Linux-Foundation

Exam Questions KCNA

Kubernetes and Cloud Native Associate (KCNA)

NEW QUESTION 1
kubeadm is an administrative dashboard for kubernetes

A. False
B. True

Answer: A

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

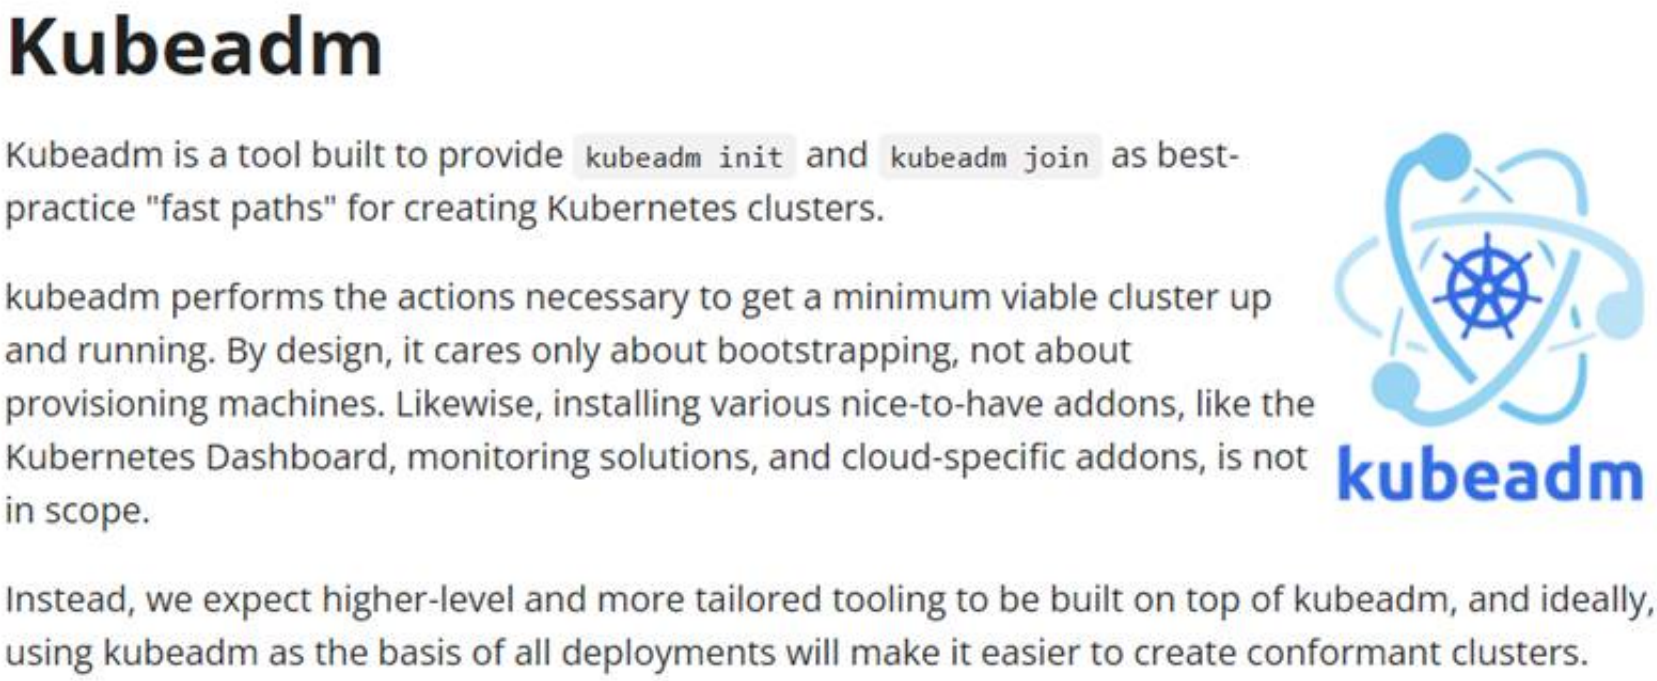

NEW QUESTION 2
There are three Nodes in a cluster, and want to run exactly one replica of a Pod on each Node. Pre-fer to automatically create a replica on any new Nodes when they are added. Which Kubernetes re-source should you use?

