Linux-Foundation

Exam Questions KCNA

Kubernetes and Cloud Native Associate (KCNA)

NEW QUESTION 1
How to create deployment name app-dep, image=nginx, and replicas 5 using imperative command?

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

Answer: B

Explanation:
https://kubernetes.io/docs/reference/generated/kubectl/kubectl-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

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

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 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

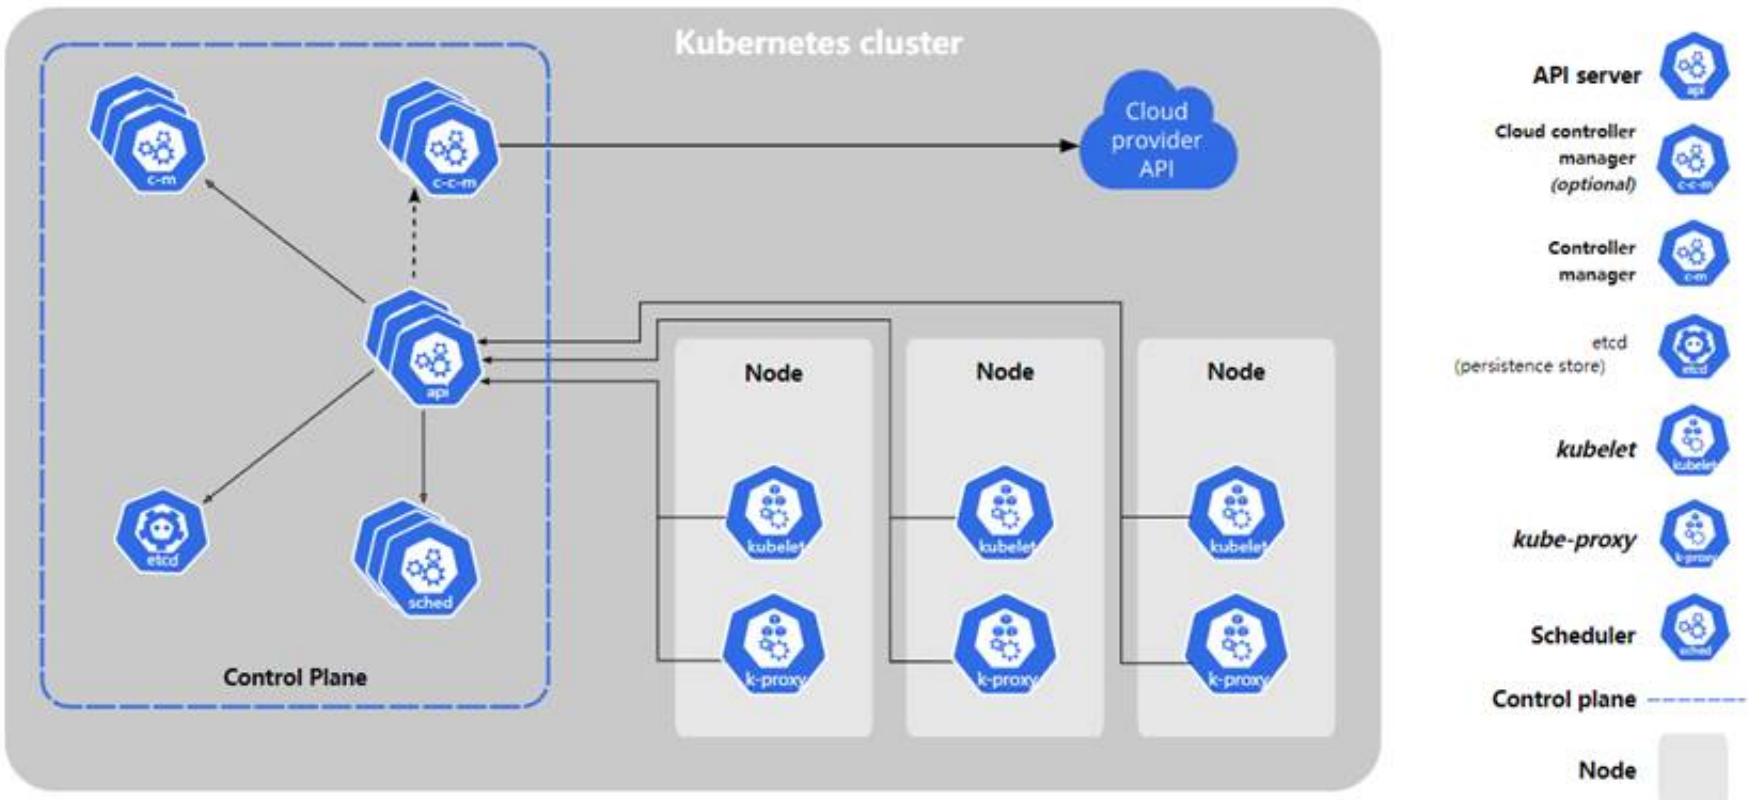

NEW QUESTION 4
To specify a Kubernetes object which language is used?

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

NEW QUESTION 5
There are three Nodes in a cluster, and want to run exactly one replica of a Pod on each Node. Pre-fer to automatically create a replica on any new Nodes when they are added. Which Kubernetes re-source 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

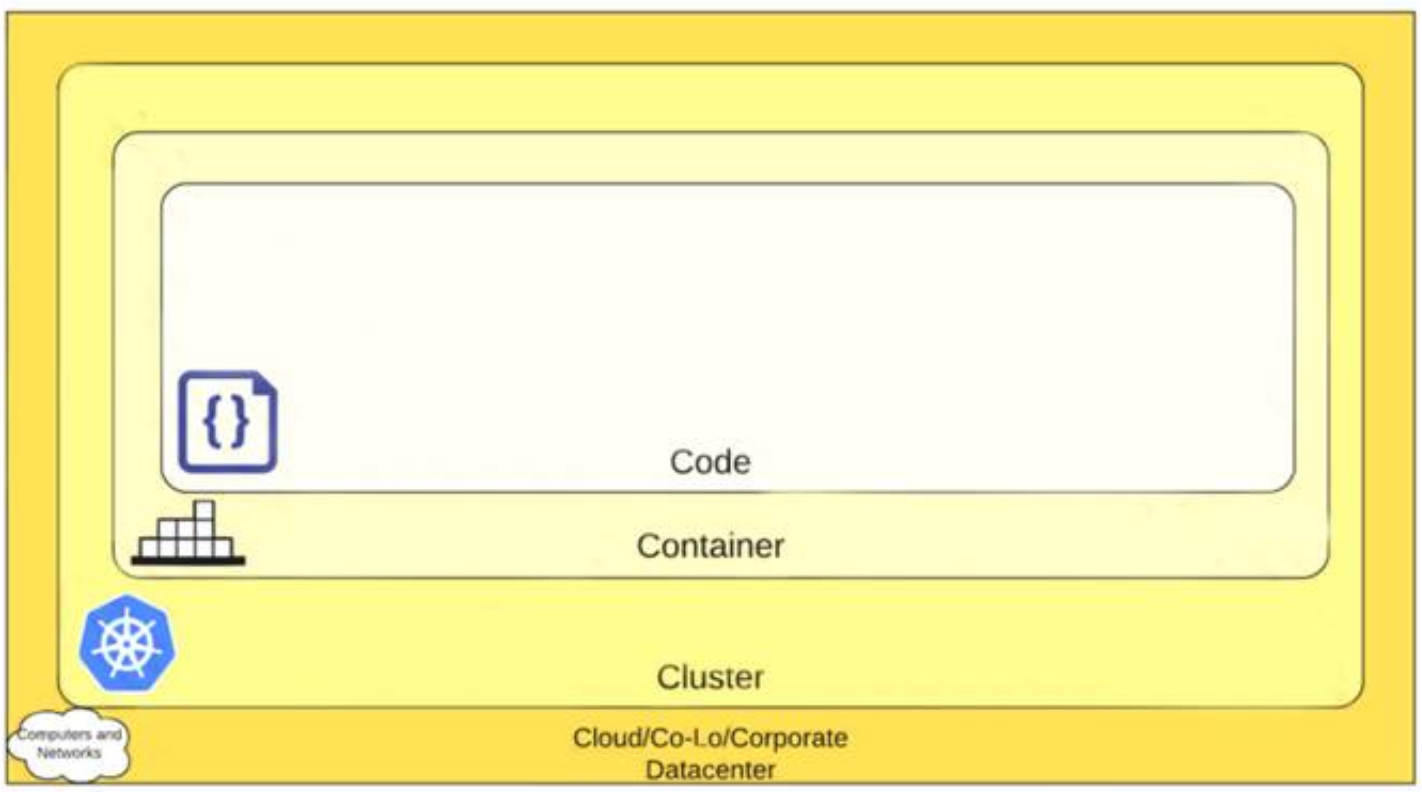

NEW QUESTION 6
Which kubernetes resource type allows defining which pods are isolated when it comes to network-ing?

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 7
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 8
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 9
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

