Linux-Foundation

Exam Questions KCNA

Kubernetes and Cloud Native Associate (KCNA)

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

  All examinations will be up to date.

* 24/7 Quality Support

  We will provide service round the clock.

* 100% Pass Rate

  Our guarantee that you will pass the exam.

* Unique Gurantee

  If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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

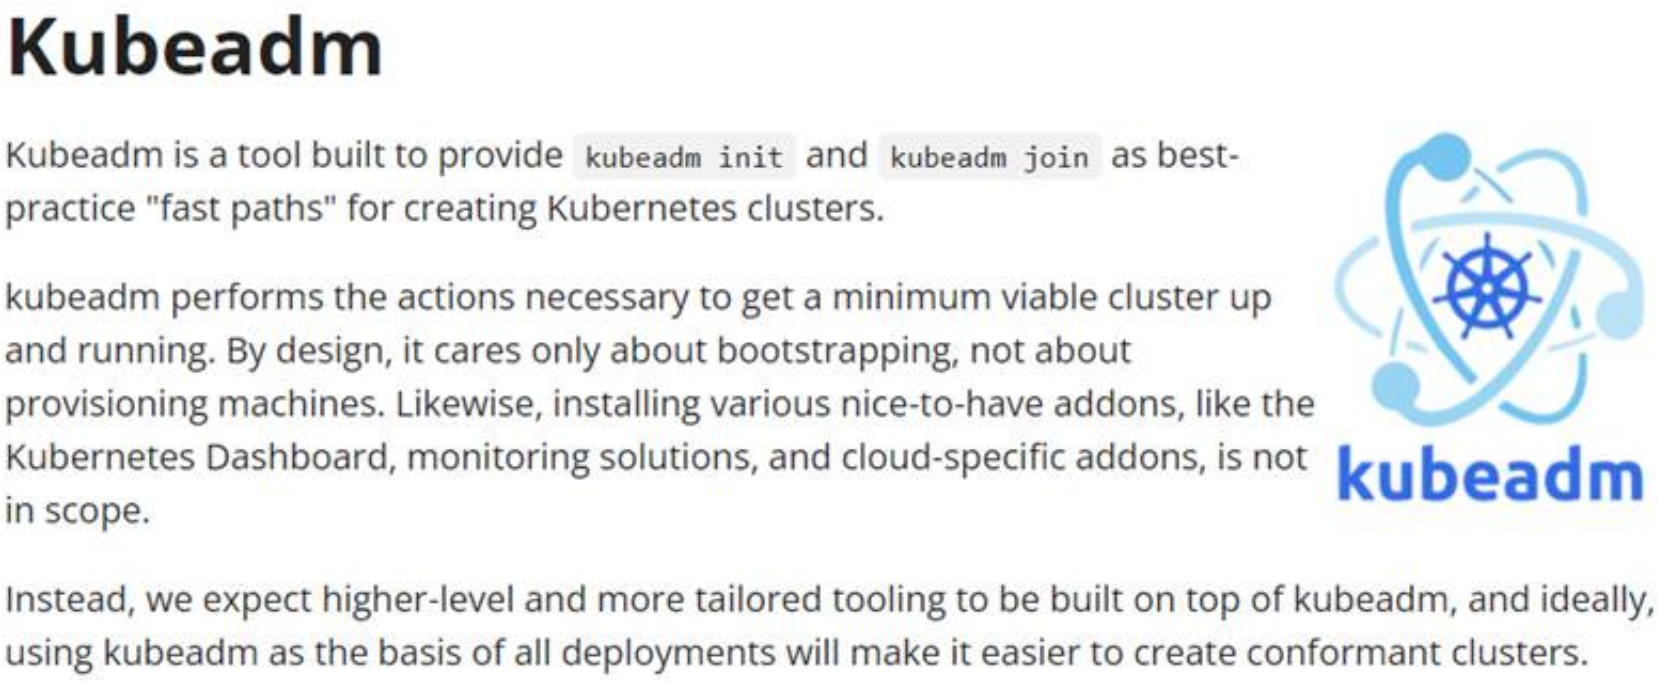

NEW QUESTION 2
What is the most common way to scale the application in the cloud environment?

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
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 4
What is the smallest possible unit in Kubernetes to run a container?

A. pod
B. docker
C. service
D. container

Answer: A

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

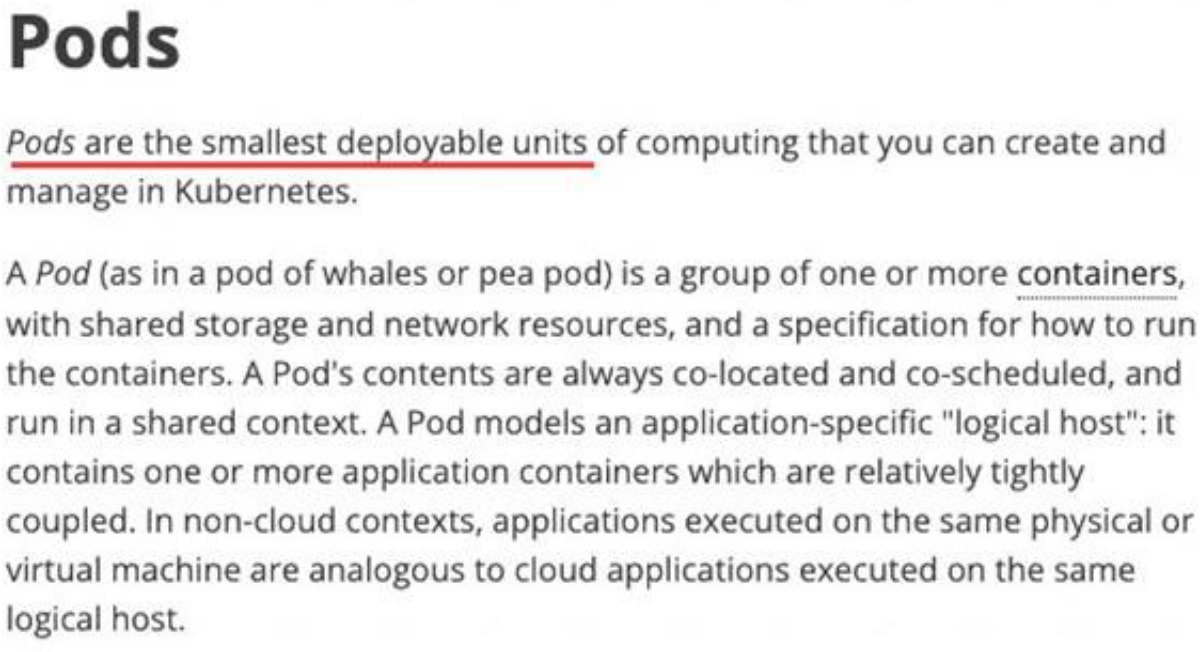

NEW QUESTION 5
There are three Nodes in a cluster, and want to run exactly one replica of a Pod on each Node. Pre-fer to automatically create a replica on any new Nodes when they are added. Which Kubernetes re-source should you use?

