort` Service routes, is automatically created. You'll be able to contact the `NodePort` Service, from outside the cluster, by requesting `<NodeIP>:<NodePort>`.
- `LoadBalancer`: Exposes the Service externally using a cloud provider's load balancer. `NodePort` and `ClusterIP` Services, to which the external load balancer routes, are automatically created.
- `ExternalName`: Maps the Service to the contents of the `externalName` field (e.g. `foo.bar.example.com`), by returning a `CNAME` record with its value. No proxying of any kind is set up.

NEW QUESTION 73

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 75

Which of the following factors does scheduling take into account when selecting a Node?

- A. How many replicas there are in a Deployment
- B. Services
- C. Resource requirements
- D. The number of existing Pods on a Node

Answer: C

Explanation:
Scheduling takes resource requirements into account in the form of resource requests.

NEW QUESTION 76

What is etcd used for in Kubernetes?

- A. Integration with cloud platforms
- B. Network routing for the cluster
- C. Kubernetes API security
- D. Backend object storage for the Kubernetes API

Answer: D

Explanation:
etcd serves as a distributed object store that backs the Kubernetes API.

NEW QUESTION 81

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 82

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 85

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 86

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



NEW QUESTION 89

Continuous delivery is .

- A. Manually deploying the code
- B. Coding, Building and Testing the code
- C. Automatically deploying code to [container or server] environment

Answer: C

NEW QUESTION 93

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 98

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.

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 101

What are default kubernetes namespaces?

- A. default, kube-public, kube-system, kube-node-lease
- B. kube-default, kube-public, kube-system, kube-node-lease
- C. default, kube-public, kube-systems, kube-node-lease
- D. default, kube-public, kube-system, kube-node-leases

Answer: A

Explanation:

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

You can list the current namespaces in a cluster using:

```
kubectl get namespace
```

NAME	STATUS	AGE
default	Active	1d
kube-node-lease	Active	1d
kube-public	Active	1d
kube-system	Active	1d

Kubernetes starts with four initial namespaces:

- `default` The default namespace for objects with no other namespace
- `kube-system` The namespace for objects created by the Kubernetes system
- `kube-public` This namespace is created automatically and is readable by all users (including those not authenticated). This namespace is mostly reserved for cluster usage, in case that some resources should be visible and readable publicly throughout the whole cluster. The public aspect of this namespace is only a convention, not a requirement.
- `kube-node-lease` This namespace holds [Lease](#) objects associated with each node. Node leases allow the kubelet to send [heartbeats](#) so that the control plane can detect node failure.

NEW QUESTION 106

What cloud-native construct does a kubernetes pod wrap?

- A. Container
- B. Virtual Machine (VM)
- C. side car process
- D. Docker image

Answer: A

Explanation:

Kubernetes is an orchestrator of containerized apps. However, containers must be wrapped in pods before they can be deployed on kubernetes.

NEW QUESTION 110

Which control plane component is responsible for scheduling pods?

- A. kube-proxy
- B. kube scheduler
- C. kubelet
- D. kube api-server

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/overview/components/>

Graphical user interface, text, application Description automatically generated

kube-scheduler

Control plane component that watches for newly created Pods with no assigned node, and selects a node for them to run on.

Factors taken into account for scheduling decisions include: individual and collective resource requirements, hardware/software/policy constraints, affinity and anti-affinity specifications, data locality, inter-workload interference, and deadlines.

NEW QUESTION 113

Which style of operations are preferred for K8S and cloud native applications?

- A. JSON
- B. Declarative
- C. Imperative

Answer: B

Explanation:

<https://kubernetes.io/docs/tasks/manage-kubernetes-objects/declarative-config/#trade-offs>

NEW QUESTION 115

What is FinOps?

- A. The first step in any cloud transformation
- B. Stage beyond DevOps or DevSecOps, where organization transition to serverless technologies
- C. Using data to make cost savings decisions about cloud usage
- D. Specialized cloud features used by financial industries (example: banks, insurance, etc)

Answer: C

Explanation:

<https://www.servicenow.com/products/it-asset-management/what-is-finops.html> Text, letter Description automatically generated

What is the origin of FinOps?

Unlike many modern tech movements, FinOps is not a single advancement or policy change pioneered by any specific company or organization; it's a natural evolution of technology management to account for on-demand cloud resources.