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

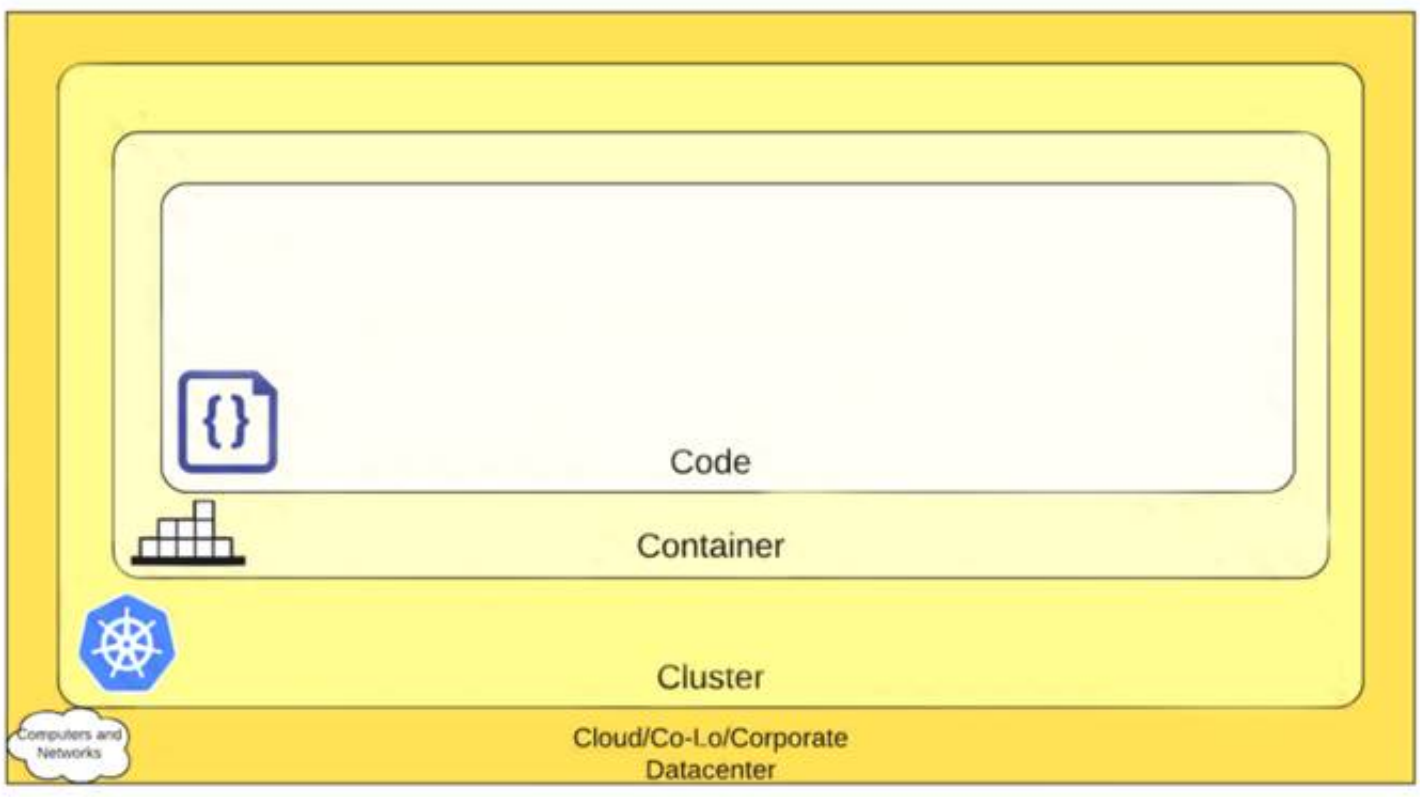

NEW QUESTION 3
What is the main difference between Argo vs. Flux CD?

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 4
A is a ready-to-run software package, containing everything needed to run an application.

A. Container Repository
B. Container Runtime
C. Docker
D. Container Image

Answer: D

Explanation:
https://kubernetes.io/docs/concepts/containers/#container-images Text, letter Description automatically generated