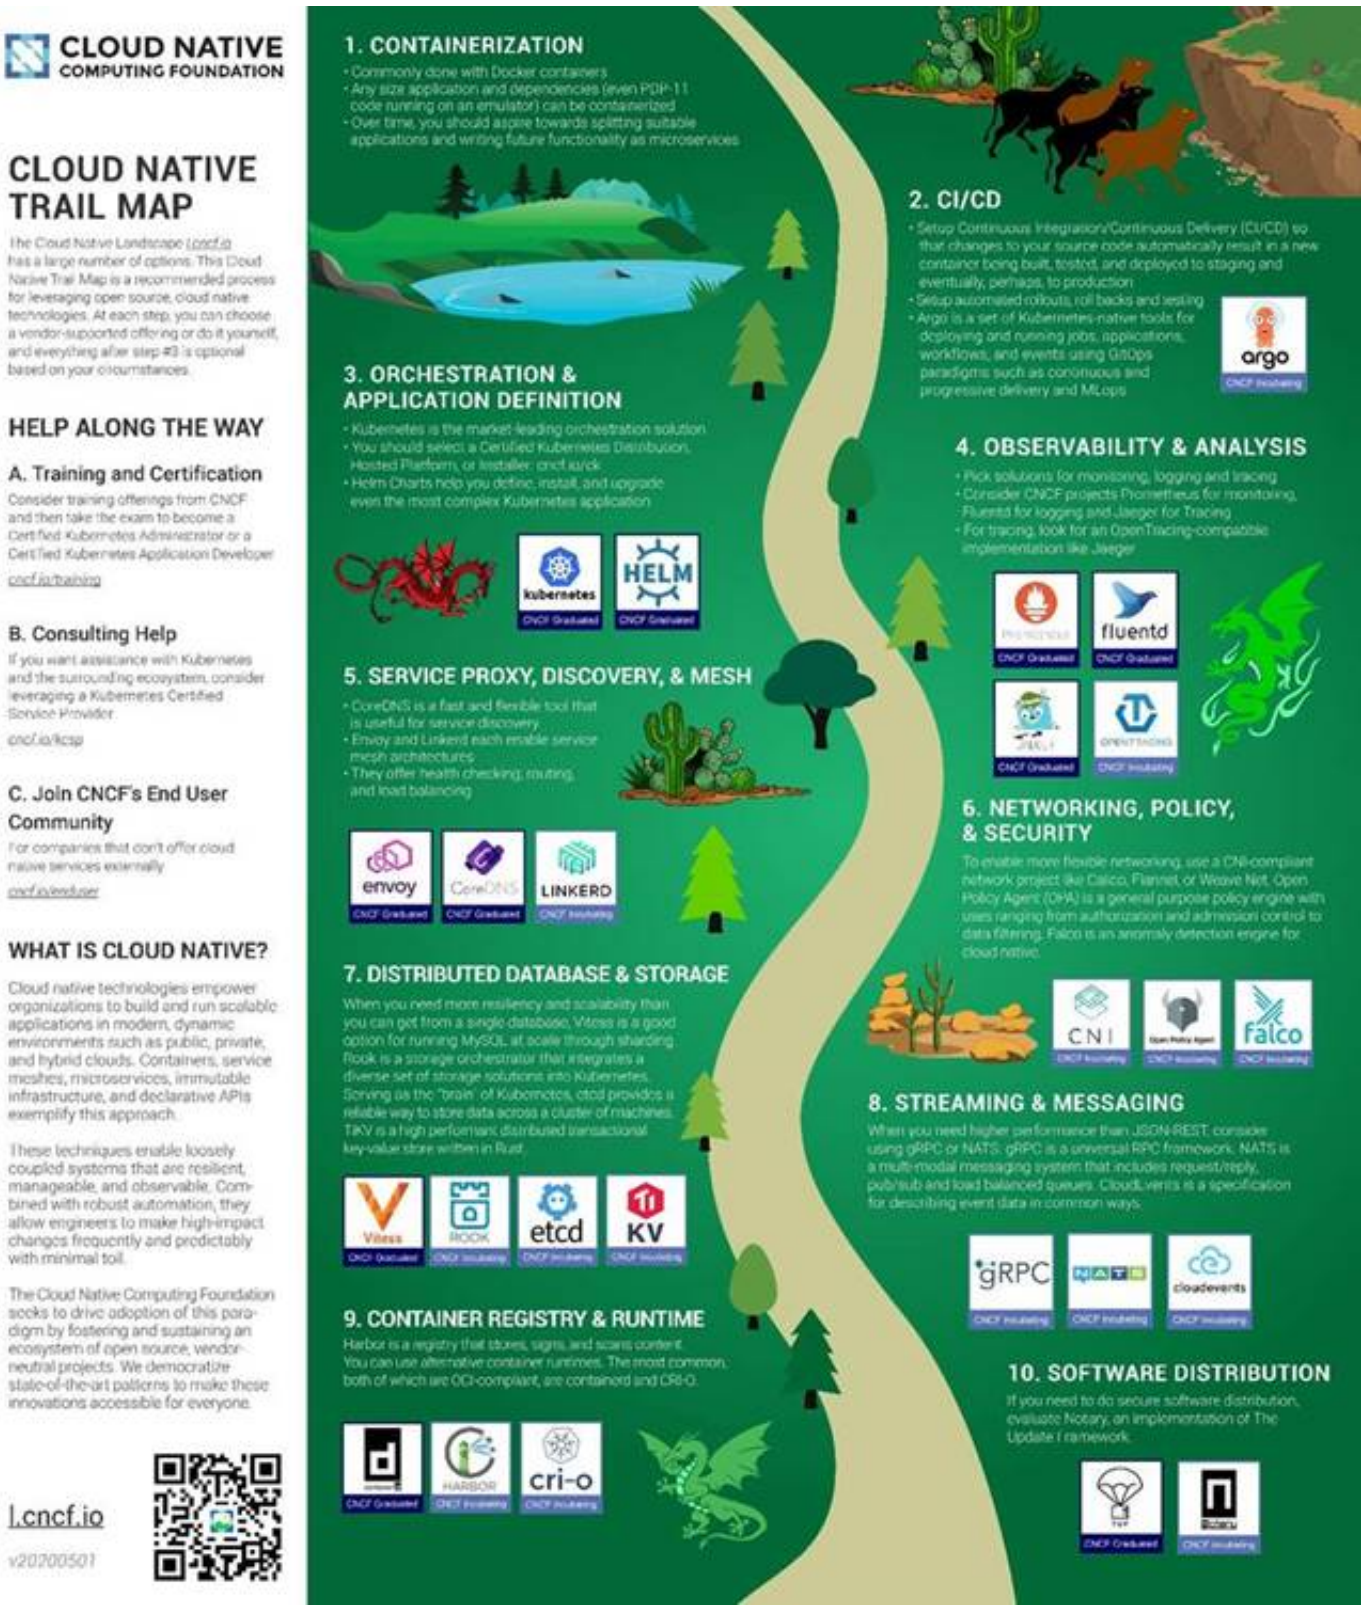

NEW QUESTION 10
Which prometheus metric type represents a single number value that can increase and decrease over time?

A. Gauge
B. Histogram
C. Summary
D. Counter

Answer: A

Explanation:
https://prometheus.io/docs/concepts/metric_types/#gauge Graphical user interface, text Description automatically generated

Gauge

A *gauge* is a metric that represents a single numerical value that can arbitrarily go up and down.

Gauges are typically used for measured values like temperatures or current memory usage, but also "counts" that can go up and down, like the number of concurrent requests.

NEW QUESTION 10
Which of the following best describes a cloud-native app?

A. An application where all logic is coded into a single large binary.
B. An application that publishes an HTTPS web front-end.
C. An application that takes advantages of cloud computing fromworks and their loosely coupled cloud services.
D. An application that leverages services that are native to public cloud platforms such as Azure, GCP,and/or AW

Answer: C

Explanation:
Cloud-native apps leverage cloud computing frameworks and tend to be microservices based, where individual components of the app are coded as individual.

NEW QUESTION 13
To run a startup task before a Pod's container starts up. What Kubernetes feature can help you ac-complish this?

A. Init container
B. Sidecar container
C. Startup probe
D. DaemonSet

Answer: A

Explanation:
https://kubernetes.io/docs/concepts/workloads/pods/init-containers/ Text, letter Description automatically generated

A Pod can have multiple containers running apps within it, but it can also have one or more init containers, which are run before the app containers are started.

Init containers are exactly like regular containers, except:

- Init containers always run to completion.
- Each init container must complete successfully before the next one starts.

If a Pod's init container fails, the kubelet repeatedly restarts that init container until it succeeds. However, if the Pod has a `restartPolicy` of Never, and an init container fails during startup of that Pod, Kubernetes treats the overall Pod as failed.

To specify an init container for a Pod, add the `initContainers` field into the Pod specification, as an array of `container` items (similar to the app `containers` field and its contents). See Container in the API reference for more details.

NEW QUESTION 17
What do control groups provide when it come to containers

A. Permission
B. Image Storage
C. Isolation
D. Logging

Answer: C

Explanation:
Text Description automatically generated

What is the use of kernel control groups in container technology?

A control group (cgroup) is a Linux kernel feature that **limits, accounts for, and isolates the resource usage (CPU, memory, disk I/O, network, and so on) of a collection of processes.** Jul 21, 2021