With the rise and proliferation of cloud computing in the new millennium, many companies began to see a shift from standard, traditional pricing to usage-based pricing models. And, while this allowed businesses to take a more cost effective approach to technology—paying only for the time and resources they used, rather than paying a set rate—it created a crisis for CFOs. After all, it's next to impossible to predict tool usage with any degree of accuracy, which can make budgeting an exercise in futility.

To address this issue, prevent runaway expenses, and promote business profitability, organizations around the world began to develop the concept of financial operations, (FinOps). This revolution was guided by respected technology companies around the world, first taking shape as cloud cost management, developing into cloud cost optimization, and then into cloud financial management.

Finally, taking inspiration from the success of DevOps, FinOps was born, bringing cross-functionality and agility to financial management of cloud technologies.

NEW QUESTION 117

What is a benefits of Kubernetes federation?

- A. Avoids scalability limits on pods and nodes
- B. Creates highly available clusters in different regions
- C. Low latency

Answer: ABC

NEW QUESTION 122

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 126

In distributed system tracing, is the term used to refer to a request as it passes through a single component of the distributed system?

- A. Log
- B. Span
- C. Trace
- D. Bucket

Answer: B

Explanation:

https://www.splunk.com/en_us/data-insider/what-is-distributed-tracing.html Text, letter Description automatically generated

How does distributed tracing work?

To quickly grasp how distributed tracing works, it's best to look at how it handles a single request. Tracing starts the moment an end user interacts with an application. When the user sends an initial request — an HTTP request, to use a common example — it is assigned a unique trace ID. As the request moves through the host system, every operation performed on it (called a “span” or a “child span”) is tagged with that first request's trace ID, as well as its own unique ID, plus the ID of the operation that originally generated the current request (called the “parent span”).

Each span is a single step on the request's journey and is encoded with important data relating to the microservice process that is performing that operation. These include:

- The service name and address of the process handling the request.
- Logs and events that provide context about the process's activity.
- Tags to query and filter requests by session ID, database host, HTTP method, and other identifiers.
- Detailed stack traces and error messages in the event of a failure.

A distributed tracing tool like Zipkin or Jaeger (both of which we will explore in more detail in a bit) can correlate the data from all the spans and format them into visualizations that are available on request through a web interface.

Now think of a popular online video game with millions of users, the epitome of a modern microservices-driven app. It must track each end user's location, each interaction with other players and the environment, every item the player acquires, end time, and a host of other in-game data. Keeping the game running smoothly would be unthinkable with traditional tracing methods. But distributed request tracing makes it possible.

NEW QUESTION 130

Which of the following container runtime is planned to be deprecated in Kubernetes 1.20 and high-er?

- A. cri-o
- B. None of the options
- C. docker
- D. podman
- E. containerd

Answer: C

Explanation:

<https://kubernetes.io/blog/2020/12/02/dont-panic-kubernetes-and-docker/>

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

Wednesday, December 02, 2020

Update: *Kubernetes support for Docker via `dockershim` is now removed. For more information, read the [removal FAQ](#). You can also discuss the deprecation via a dedicated [GitHub issue](#).*

Authors: Jorge Castro, Duffie Cooley, Kat Cosgrove, Justin Garrison, Noah Kantrowitz, Bob Killen, Rey Lejano, Dan “POP” Papandrea, Jeffrey Sica, Davanum “Dims” Srinivas

Kubernetes is [deprecating Docker](#) as a container runtime after v1.20.

You do not need to panic. It's not as dramatic as it sounds.

TL;DR Docker as an underlying runtime is being deprecated in favor of runtimes that use the [Container Runtime Interface \(CRI\)](#) created for Kubernetes. Docker-produced images will continue to work in your cluster with all runtimes, as they always have.

NEW QUESTION 134

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 137

The 4C's of Cloud Native security

- A. Chroot, Compute, Cluster and Container
- B. Cluster, Cloud, Compute, and Containers
- C. Code, Containers, Compute, and Cloud

D. Cloud, Clusters, Containers, and Code

Answer: D

Explanation:

<https://kubernetes.io/docs/concepts/security/overview/>

NEW QUESTION 142

Fluentd is the leading project in the CNCF space for logging?