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

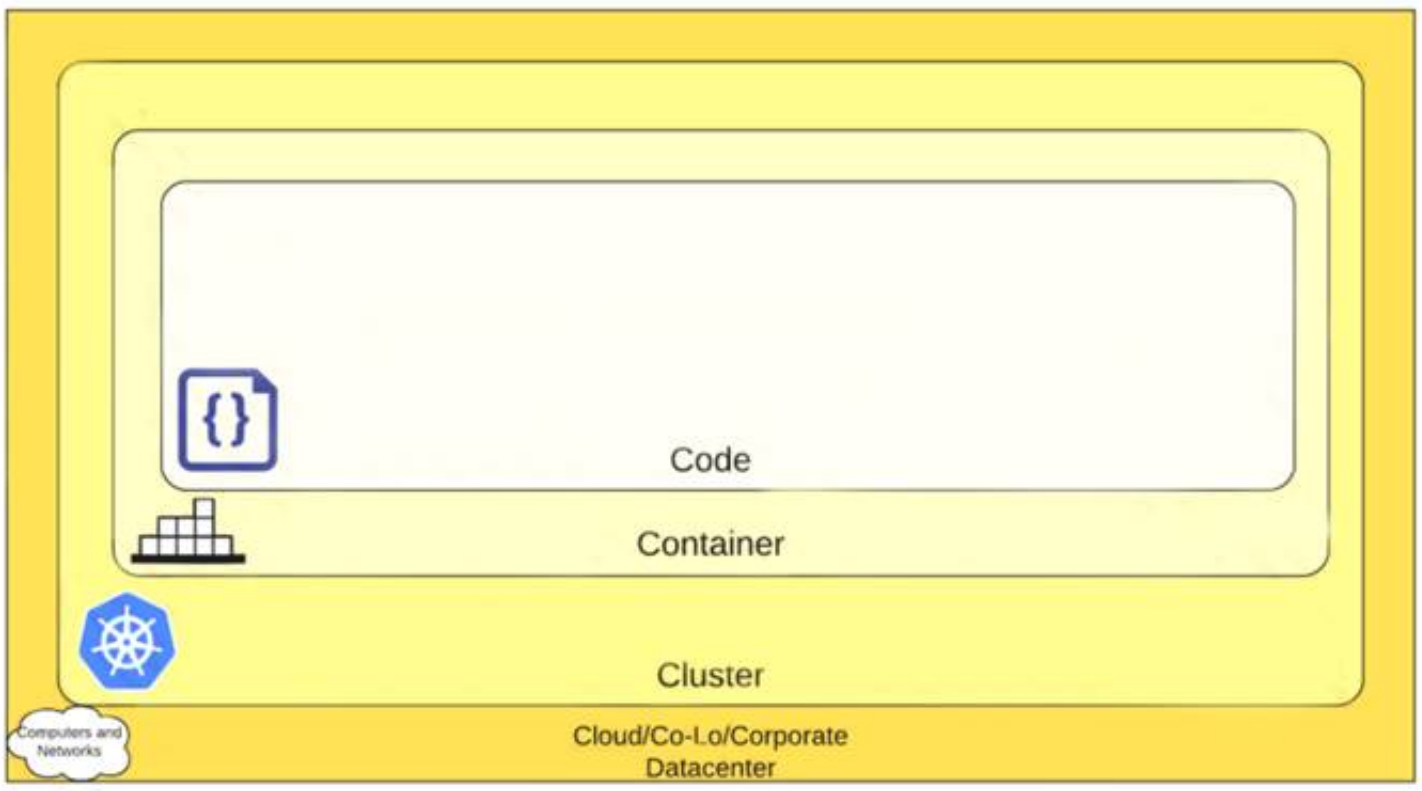

NEW QUESTION 6
What tool allows us to build useful visual representations of prometheus data?

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

NEW QUESTION 7
What is the main difference between Argo vs. Flux CD?

A. Argo is pull-based, and Flux is push-based
B. No difference; both are pull-based
C. Argo is push-based, and Flux is pull-based
D. No difference; both are push-based

Answer: C

Explanation:
AgroCD:
https://argo-cd.readthedocs.io/en/stable/developer-guide/ci/#can-i-retrigger-the-checks-without-pushing-a-new-c
FluxCD: https://fluxcd.io/

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 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

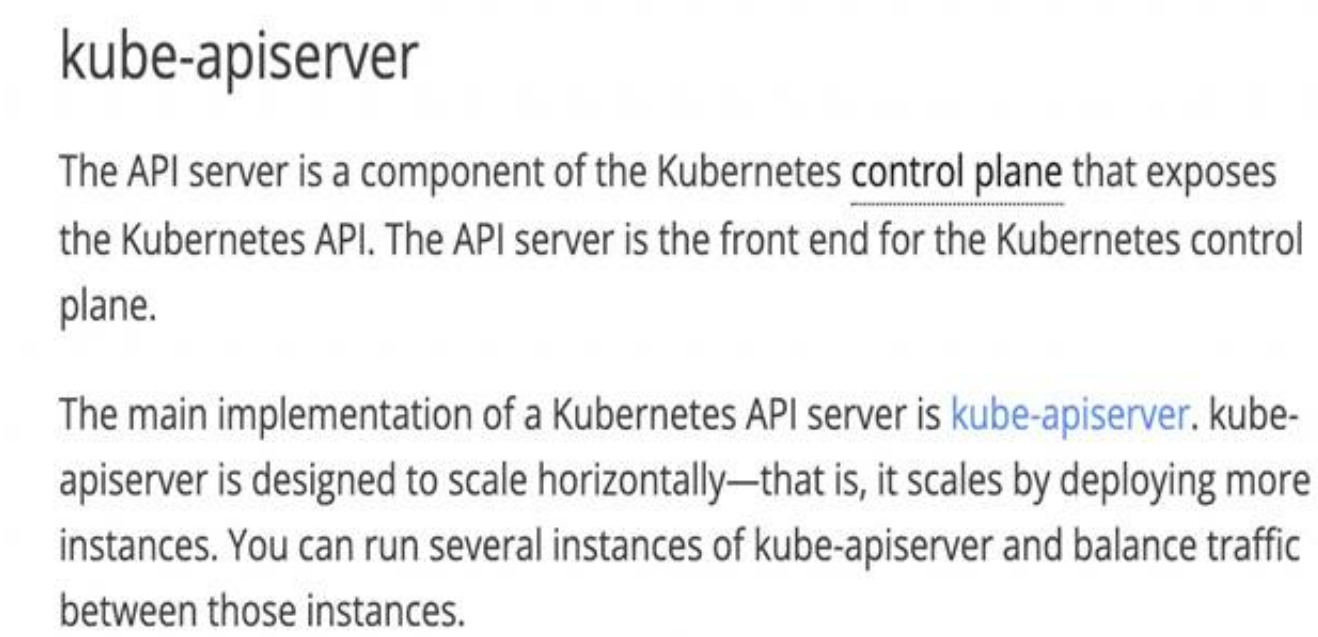

NEW QUESTION 10
The Kubernetes rolling update is used for _____ .

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 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 11
Which authentication method allows JWTs to authenticate?

A. OpenId connect
B. Client 'TLS' certificates
C. OPA gatekeeper
D. Anonymous

Answer: A

NEW QUESTION 13
Which of the following is not the part of Kubernetes Control Plane?

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

Answer: D

Explanation:
https://kubernetes.io/docs/concepts/overview/components/
Graphical user interface, diagram, application Description automatically generated