NEW QUESTION 22
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

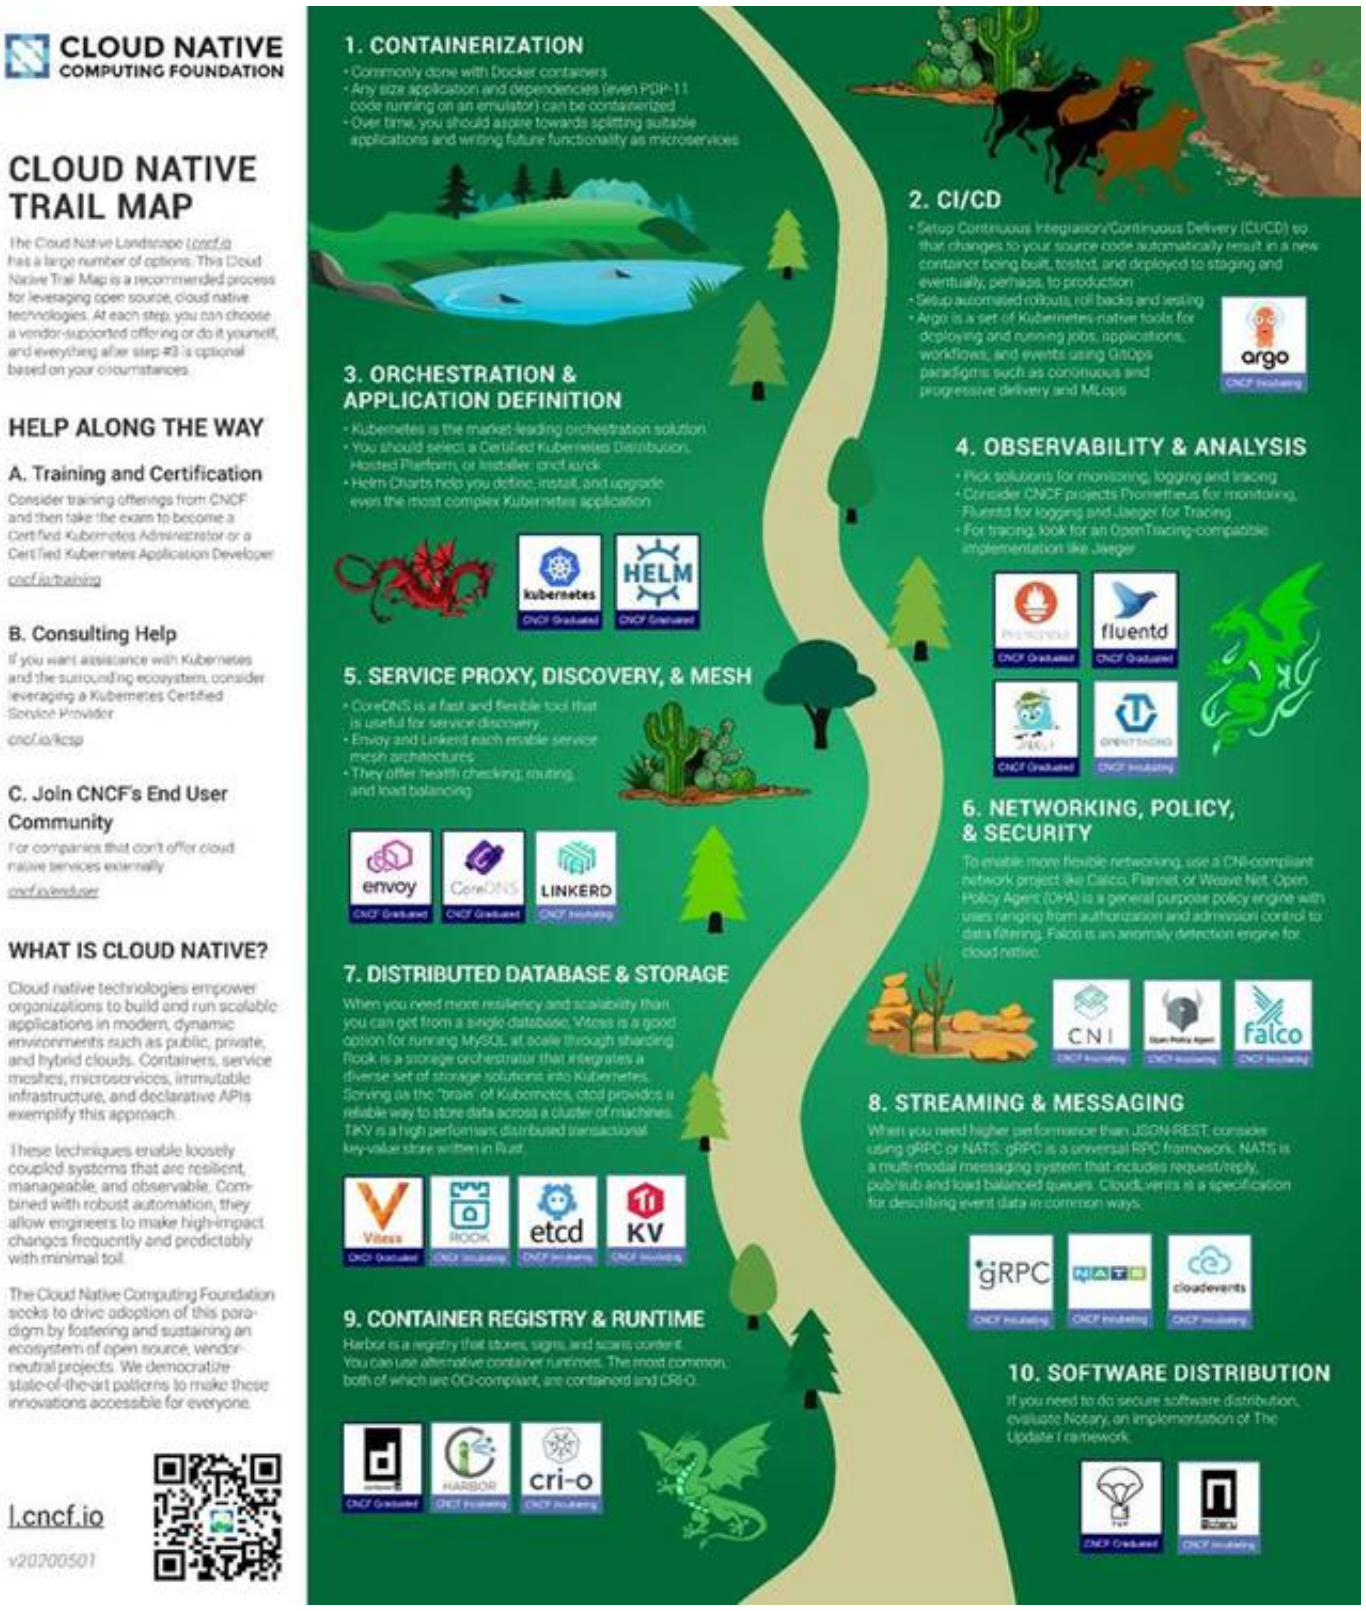

NEW QUESTION 25
What is OPA?

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

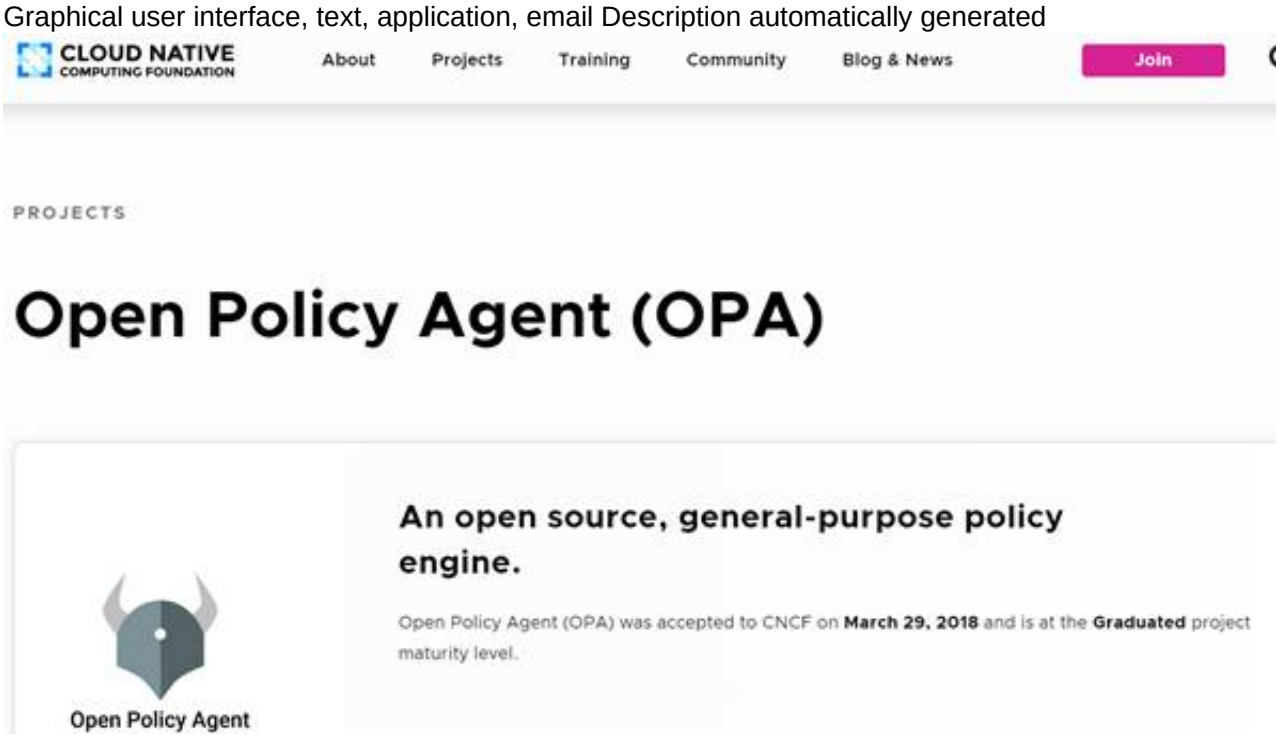

NEW QUESTION 27
What is the command to list all the available objects in your Kubernetes cluster?

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

NEW QUESTION 31
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

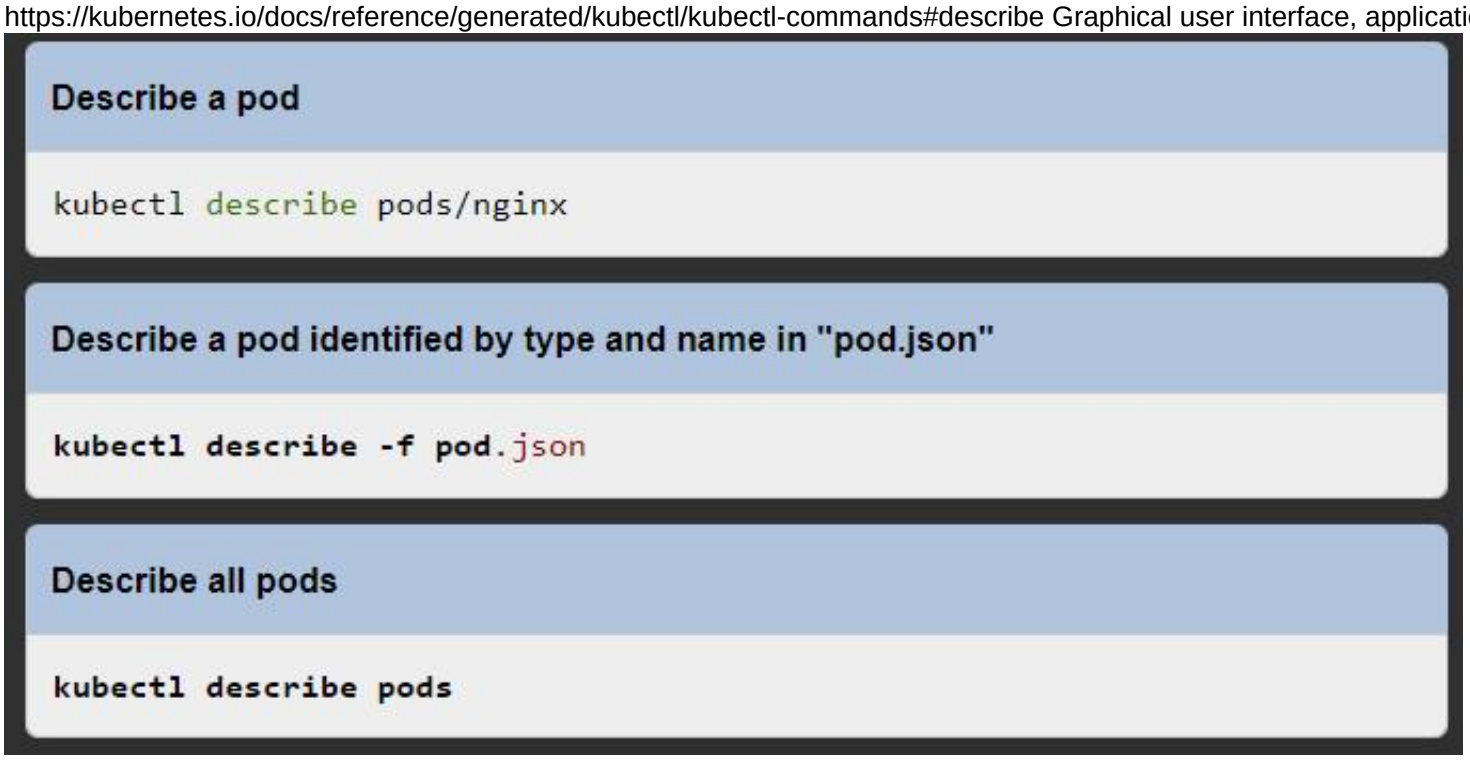

NEW QUESTION 35
Which of the following is not the Kubernetes AutoScaling Strategy?