A. TRUE

B. FALSE

Answer: A

Explanation:

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



NEW QUESTION 147

Which access control component of Kubernetes is responsible for authorization and decides what requestor is allowed to do?

A. Service Account

B. Role-based access control 'RBAC'

C. Deployment

Answer: B

Explanation:

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

Authorization Modes

The Kubernetes API server may authorize a request using one of several authorization modes:

- **Node** - A special-purpose authorization mode that grants permissions to kubelets based on the pods they are scheduled to run. To learn more about using the Node authorization mode, see [Node Authorization](#).
- **ABAC** - Attribute-based access control (ABAC) defines an access control paradigm whereby access rights are granted to users through the use of policies which combine attributes together. The policies can use any type of attributes (user attributes, resource attributes, object, environment attributes, etc). To learn more about using the ABAC mode, see [ABAC Mode](#).
- **RBAC** - Role-based access control (RBAC) is a method of regulating access to computer or network resources based on the roles of individual users within an enterprise. In this context, access is the ability of an individual user to perform a specific task, such as view, create, or modify a file. To learn more about using the RBAC mode, see [RBAC Mode](#)
 - When specified RBAC (Role-Based Access Control) uses the `rbac.authorization.k8s.io` API group to drive authorization decisions, allowing admins to dynamically configure permission policies through the Kubernetes API.
 - To enable RBAC, start the apiserver with `--authorization-mode=RBAC`.

NEW QUESTION 152

Various Container Orchestrator Systems (COS)?

- A. Apache Mesos
- B. None of the options
- C. Docker Swarm
- D. Kubernetes

Answer: ACD

NEW QUESTION 155

'kubectl delete -n my-ns po,svc --all' will delete pods and services including uninitialized ones in the namespace 'my-ns'

- A. FALSE
- B. TRUE

Answer: B

Explanation:

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

IMPORTANT: Force deleting pods does not wait for confirmation that the pod's processes have been terminated, which can leave those processes running until the node detects the deletion and completes graceful deletion. If your processes use shared storage or talk to a remote API and depend on the name of the pod to identify themselves, force deleting those pods may result in multiple processes running on different machines using the same identification which may lead to data corruption or inconsistency. Only force delete pods when you are sure the pod is terminated, or if your application can tolerate multiple copies of the same pod running at once. Also, if you force delete pods, the scheduler may place new pods on those nodes before the node has released those resources and causing those pods to be evicted immediately.

Note that the delete command does NOT do resource version checks, so if someone submits an update to a resource right when you submit a delete, their update will be lost along with the rest of the resource.

After a CustomResourceDefinition is deleted, invalidation of discovery cache may take up to 10 minutes. If you don't want to wait, you might want to run "kubectl api-resources" to refresh the discovery cache.

Usage

```
$ kubectl delete [{-f FILENAME} | {-k DIRECTORY} | TYPE [(NAME | -l label | --all))]
```

Delete a pod based on the type and name in the JSON passed into stdin
<code>cat pod.json kubectl delete -f -</code>
Delete pods and services with same names "baz" and "foo"
<code>kubectl delete pod,service baz foo</code>
Delete pods and services with label name=myLabel
<code>kubectl delete pods,services -l name=myLabel</code>
Delete a pod with minimal delay
<code>kubectl delete pod foo --now</code>
Force delete a pod on a dead node
<code>kubectl delete pod foo --force</code>
Delete all pods
<code>kubectl delete pods --all</code>

NEW QUESTION 157

Which of the following computing model doesn't require you to provision infrastructure?

- A. None of the above
- B. Bare Metal
- C. Compute Engine
- D. Virtual Machines
- E. Serverless

Answer: E

NEW QUESTION 159

In Kubernetes, what is considered the primary cluster data source?

- A. etcd (pronounce: esty-d)
- B. api server
- C. kubelet
- D. scheduler

Answer: A

Explanation:

<https://kubernetes.io/docs/concepts/overview/components/#etcd>

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

etcd

Consistent and highly-available key value store used as Kubernetes' backing store for all cluster data.

If your Kubernetes cluster uses etcd as its backing store, make sure you have a **back up** plan for those data.

You can find in-depth information about etcd in the official [documentation](#).

NEW QUESTION 161

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your KCNA Exam with Our Prep Materials Via below:

<https://www.certleader.com/KCNA-dumps.html>