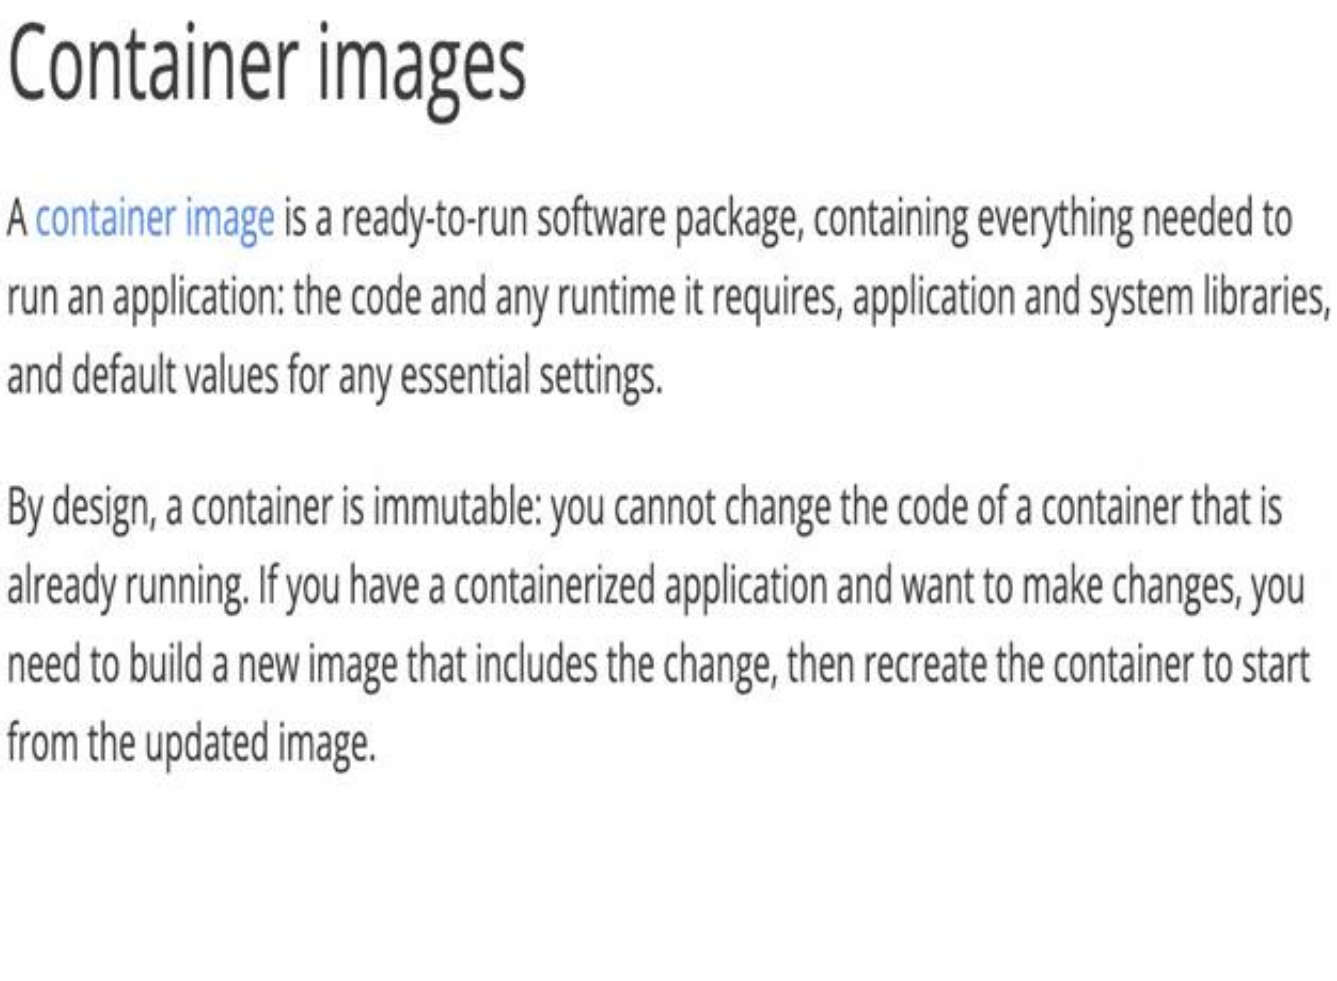

NEW QUESTION 5
Which prometheus metric type represents a single number value that can increase and decrease over time?

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

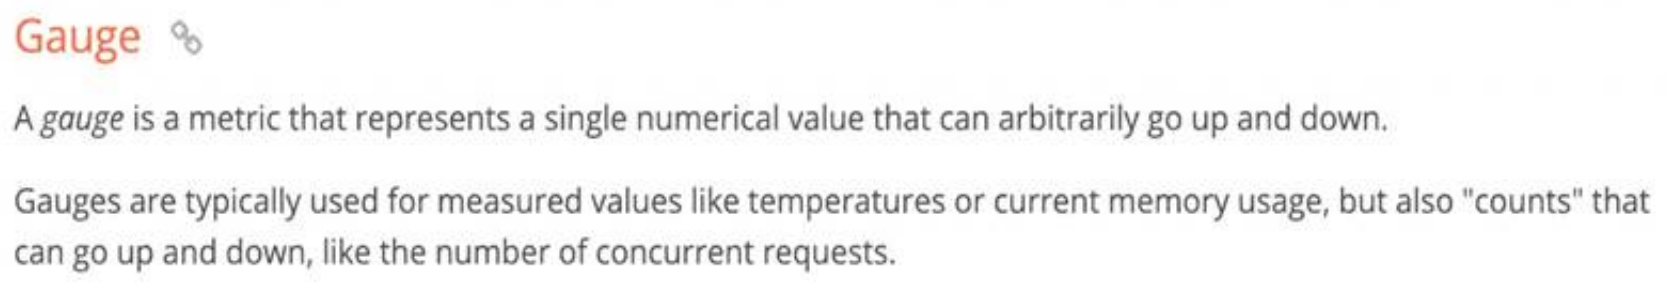

NEW QUESTION 6
What tool allows you to create self-managing, self-scaling, self-healing storage?