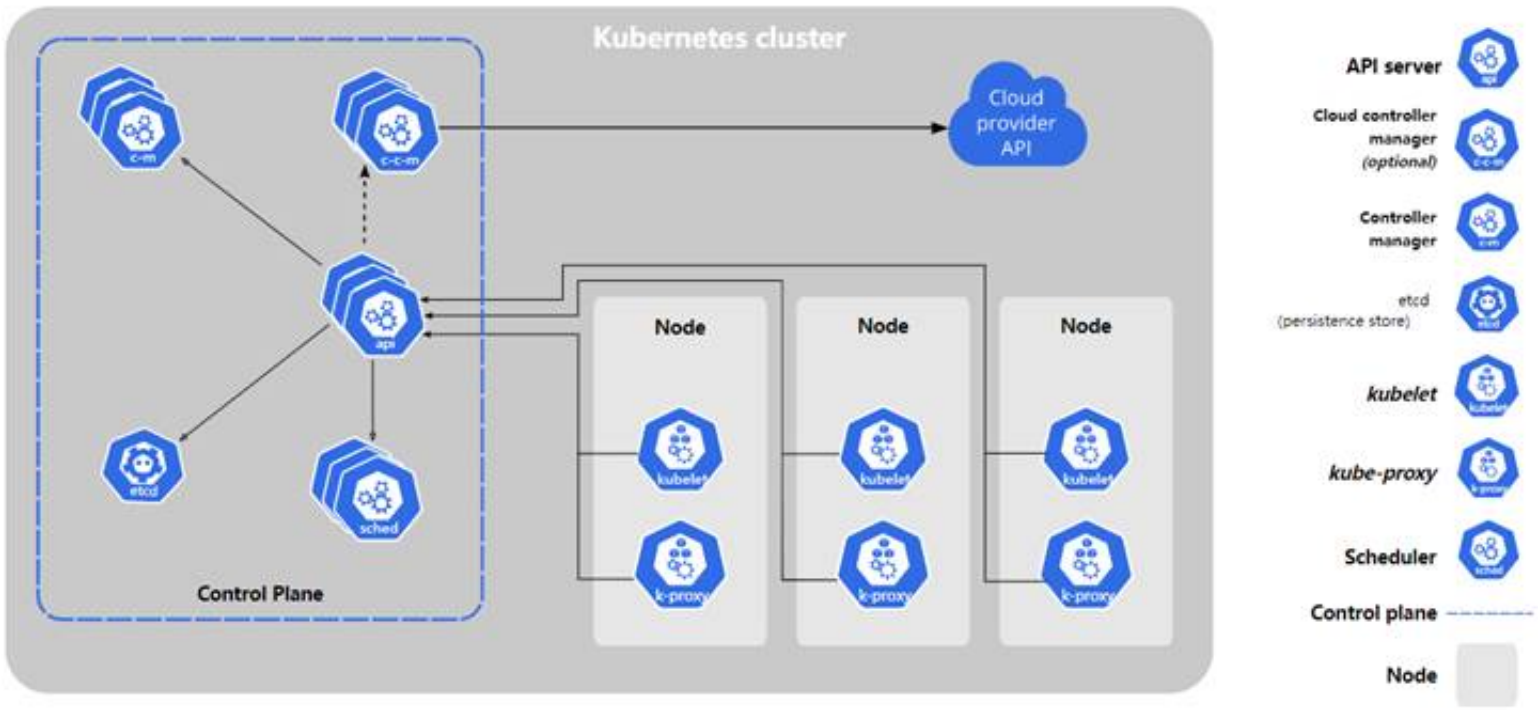

NEW QUESTION 15
Which tool is built on the GitOps toolkit?

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

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

NEW QUESTION 16
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

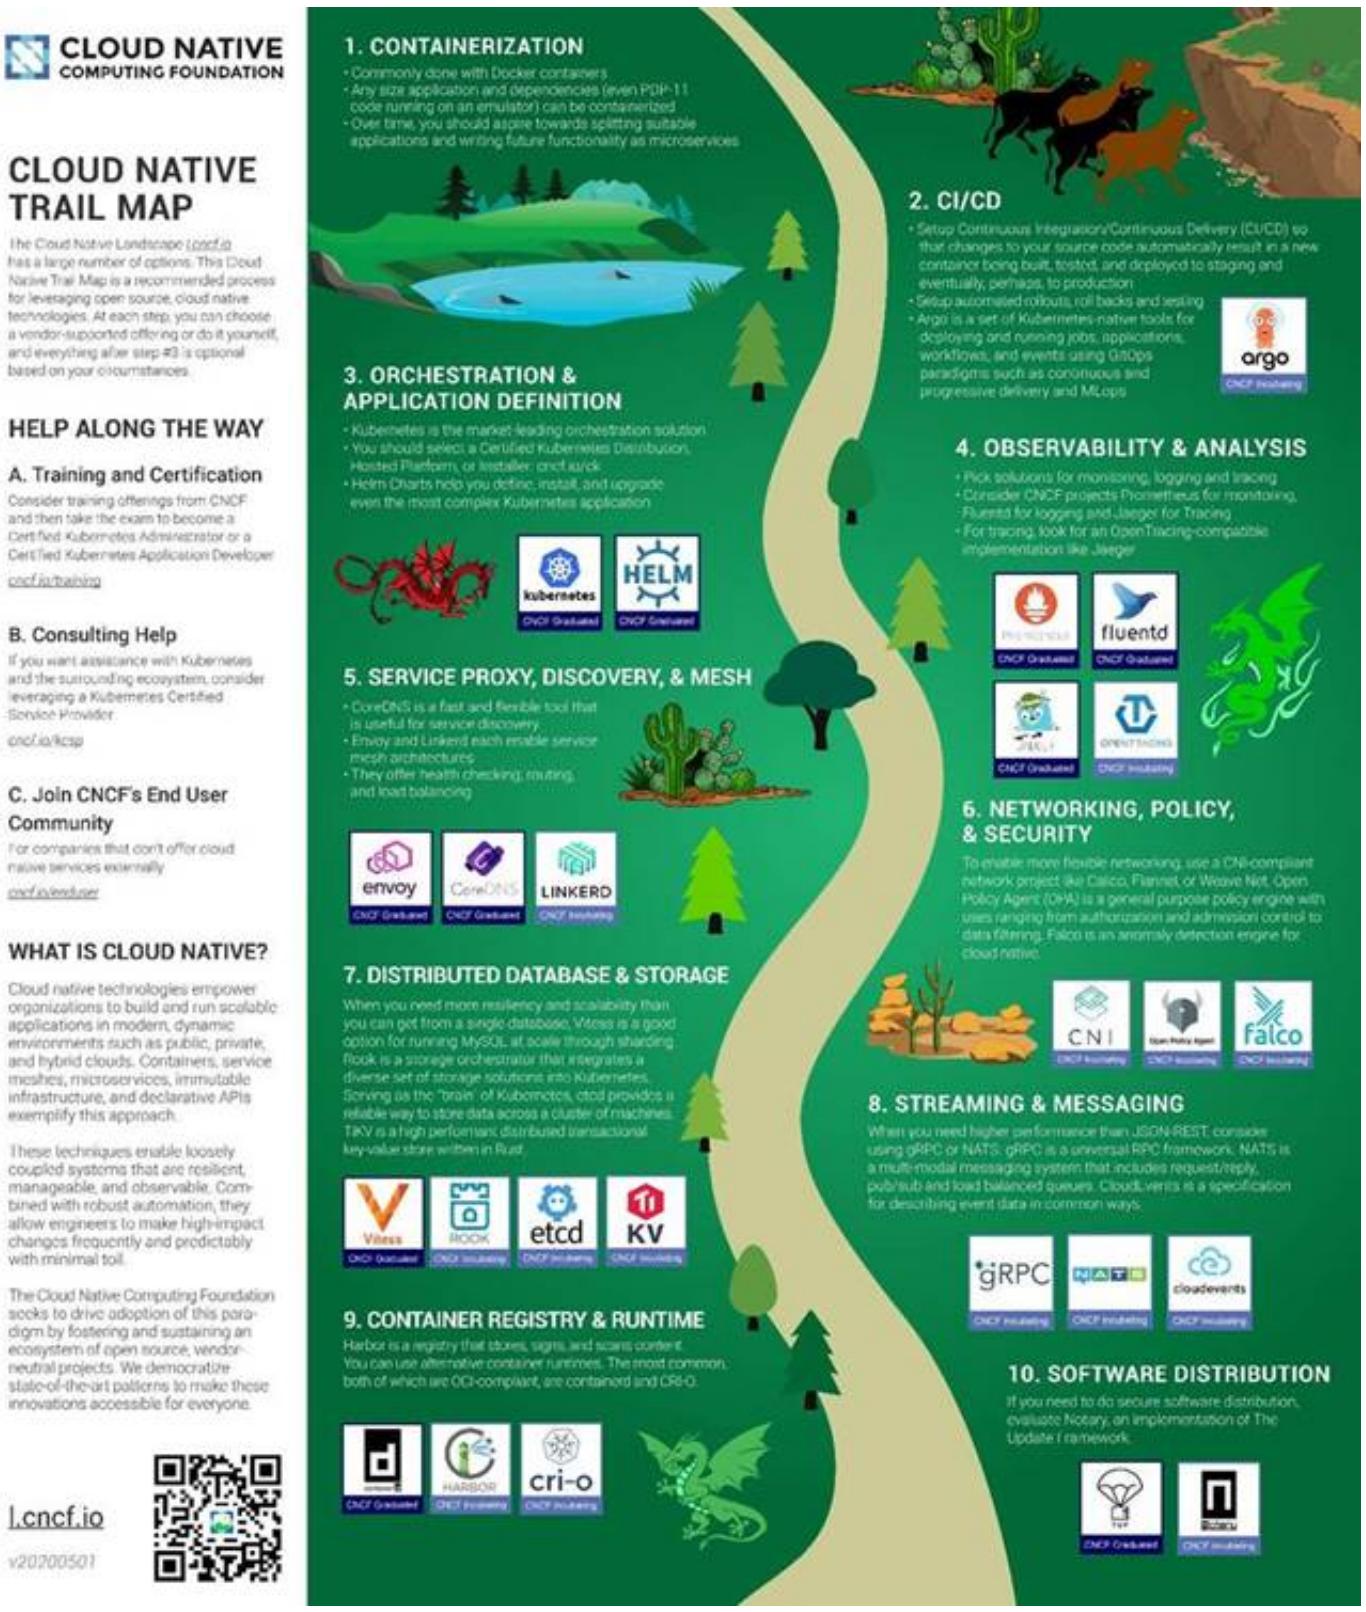

NEW QUESTION 19
What are the two major components of service mesh?