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 40
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-namspaces
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 44
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
What do GitOps tools do in kubernetes?

A. They allow us to make changes to a kubernetes cluster using a Git repository
B. They manage the source code of kubernetes itself
C. They allow us to store software code in Git
D. They allows us to store container images in repositories

Answer: A

Explanation:
https://fluxcd.io/docs/components/
Diagram Description automatically generated

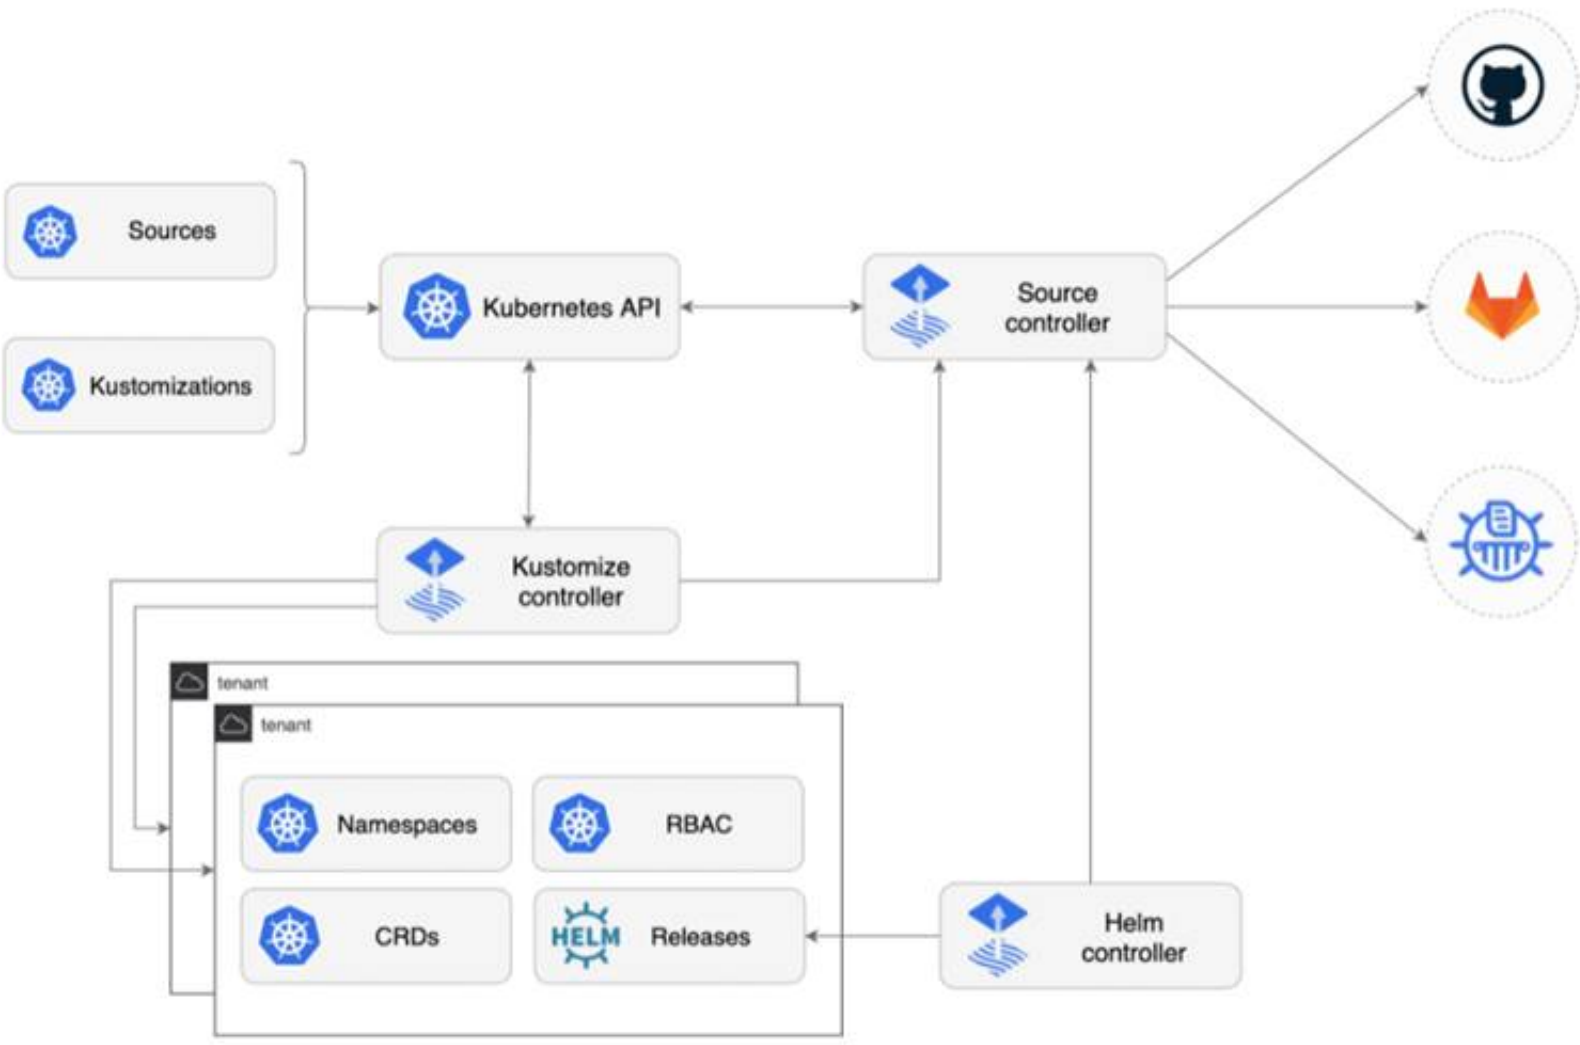

NEW QUESTION 51
Flux is built using which toolkit?

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 56
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

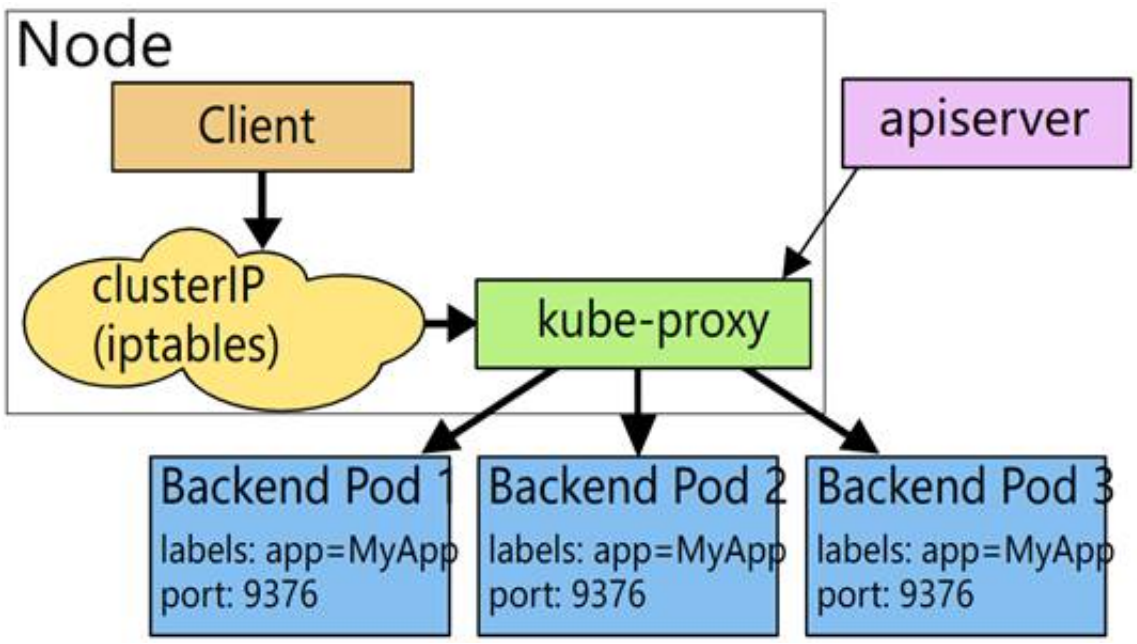

NEW QUESTION 59
An application that is nearing its usage limit. To increase the amount of users it can handle, you allo-cate additional memory resources to each instance of the application. What type of scaling is this?

A. Horizontal Scaling
B. Cluster Autoscaling
C. Recursive Scaling
D. Vertical Scaling

Answer: D

Explanation:
Graphical user interface, diagram Description automatically generated