A. Persistent Volume
B. Persistent Volume Claim

C. Storage Class
D. Rook
E. Volume

Answer: D

Explanation:
https://rook.io/
Text Description automatically generated

NEW QUESTION 7
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 8
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

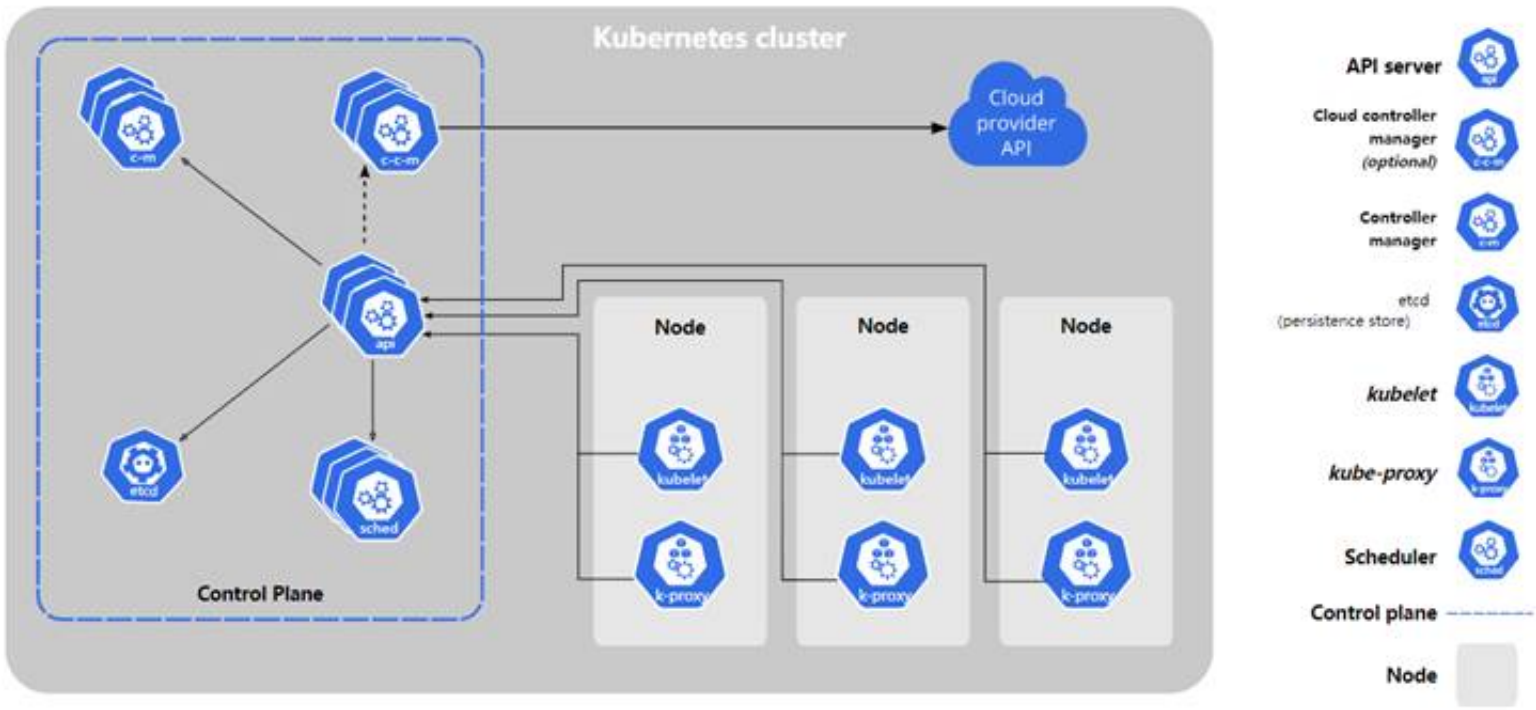

NEW QUESTION 9
What are the two major components of service mesh?

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

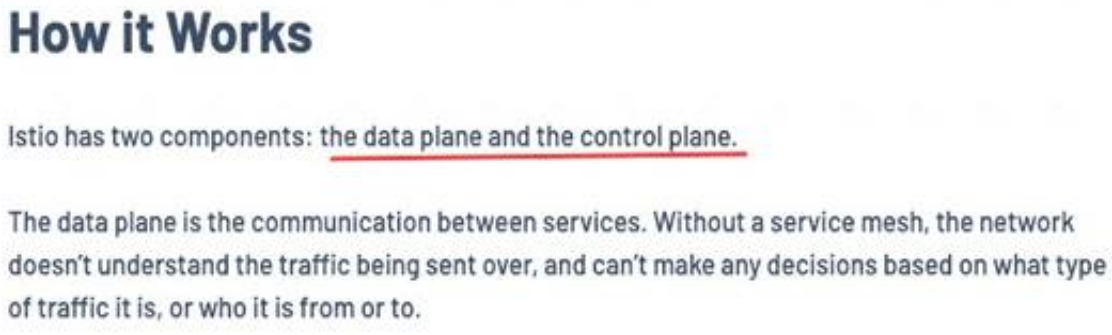
How it Works

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 10
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