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 21
Which Kubernetes resource creates Kubernetes Jobs?

A. JobFactory
B. CronJob
C. Task
D. JobDeployment

Answer: B

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

NEW QUESTION 24
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 26
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

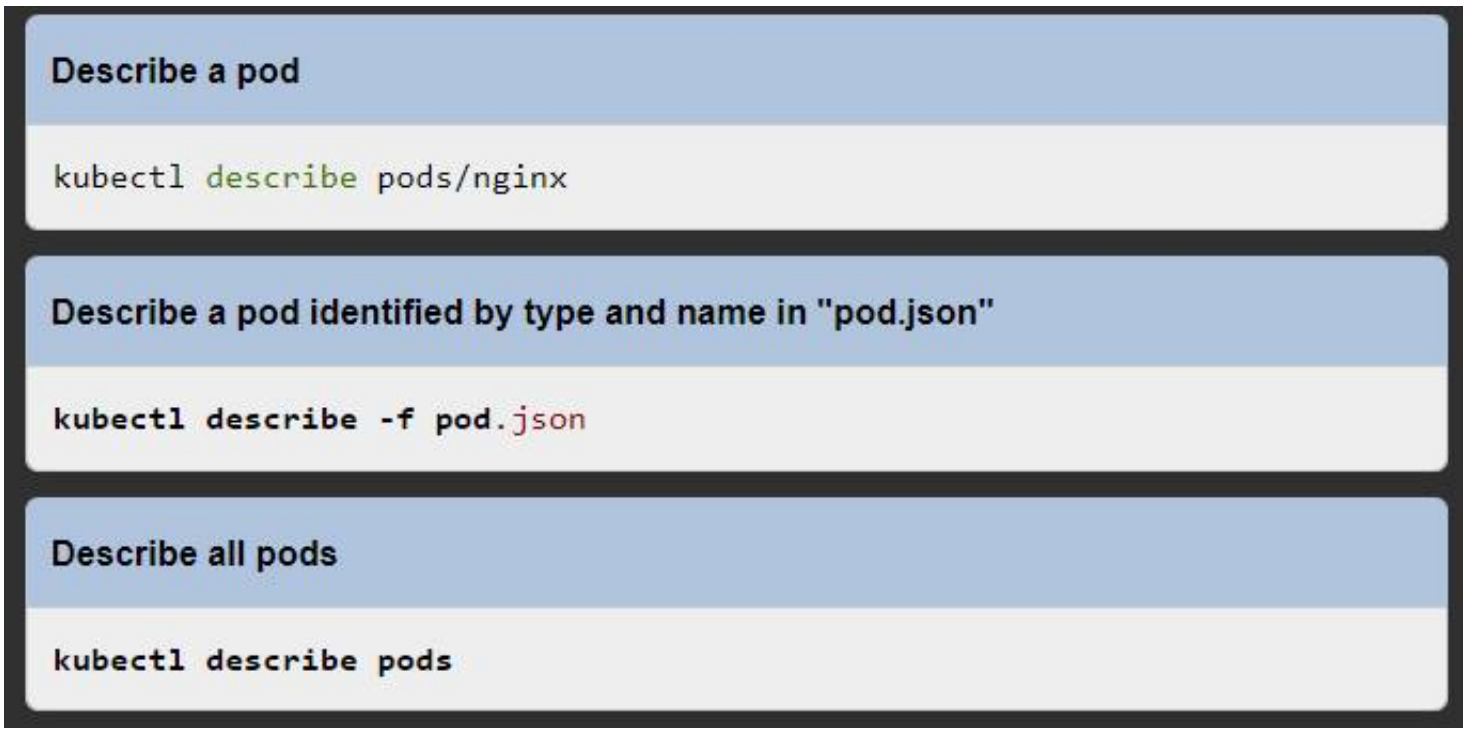

NEW QUESTION 31
What framework allows developers to write code without worrying about the servers and operating systems they will run on?

A. Virtualization
B. Docker
C. Serverless
D. Kubernetes

Answer: C

NEW QUESTION 32
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 34
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 35
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 37
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