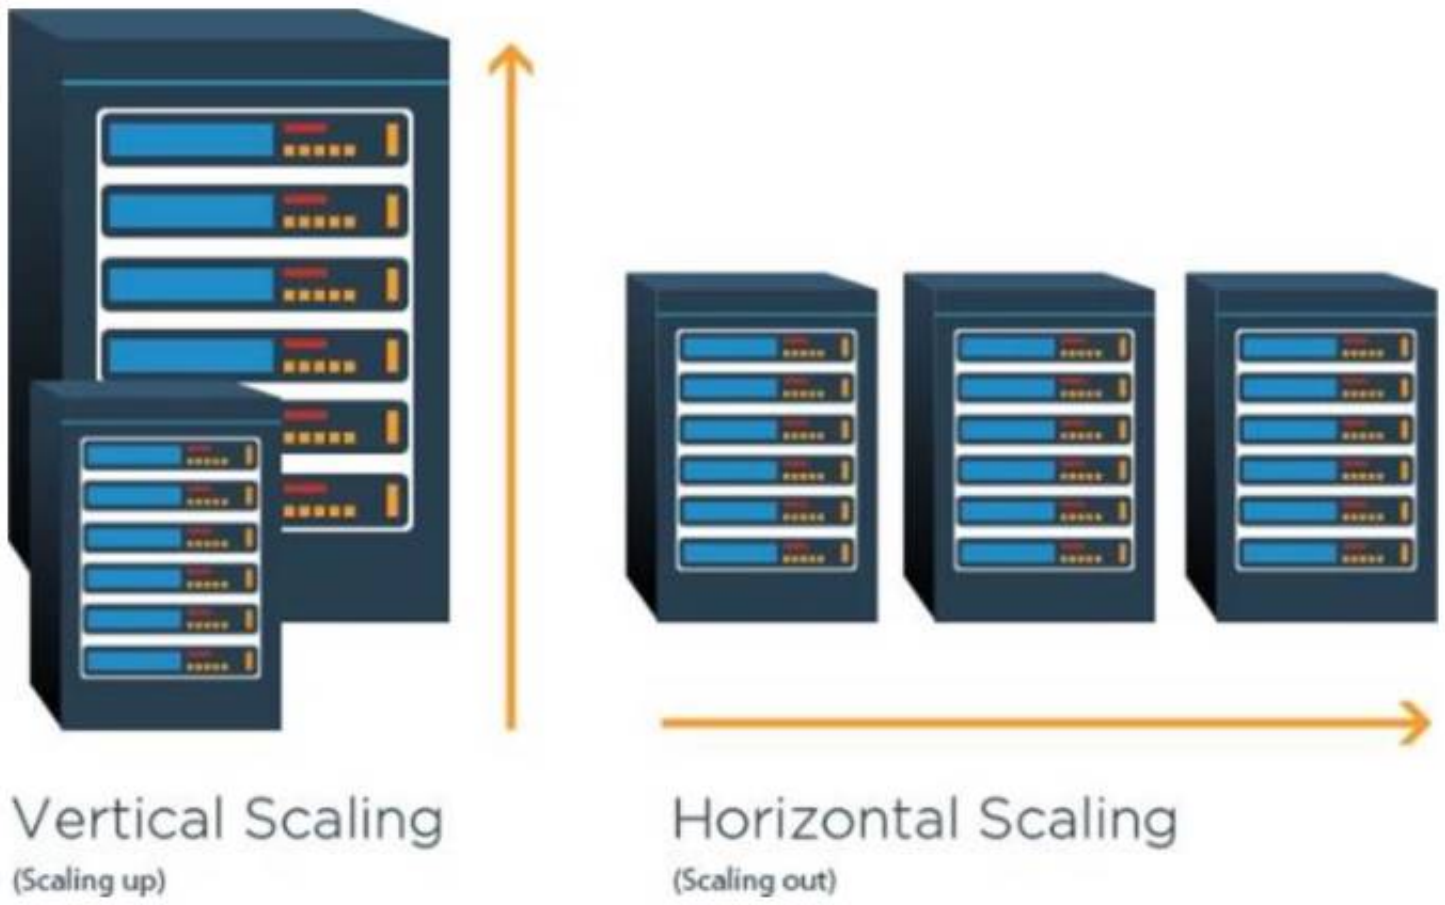

NEW QUESTION 63
What makes cloud native technology so important?

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 66
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 70
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/

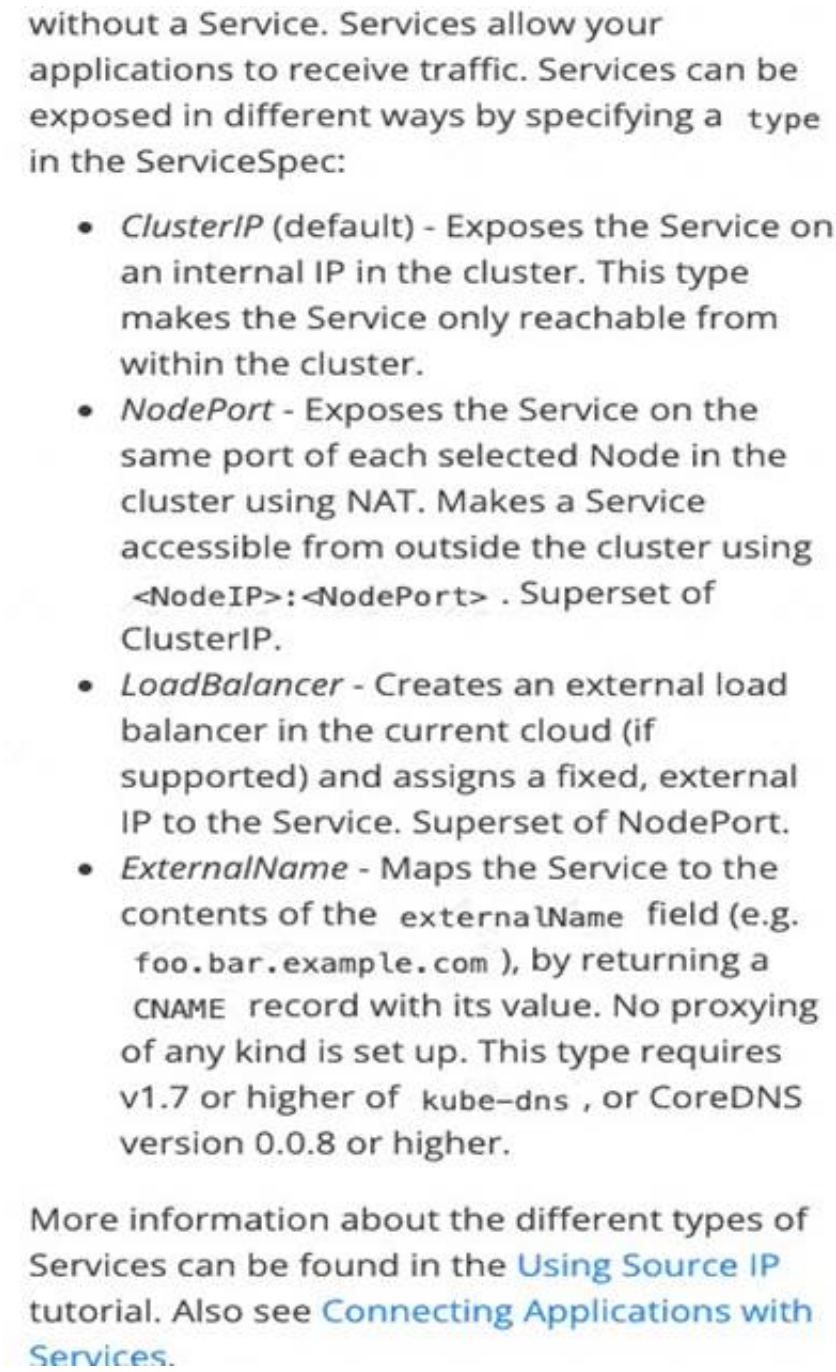

Text Description automatically generated

NEW QUESTION 74
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

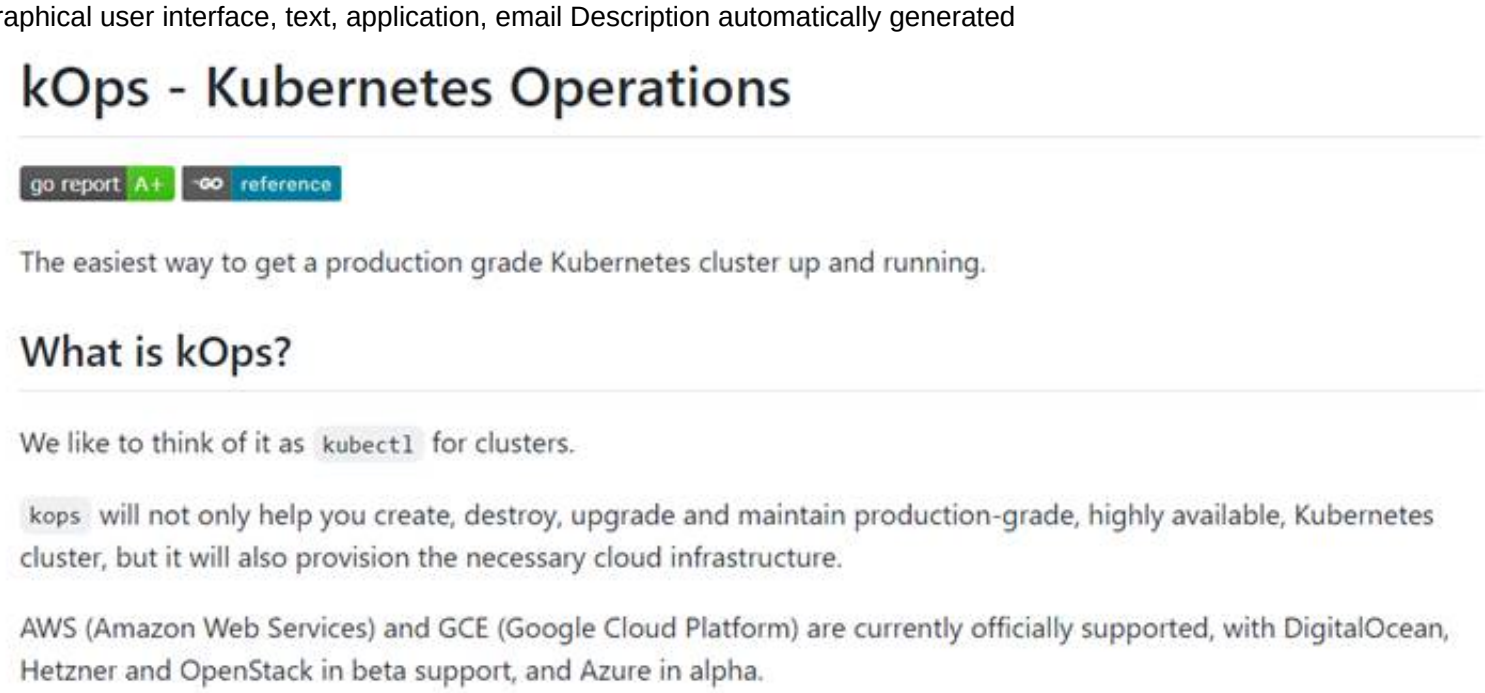

NEW QUESTION 75
Have a pod 'hello' and a container in that pod 'green'. Which of the following commands would get the logs for that container?