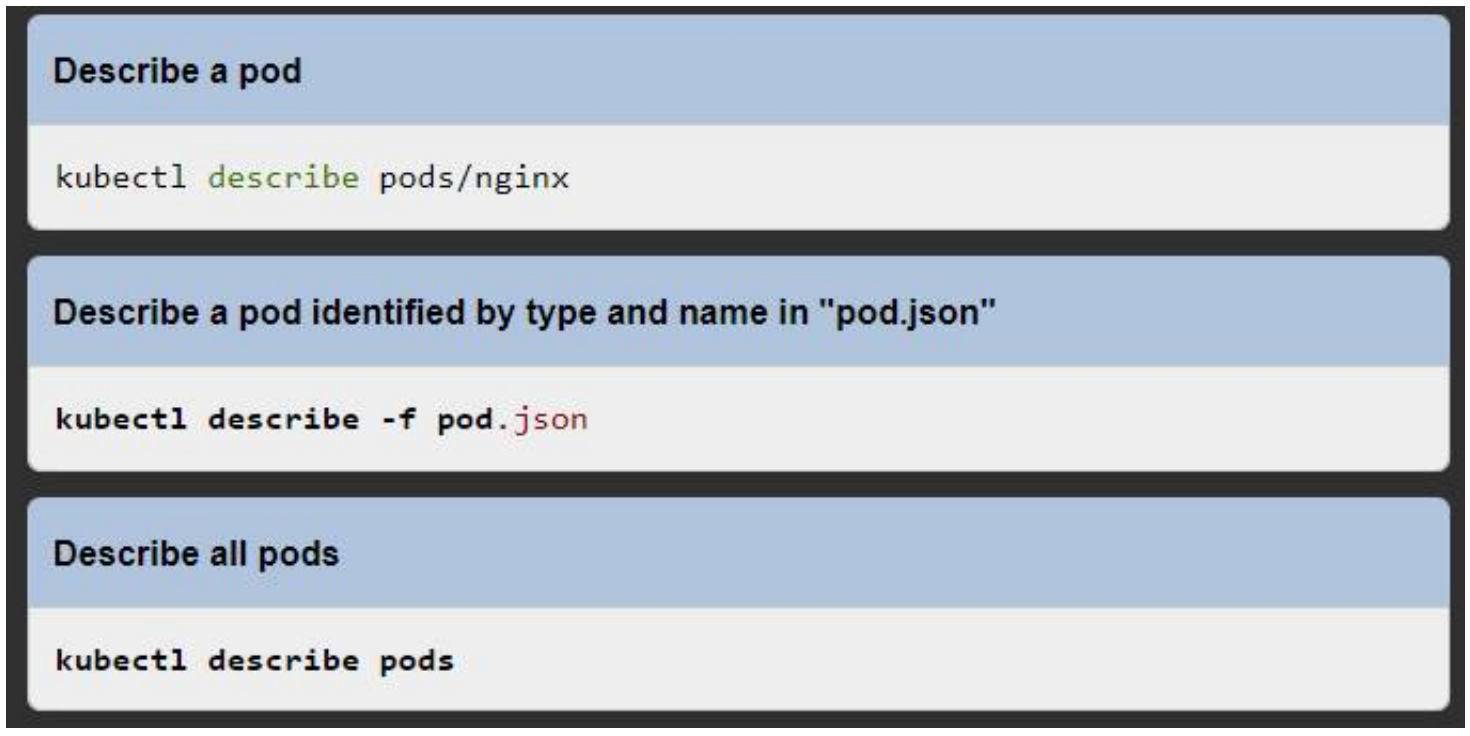

NEW QUESTION 10
Which of the following is not the Kubernetes AutoScaling Strategy?

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 12
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 16
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 17
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 20
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 23
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

4 - Namespaces

In Kubernetes, *namespaces* provides a mechanism for isolating groups of resources within a single cluster. Names of resources need to be unique within a namespace, but not across namespaces. Namespace-based scoping is applicable only for namespaced objects *(e.g. Deployments, Services, etc)* and not for cluster-wide objects *(e.g. StorageClass, Nodes, PersistentVolumes, etc)*.

NEW QUESTION 24
What feature is used for selecting the container runtime configuration?

A. RuntimeClass
B. RuntimeContainer
C. Runtime
D. RuntimeConfig

Answer: A

Explanation:
https://kubernetes.io/docs/concepts/containers/runtime-class/
Graphical user interface, text, application, email Description automatically generated

Runtime Class

FEATURE STATE: `Kubernetes v1.20 [stable]`

This page describes the RuntimeClass resource and runtime selection mechanism.