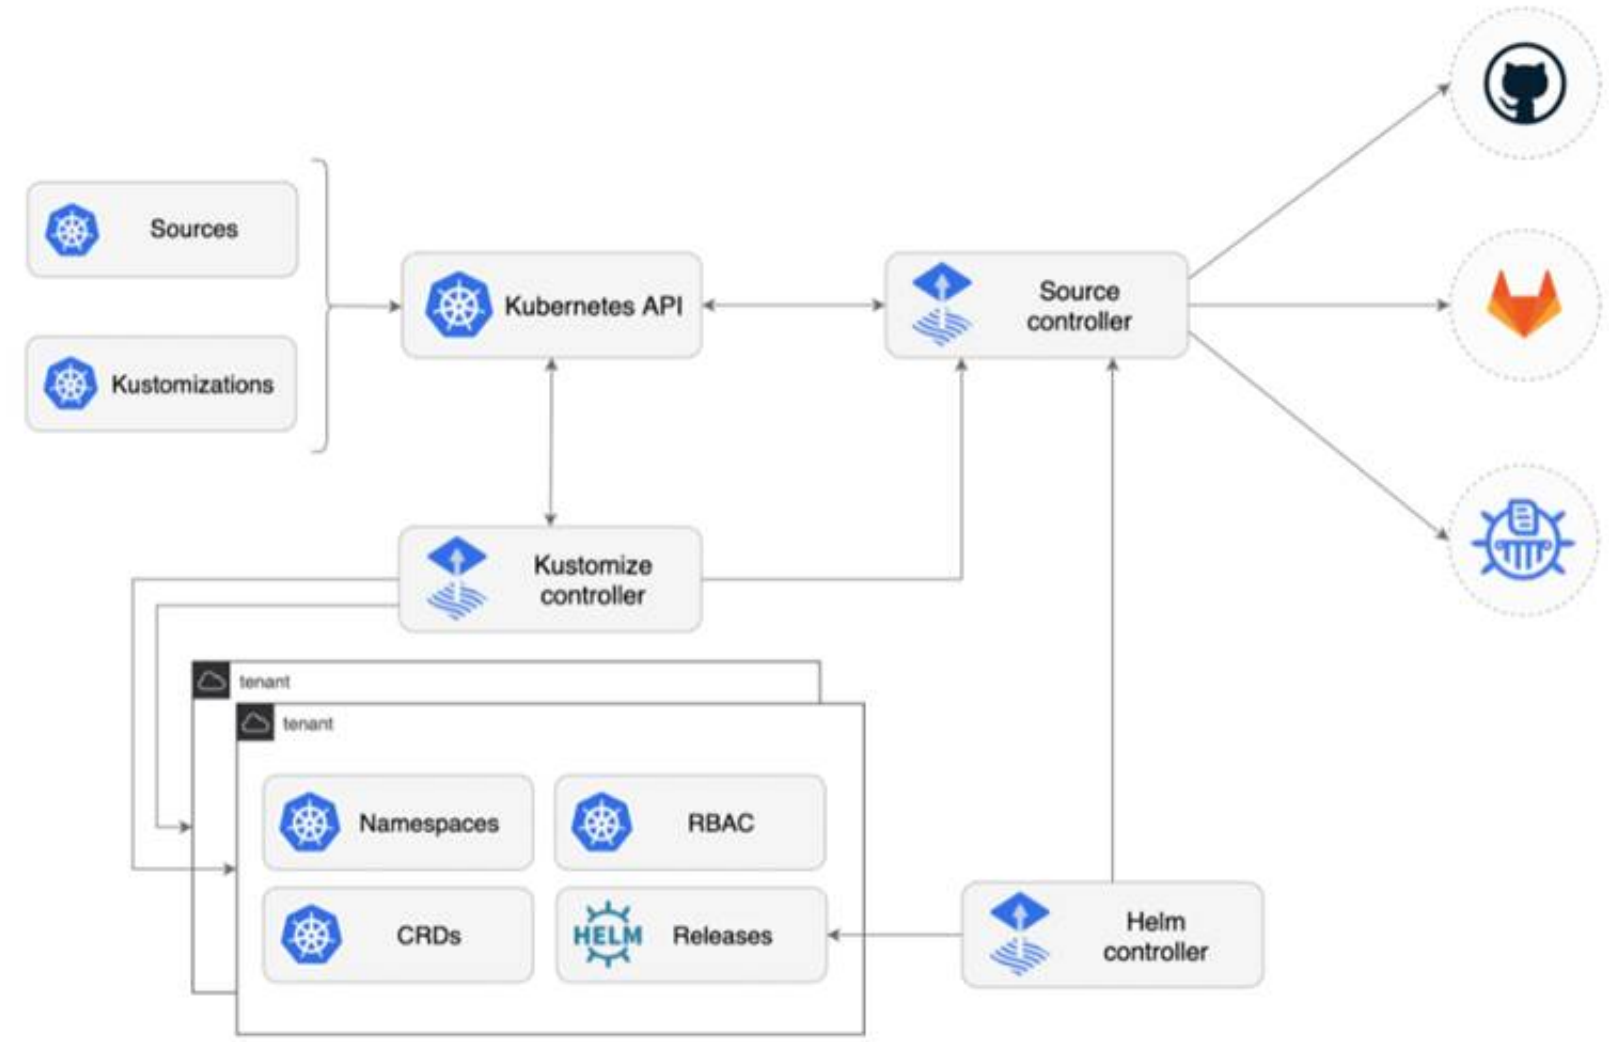

NEW QUESTION 38
Flux is built using which toolkit?

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 41
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

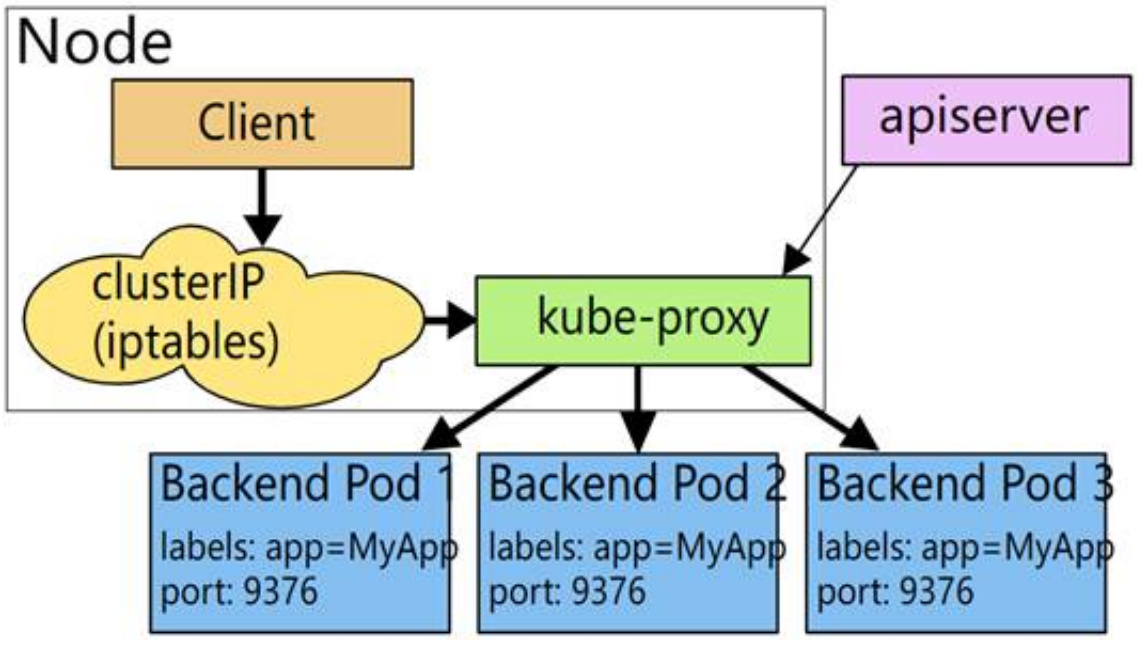

NEW QUESTION 42
What is the use of labels in Kubernetes?

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

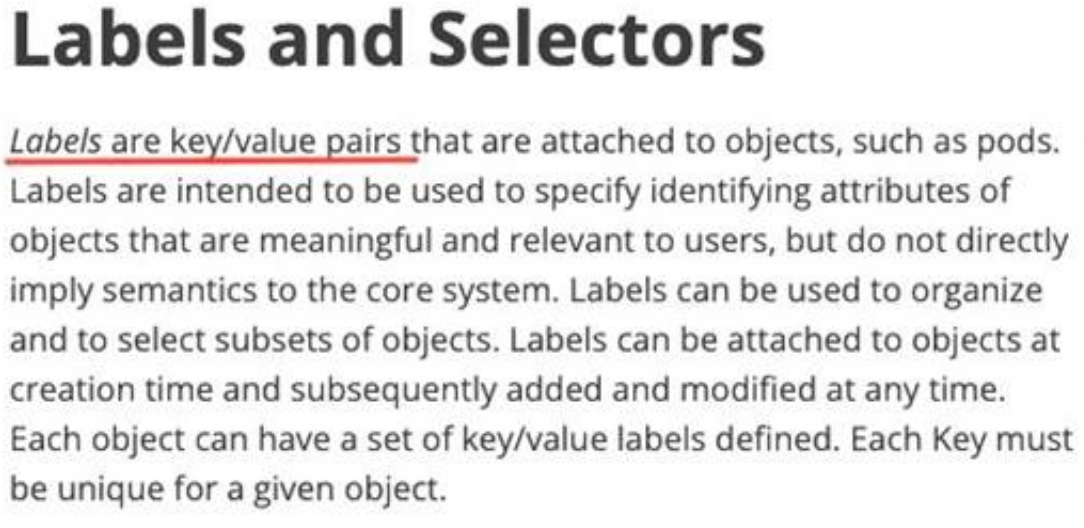

NEW QUESTION 46
What are cluster-wide objects

A. Service and Pods
B. Volumes and Nodes
C. ConfigMaps and Secrets

Answer: B

Explanation:
https://kubernetes.io/docs/concepts/overview/working-with-objects/_print/ Text, letter Description automatically generated

NEW QUESTION 49
Which part of a Kubernetes cluster is responsible for running container workloads?

A. Worker Node

B. kube-proxy
C. Control plane
D. etcd

Answer: A

Explanation:
Worker Nodes are responsible for executing containerized workloads.

NEW QUESTION 51
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 56
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 58
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/