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

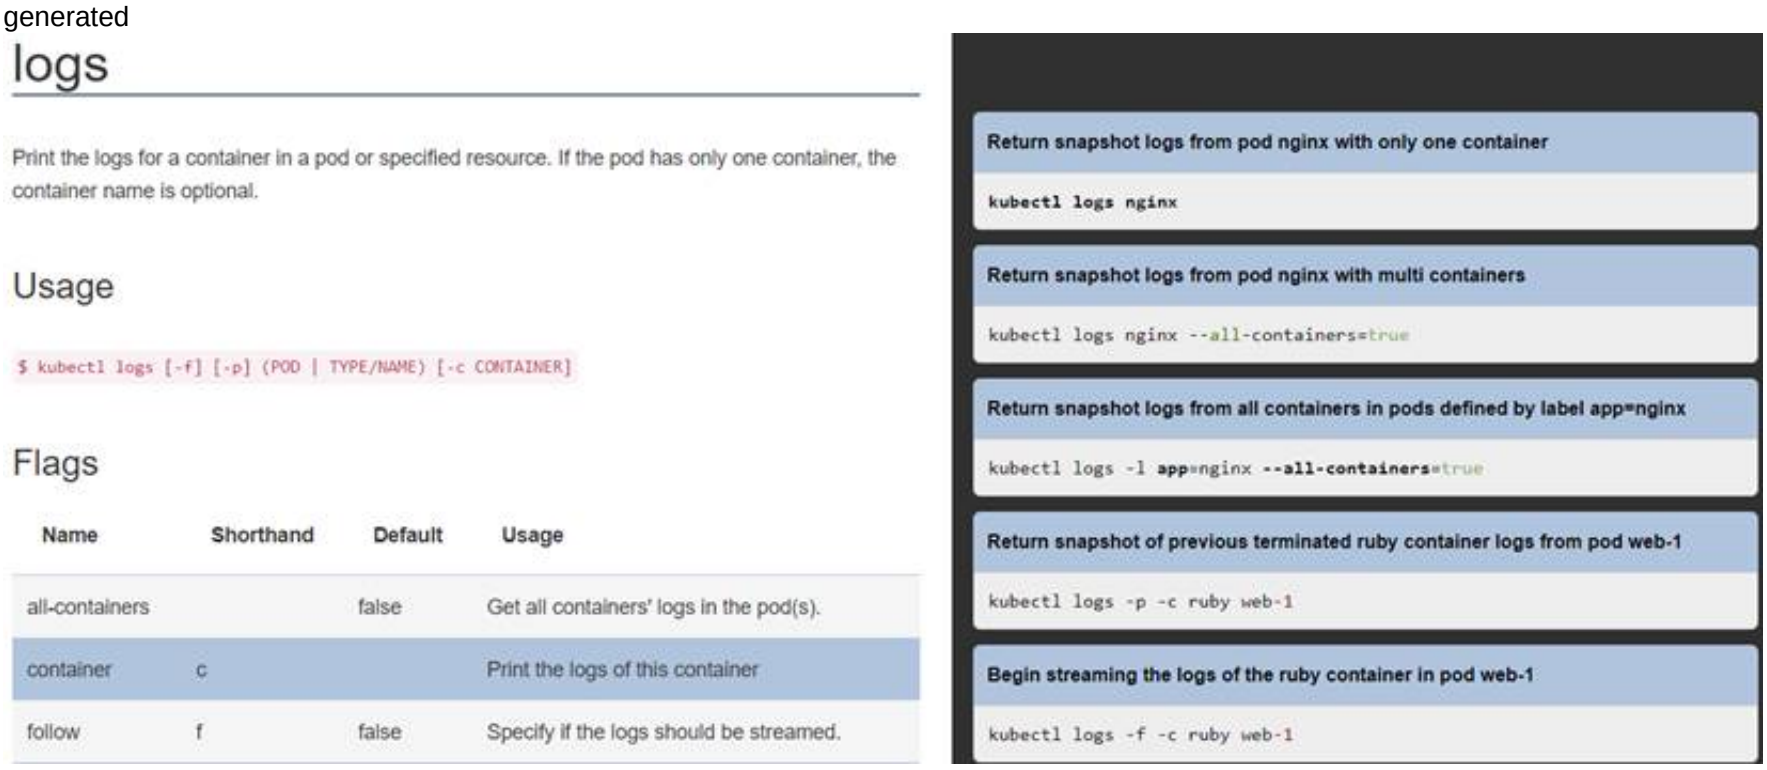

NEW QUESTION 80
Open Container Initiative set container standards for

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 82
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 87
What is autoscaling?

A. Automatically measuring resource usage
B. Automatically assigning workloads to nodes in a cluster
C. Automatically repairing broken application instances
D. Automatically adding or removing compute resources as needed

Answer: D

Explanation:
https://kubernetes.io/blog/2016/07/autoscaling-in-kubernetes/

Autoscaling means automatically scaling up or down in response to real-time usage data.

NEW QUESTION 90
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 92
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 97
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 99
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

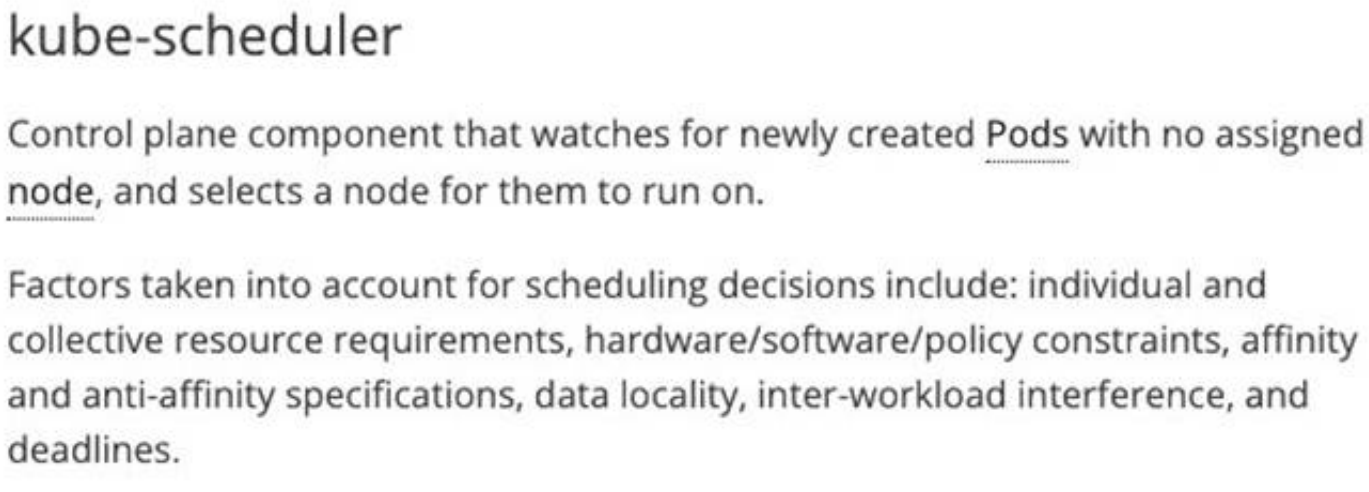

NEW QUESTION 102
Notary and the update framework leading security projects in CNCF

A. TRUE
B. FALSE

Answer: A

Explanation:
https://github.com/cncf/landscape#trail-map
A picture containing timeline Description automatically generated