RuntimeClass is a feature for selecting the container runtime configuration. The container runtime configuration is used to run a Pod's containers.

Motivation

You can set a different RuntimeClass between different Pods to provide a balance of performance versus security. For example, if part of your workload deserves a high level of information security assurance, you might choose to schedule those Pods so that they run in a container runtime that uses hardware virtualization. You'd then benefit from the extra isolation of the alternative runtime, at the expense of some additional overhead.

You can also use RuntimeClass to run different Pods with the same container runtime but with different settings.

NEW QUESTION 28
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 31
What is the name for the tool that manages communication between pods, injects a sidecar proxy container into each pod and directs network traffic through the proxy container?

A. namespace
B. Deployment
C. Network policy
D. Service mesh
E. Service

Answer: D

NEW QUESTION 34
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/

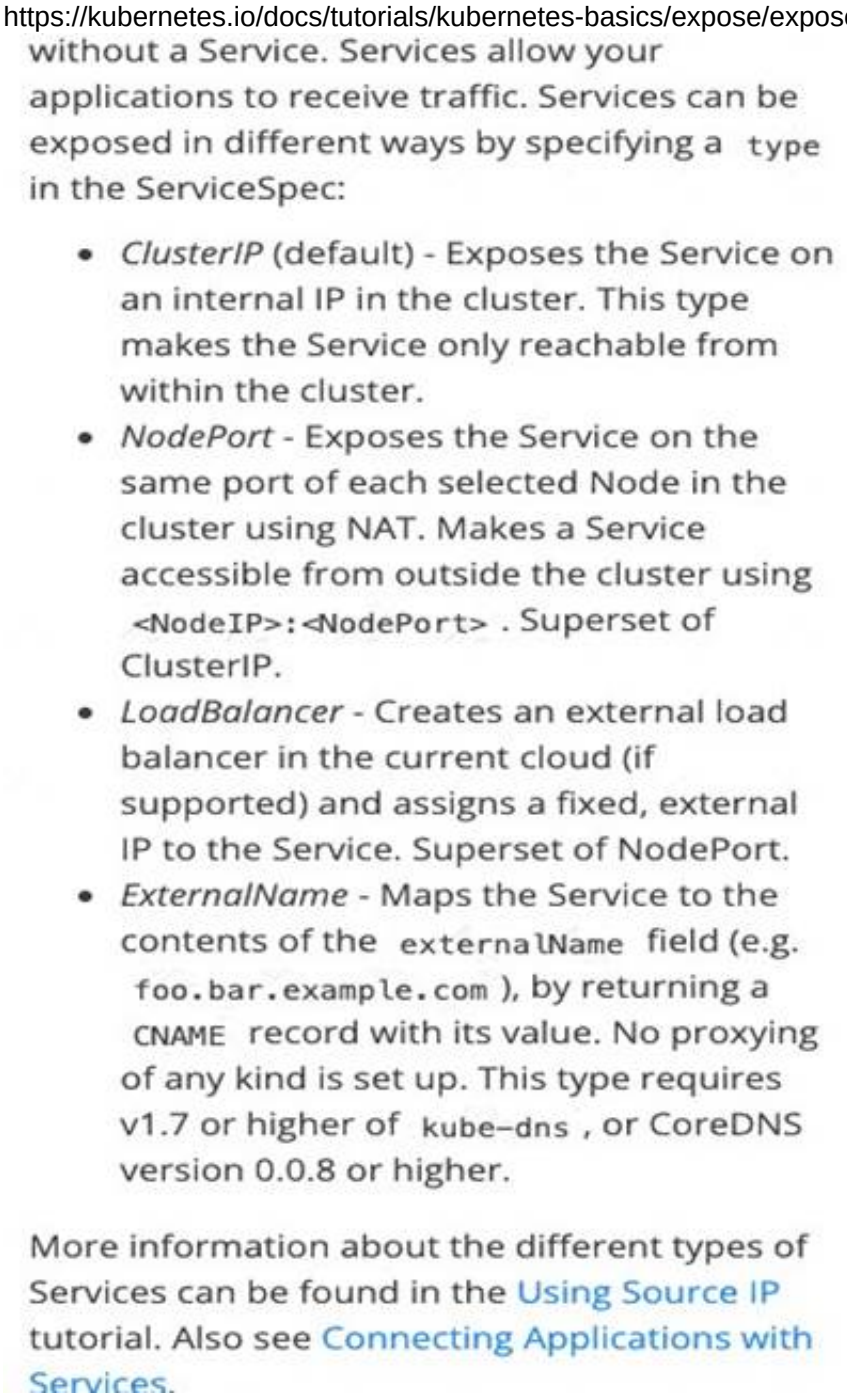
without a Service. Services allow your applications to receive traffic. Services can be exposed in different ways by specifying a `type` in the ServiceSpec:

- *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 38
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 43
Which of the following is not the required field to describe Kubernetes objects?