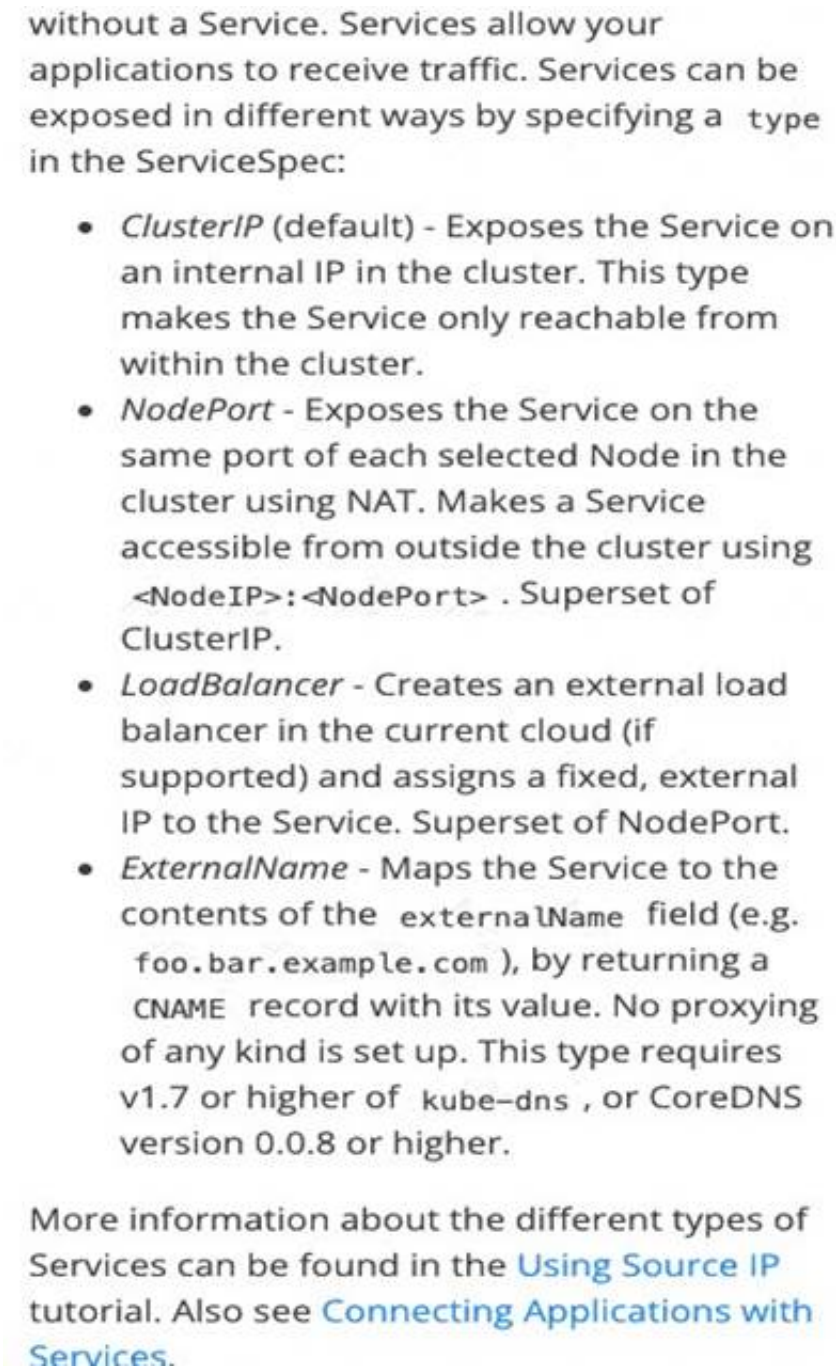

Text Description automatically generated

NEW QUESTION 63
Which of the following is an advantage a cloud-native microservices application has over monolithic applications?

A. Cloud-native microservices applications tend to be faster and more responsive than monolithic applications.
B. Cloud-native microservice applications tend to be easier to troubleshoot.
C. Cloud-native microservice applications tend to be easier to scale and perform updates on.

Answer: C

Explanation:
Cloud-native applications tend to be microservice base, they have individual services that can be independently scaled, updated and rolled back. This makes scaling and update operations simpler and less risky.

NEW QUESTION 66
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 67
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 68
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 70
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 74
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

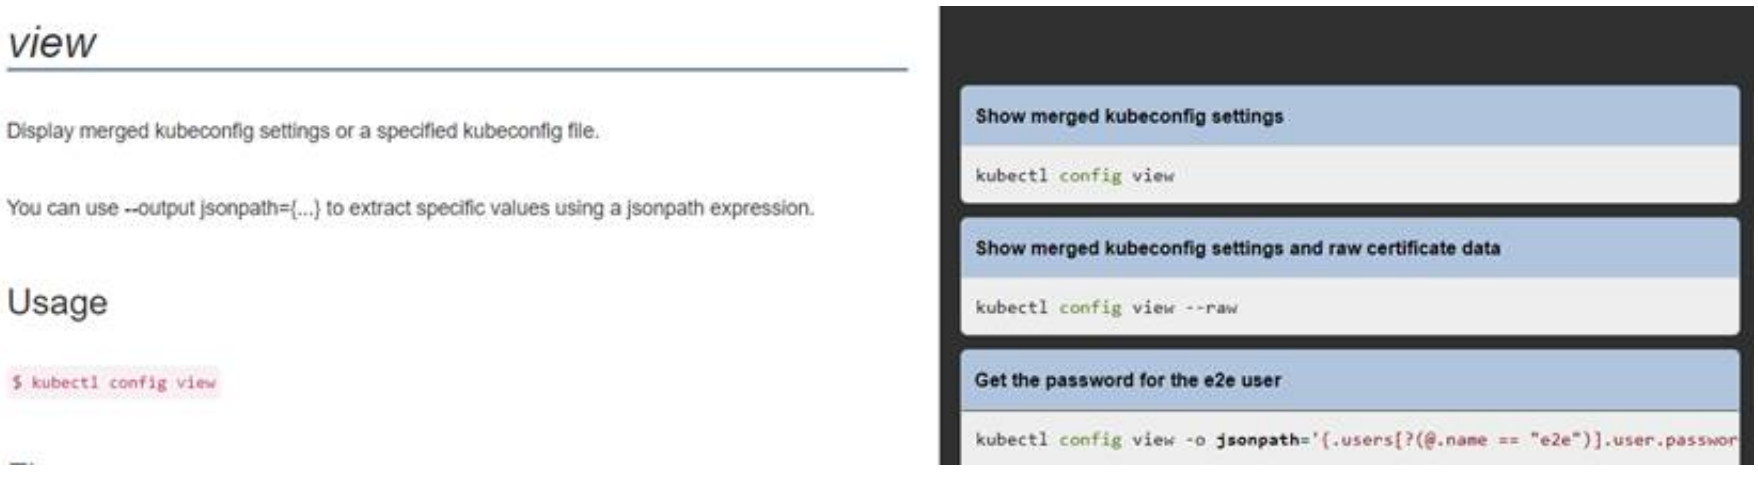

NEW QUESTION 76
The Kubernetes API provides an interface for storing objects. Which of the following describes the type of objects stored by the Kubernetes API?

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

Answer: B

Explanation:
Kubernetes objects are RESTful objects.

NEW QUESTION 80
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 82
Which of the following provides cloud-native storage orchestration?

A. Cloud Provider Specific storage (EBS, EFS, Cloud Storage)
B. Cloud Storage
C. Storage IO

Answer: A

Explanation:
https://kubernetes.io/docs/concepts/storage/persistent-volumes/#types-of-persistent-volumes Table Description automatically generated with medium confidence

Types of Persistent Volumes

PersistentVolume types are implemented as plugins. Kubernetes currently supports the following plugins:

- `awsElasticBlockStore` - AWS Elastic Block Store (EBS)
- `azureDisk` - Azure Disk
- `azureFile` - Azure File
- `cephfs` - CephFS volume
- `csi` - Container Storage Interface (CSI)
- `fc` - Fibre Channel (FC) storage
- `gcePersistentDisk` - GCE Persistent Disk
- `glusterfs` - Glusterfs volume
- `hostPath` - HostPath volume (for single node testing only; WILL NOT WORK in a multi-node cluster; consider using `local` volume instead)
- `iscsi` - iSCSI (SCSI over IP) storage
- `local` - local storage devices mounted on nodes.
- `nfs` - Network File System (NFS) storage
- `portworxVolume` - Portworx volume
- `rbd` - Rados Block Device (RBD) volume
- `vsphereVolume` - vSphere VMDK volume

The following types of PersistentVolume are deprecated. This means that support is still available but will be removed in a future Kubernetes release.

- `cinder` - Cinder (OpenStack block storage) (**deprecated** in v1.18)

NEW QUESTION 87
What CNCF project is the leading DNS project in the CNCF landscape?