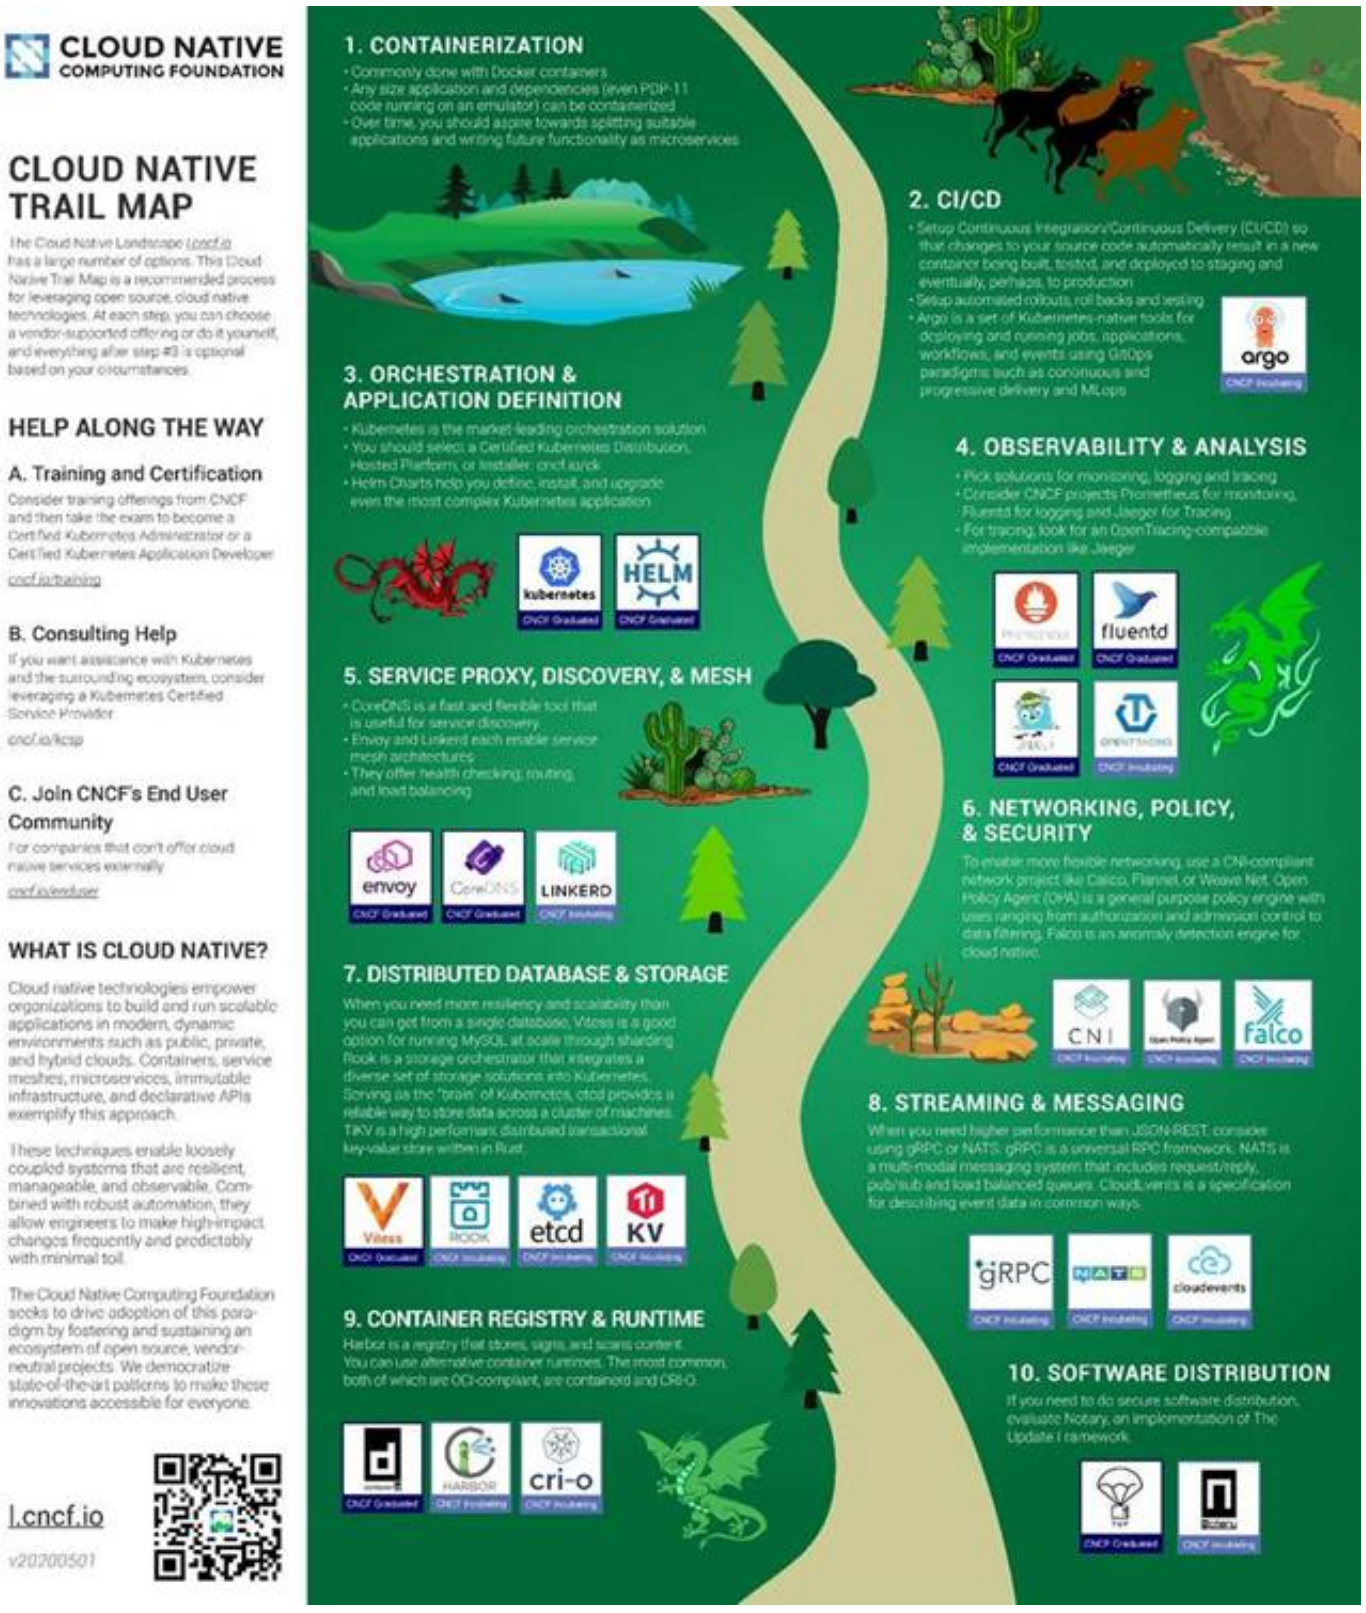

NEW QUESTION 104
What is horizontal scaling?

A. Creating a Deployment
B. Adding resources to existing apps and servers
C. Moving workloads from one server to another
D. Adding additional replicas of apps and servers

Answer: D

Explanation:
https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/ Text, letter Description automatically generated

In Kubernetes, a *HorizontalPodAutoscaler* automatically updates a workload resource (such as a Deployment or StatefulSet), with the aim of automatically scaling the workload to match demand.

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.

If the load decreases, and the number of Pods is above the configured minimum, the HorizontalPodAutoscaler instructs the workload resource (the Deployment, StatefulSet, or other similar resource) to scale back down.

Horizontal pod autoscaling does not apply to objects that can't be scaled (for example: a DaemonSet.)

The HorizontalPodAutoscaler is implemented as a Kubernetes API resource and a controller. The resource determines the behavior of the controller. The horizontal pod autoscaling controller, running within the Kubernetes control plane, periodically adjusts the desired scale of its target (for example, a Deployment) to match observed metrics such as average CPU utilization, average memory utilization, or any other custom metric you specify.

NEW QUESTION 108
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 111
What are container runtimes with Kubernetes?

A. CRI-O
B. lxd
C. containerd
D. Dockershim

Answer: AC

Explanation:
https://kubernetes.io/docs/setup/production-environment/container-runtimes/ Graphical user interface, text, application, email Description automatically generated

Container Runtimes

Note: Dockershim has been removed from the Kubernetes project as of release 1.24. Read the Dockershim Removal FAQ for further details.

You need to install a container runtime into each node in the cluster so that Pods can run there. This page outlines what is involved and describes related tasks for setting up nodes.

Kubernetes 1.25 requires that you use a runtime that conforms with the Container Runtime Interface (CRI).

See CRI version support for more information.

This page provides an outline of how to use several common container runtimes with Kubernetes.

- containerd
- CRI-O
- Docker Engine
- Mirantis Container Runtime

Note:
Kubernetes releases before v1.24 included a direct integration with Docker Engine, using a component named *dockershim*. That special direct integration is no longer part of Kubernetes (this removal was announced as part of the v1.20 release). You can read Check whether Dockershim removal affects you to understand how this removal might affect you. To learn about migrating from using dockershim, see Migrating from dockershim.

If you are running a version of Kubernetes other than v1.25, check the documentation for that version.

NEW QUESTION 114
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 117
Various Container Orchestrator Systems (COS)?

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

Answer: ACD

NEW QUESTION 121
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

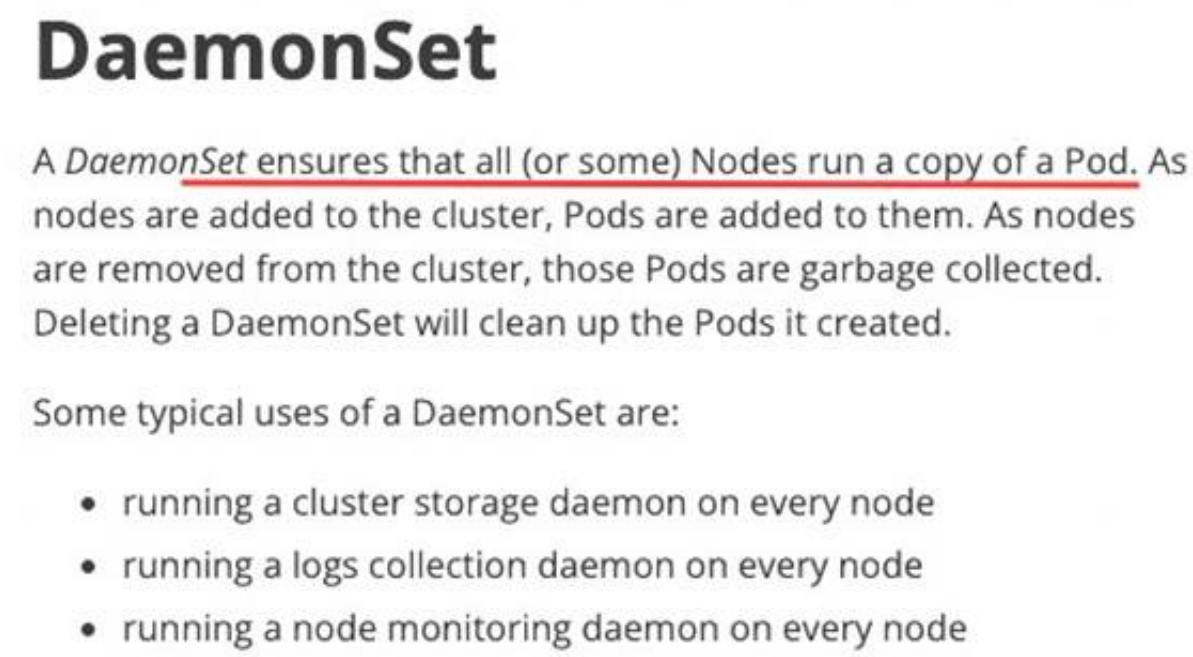

NEW QUESTION 122
What is the command used to login to the pod?

A. kubectl login
B. kubectl list
C. kubectl exec
D. kubectl get

Answer: C

Explanation:
https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#exec

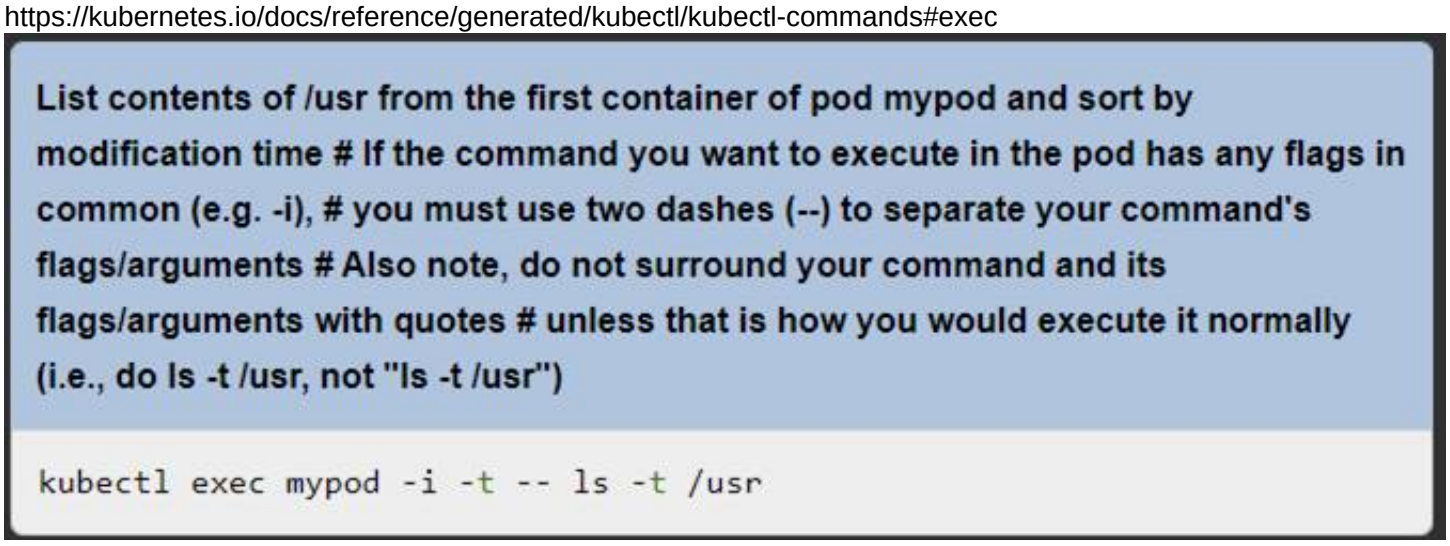

Text Description automatically generated

NEW QUESTION 123
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

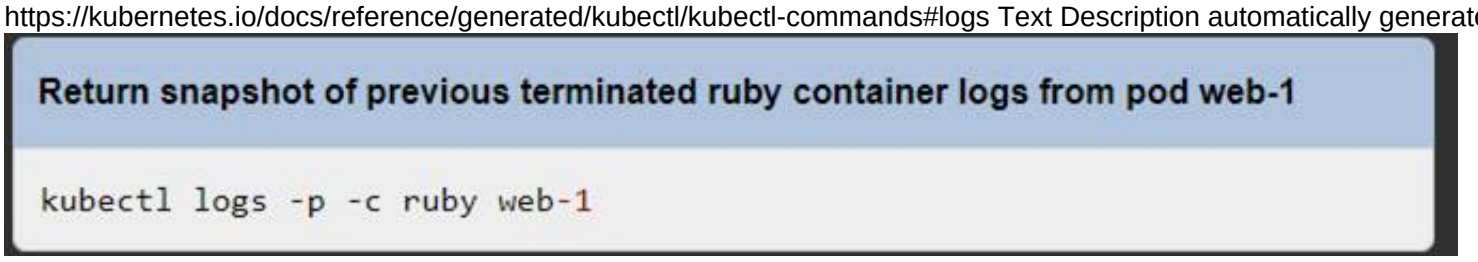

NEW QUESTION 127
Which kubernetes object do deployments use behind the scenes when they need to scale pods?