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

Answer: D

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

Required Fields

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

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

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

NEW QUESTION 48
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 50
You might need to run a stateless application in kubernetes, and you want to be able to scale easily and perform rolling updates. What kubernetes resource type can you use to do this

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

Answer: C

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

Deployments

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

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

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

NEW QUESTION 53
How can persistent volume be provisioned?

A. Automatically
B. Bootstrap
C. Dynamically

Answer: C

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

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

NEW QUESTION 54
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 56
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

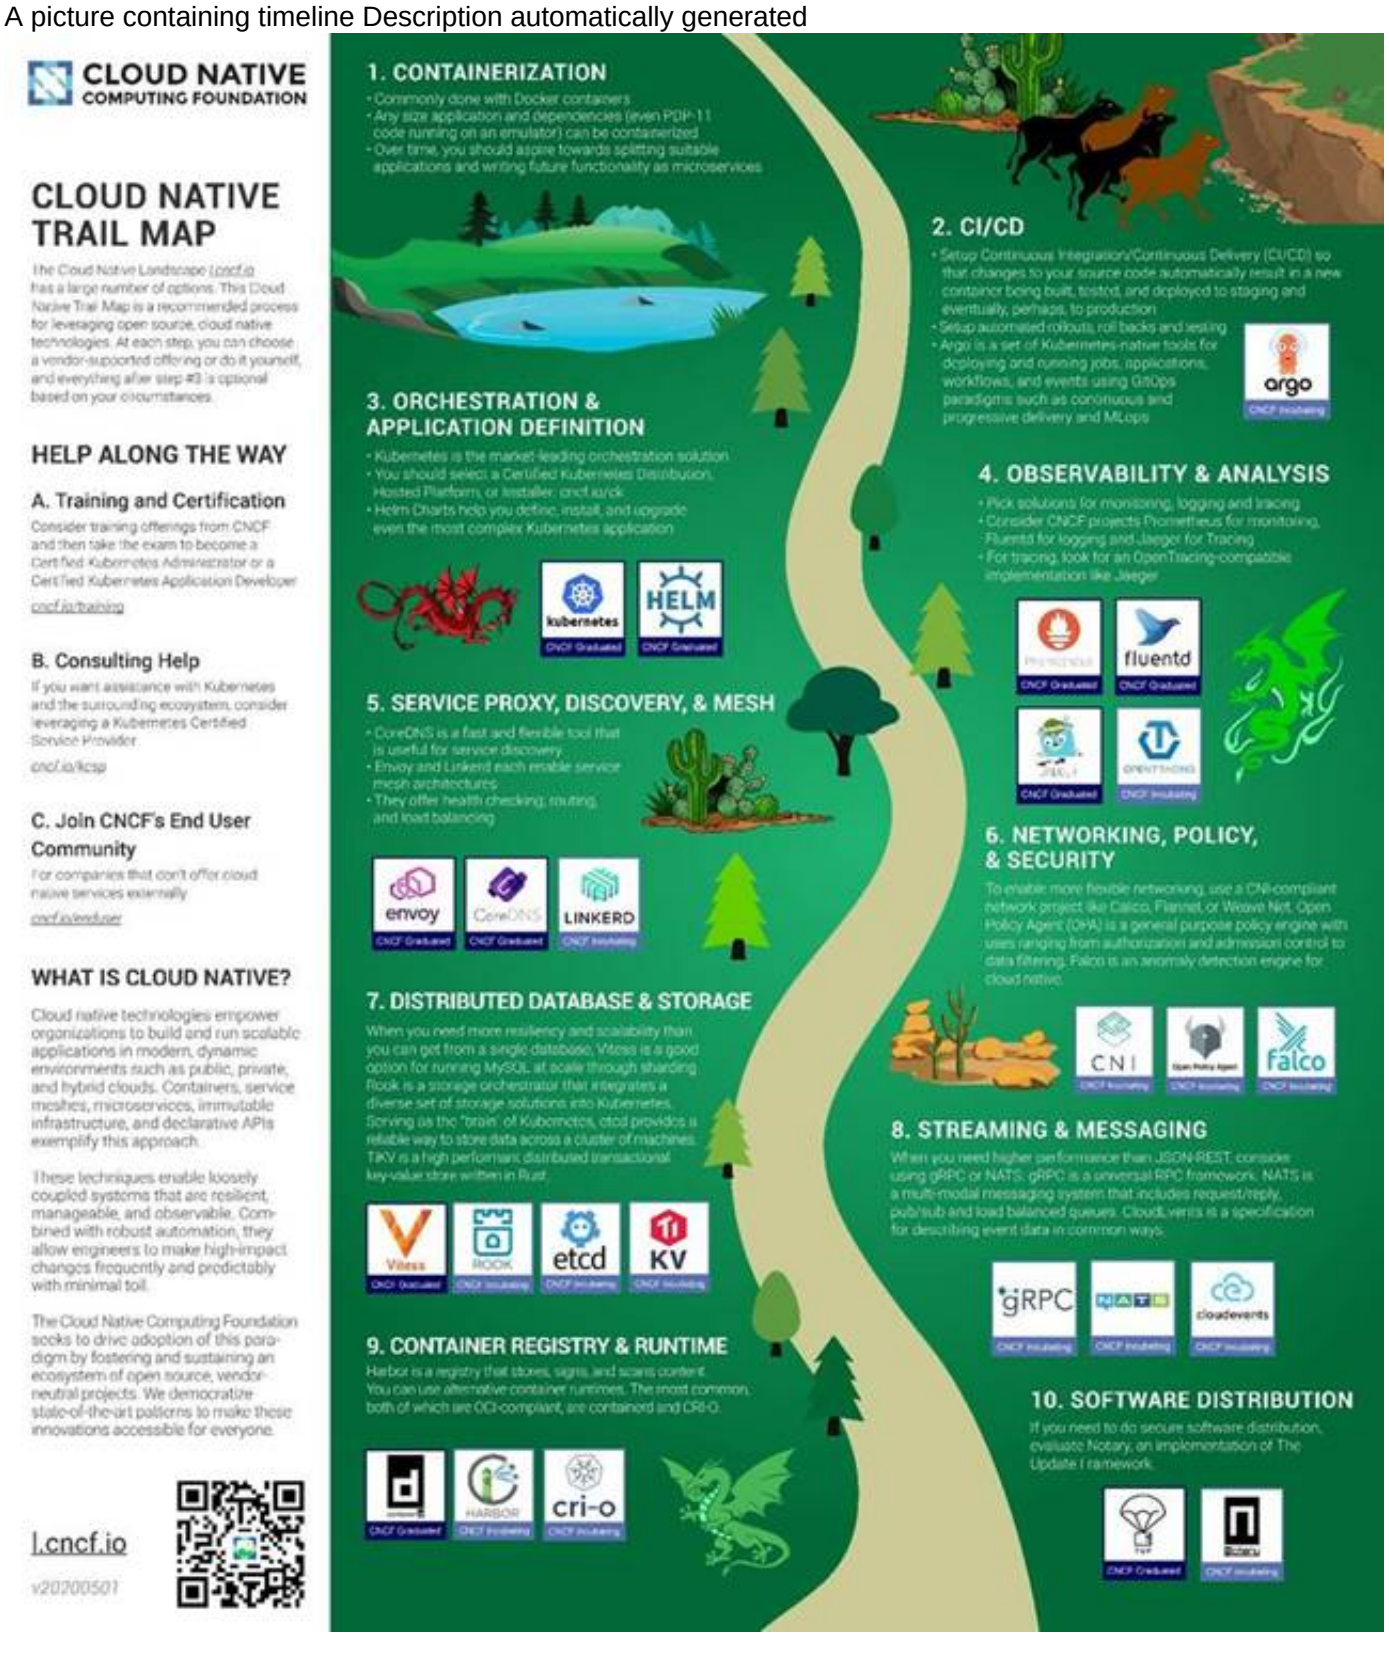

NEW QUESTION 59
Which is NOT a use case for the Kubernetes dashboard?

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 61
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 63
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 66
Which organizational persona creates Service Level Agreements 'SLA', Service Level Objectives 'SLO', and Service Level Indicator 'SLI'?

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

Answer: C

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

NEW QUESTION 70
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 74
Various Container Orchestrator Systems (COS)?

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

Answer: ACD

NEW QUESTION 77
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 78
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 83
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 87
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

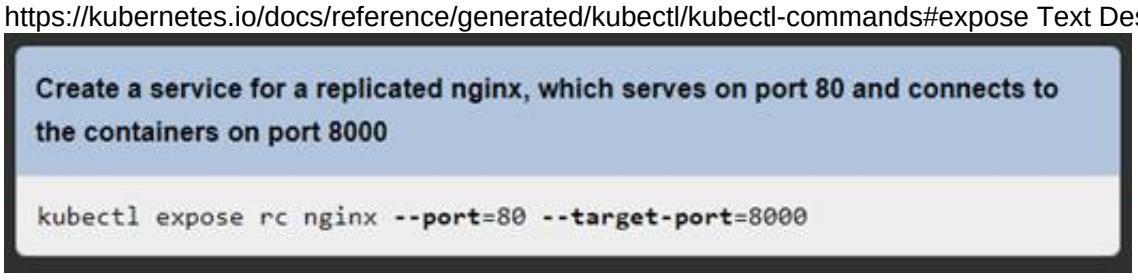

NEW QUESTION 91
.....

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