A. Kubernetes
B. gRPC
C. KubeDNS
D. CoreDNS

Answer: D

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

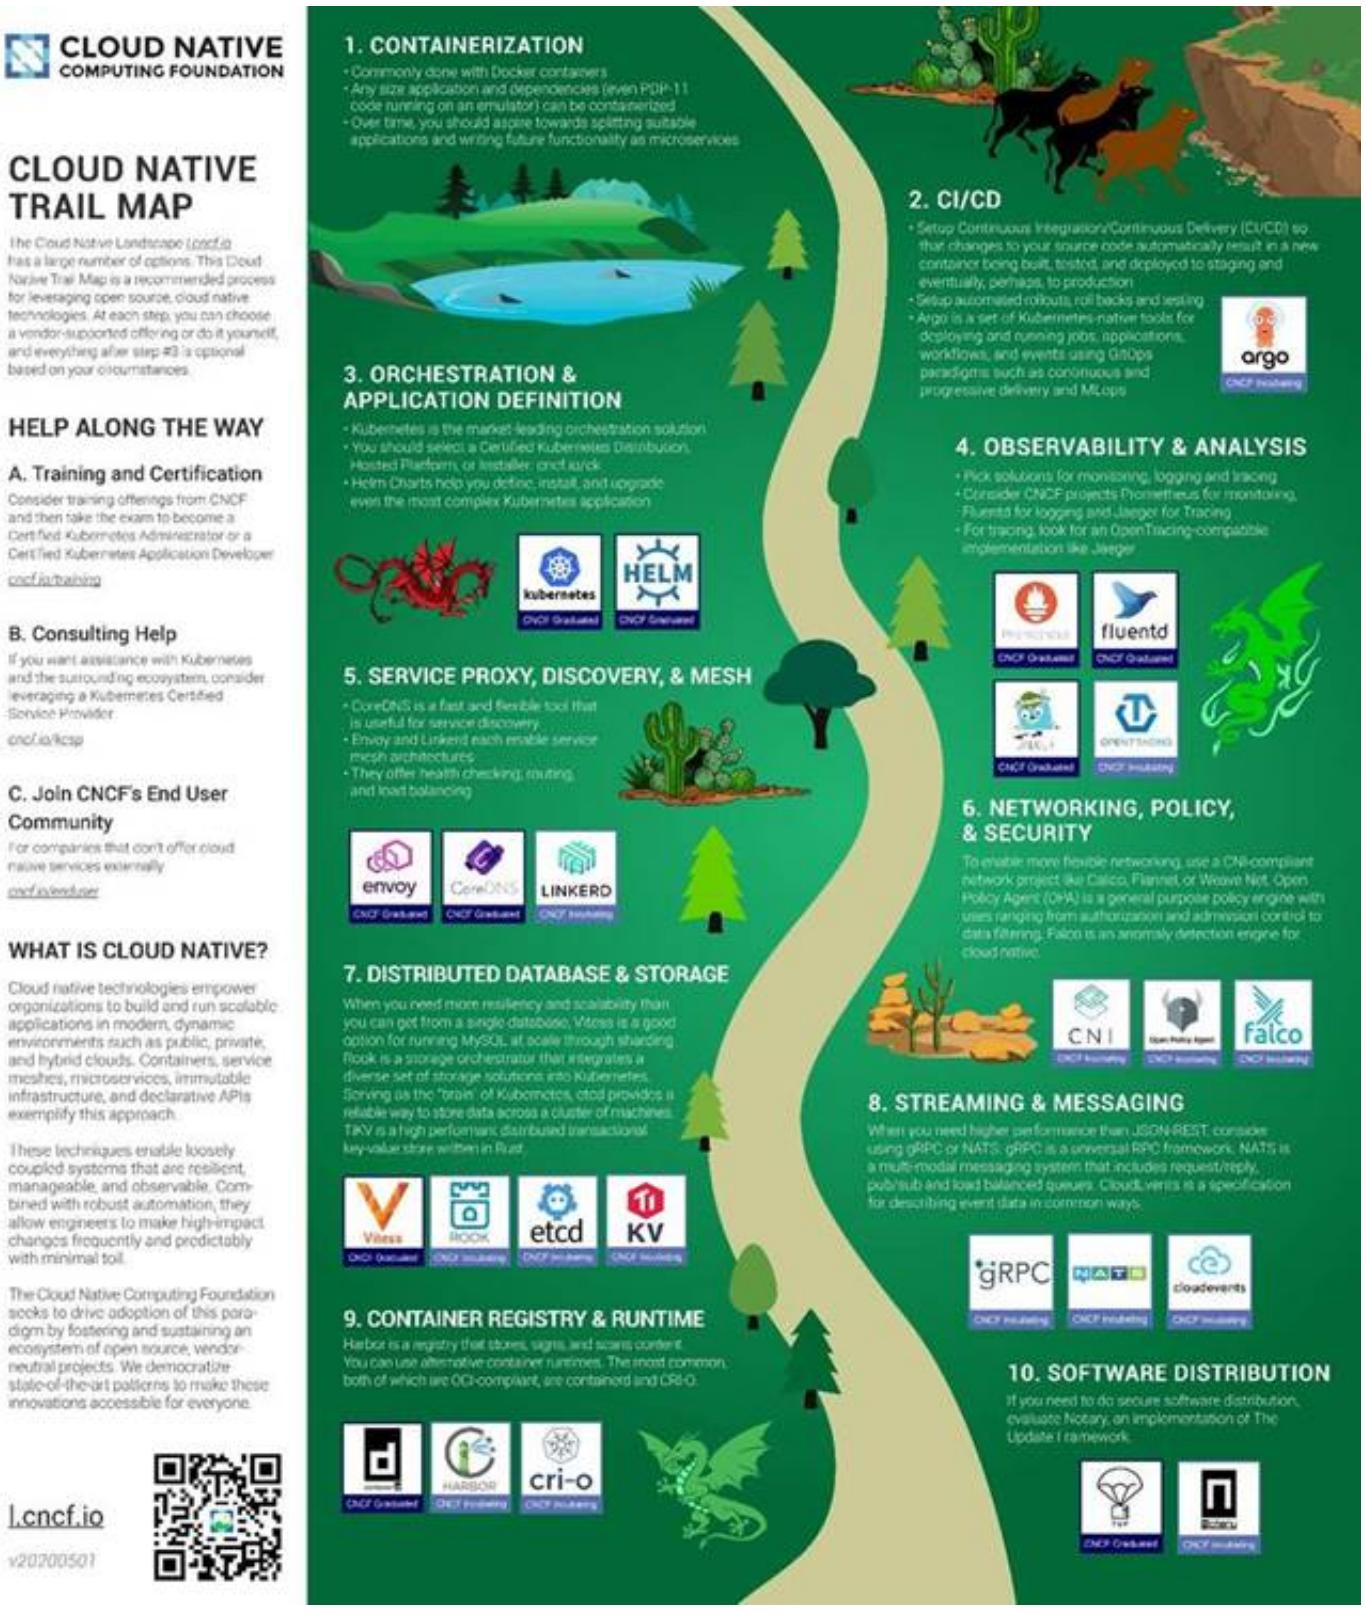

NEW QUESTION 88
Which is NOT a use case for the Kubernetes dashboard?

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 93
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 96
What is a commonly used package manager for kubernetes applications?

A. npm
B. apt
C. helm
D. kubernetes manifest

Answer: C

Explanation:
https://helm.sh/

NEW QUESTION 97
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 99
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 104
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 109
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 110
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 115
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

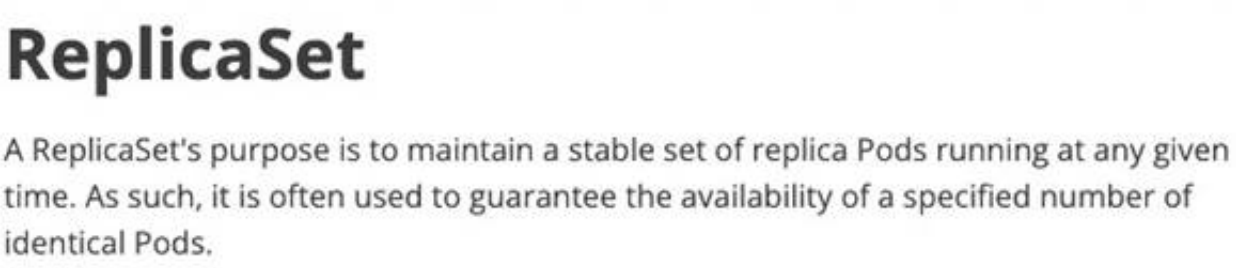

NEW QUESTION 117
Which access control component of Kubernetes is responsible for authorization and decides what requestor is allowed to do?