A. POD
B. Deployment
C. Horizontal pod autoscaler
D. Api Scheduler
E. Replicasets

Answer: E

Explanation:
https://kubernetes.io/docs/concepts/workloads/controllers/replicaset/
Graphical user interface, text, application, email Description automatically generated

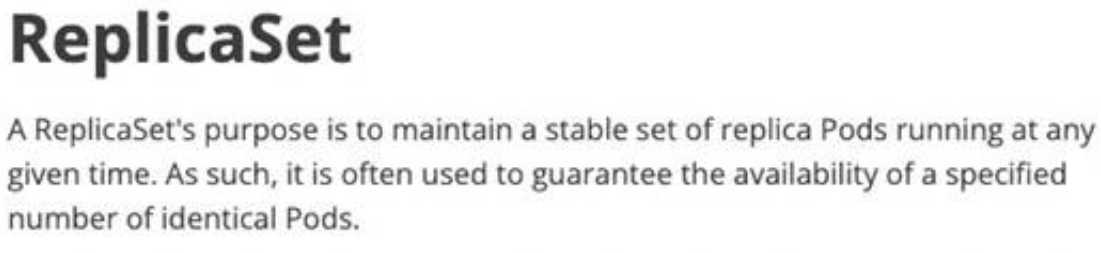

NEW QUESTION 129
'kubectl delete -n my-ns po,svc --all' will delete pods and services including uninitialized ones in the namespace 'my-ns'

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

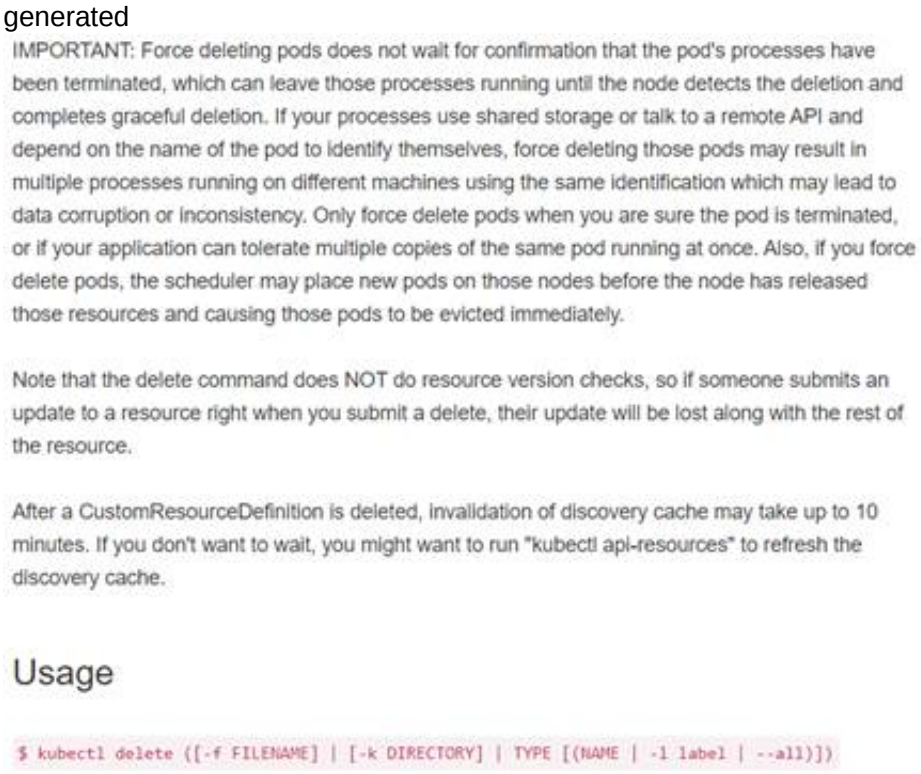

NEW QUESTION 130
In Kubernetes, what is considered the primary cluster data source?

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

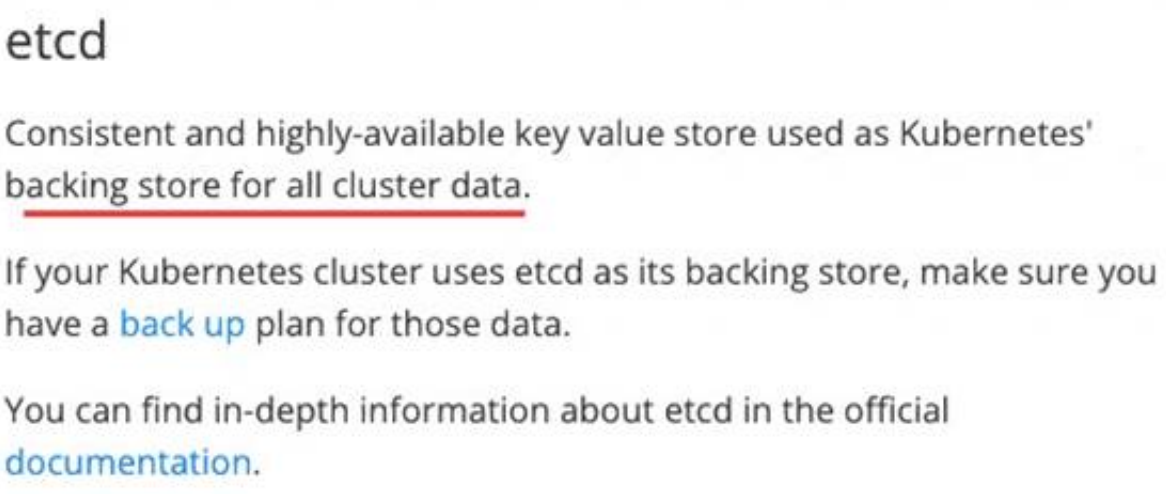

NEW QUESTION 134
Which component of the kubernetes control-plane (master) are all requests to deploy and manage objects posted to?

A. ETCD
B. Controller Manager
C. Kube-proxy
D. API Server
E. Kubelet

Answer: D

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

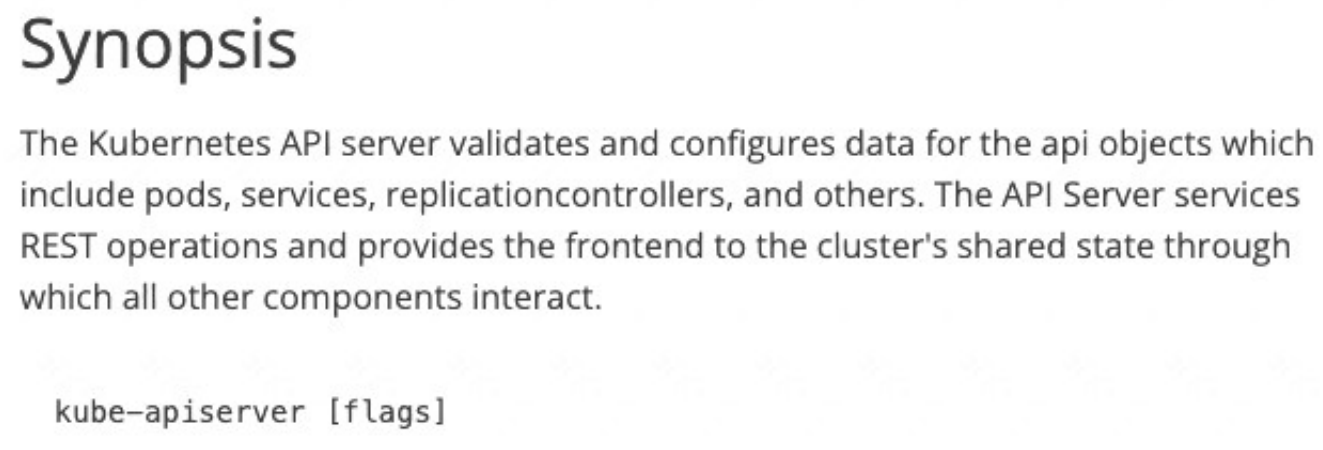

NEW QUESTION 138
Which command is used to expose Kubernetes service

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

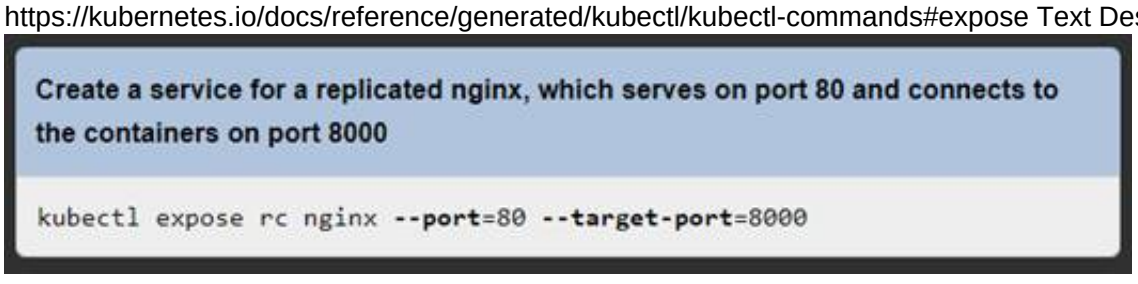

NEW QUESTION 140
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry

* KCNA Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The KCNA Practice Test Here