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

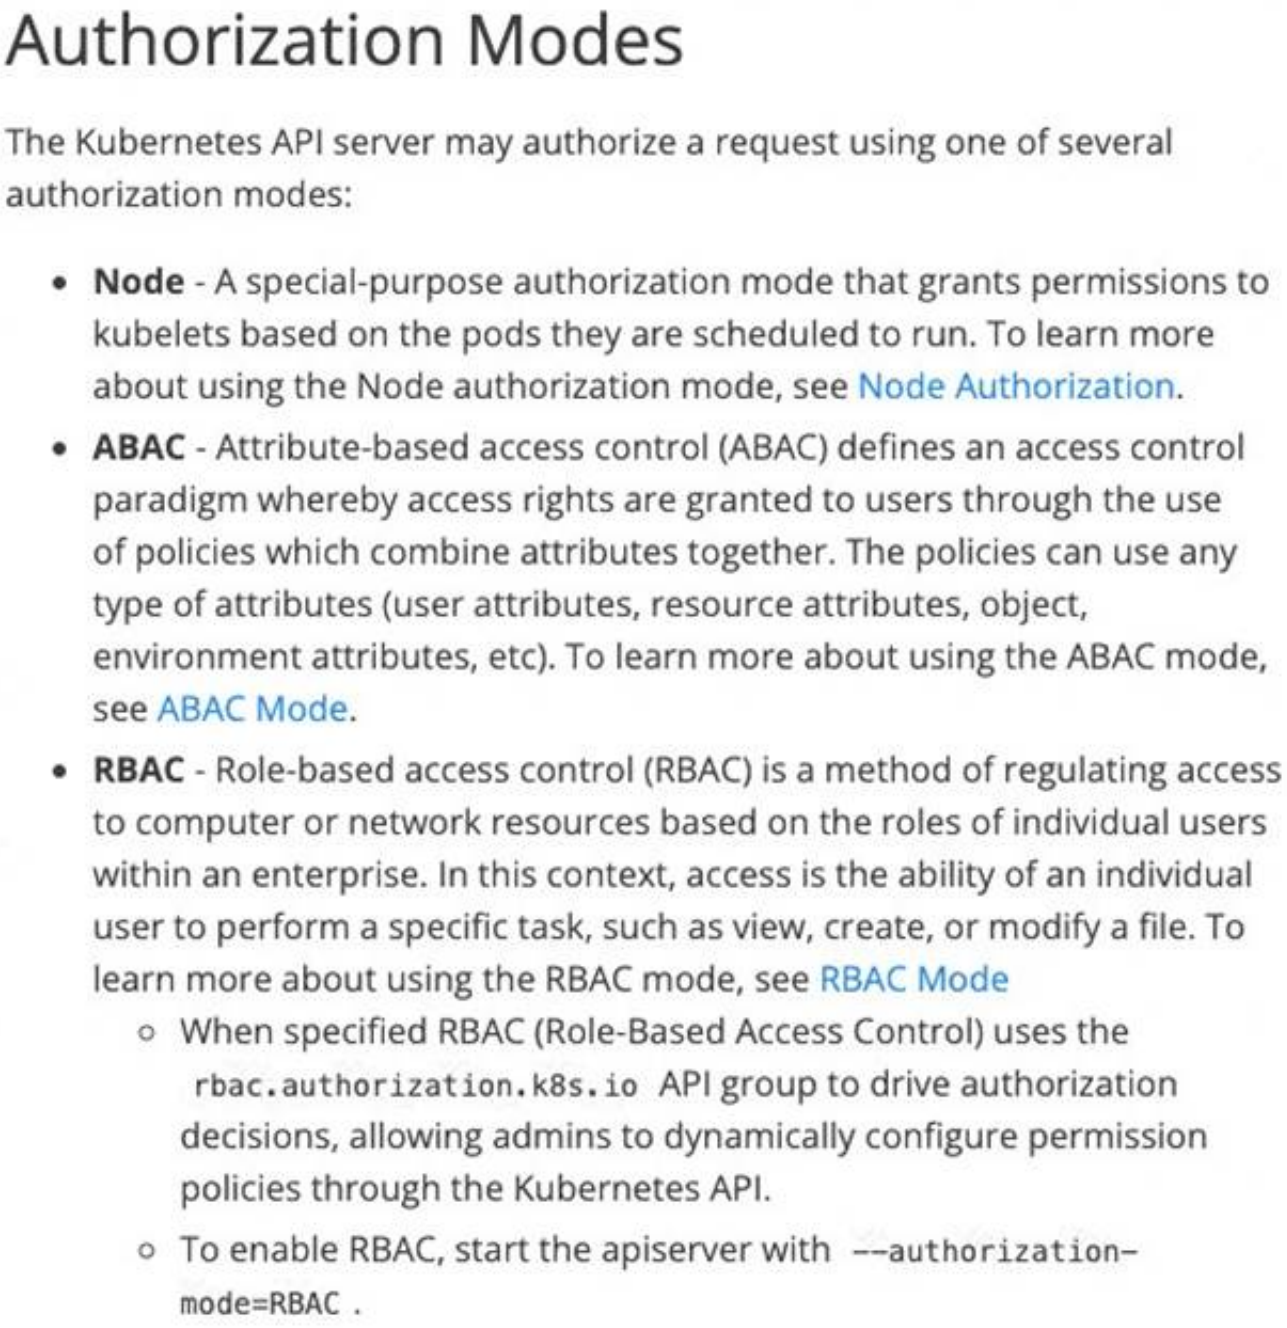

NEW QUESTION 120
Which statement is true about Pod Networking?

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 124
Which of the following is NOT a Kubernetes component?

A. Scheduler
B. Docker
C. Cloud Controller manager
D. Kube-proxy

Answer: B

Explanation:
Docker is not a Kubernetes component.

NEW QUESTION 127
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 129
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

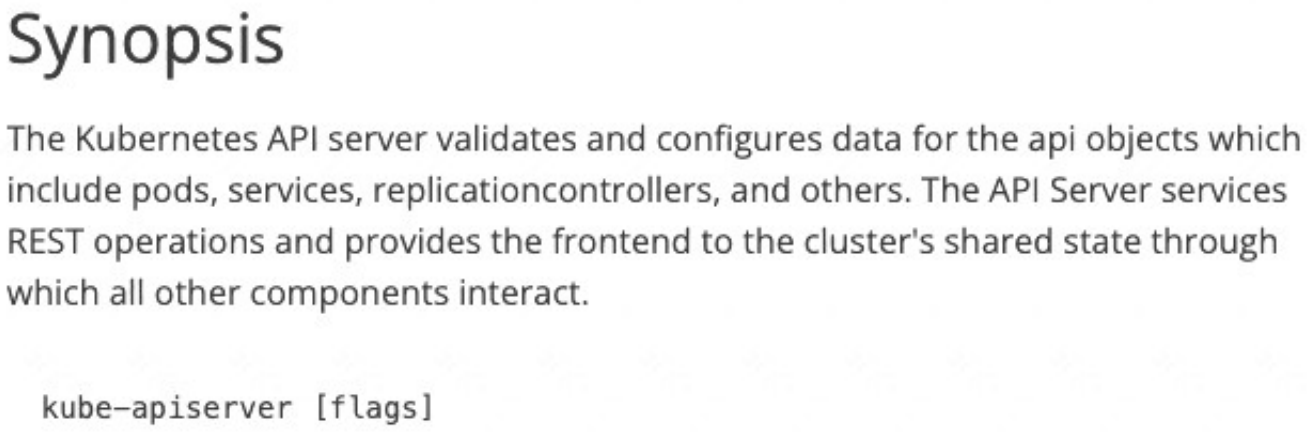

NEW QUESTION 131
What is scheduling in Kubernetes

A. Determining when to execute a cron-job
B. Assigning pods to nodes
C. Joining a new nodes to the clusters
D. Setting a time for automated tasks

Answer: B

Explanation:
https://kubernetes.io/docs/concepts/scheduling-eviction/
Graphical user interface, application Description automatically generated

NEW QUESTION 136
.....

Relate Links

100% Pass Your KCNA Exam with Exambible Prep Materials

https://www.exambible.com/KCNA-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/