Linux-Foundation

Exam Questions KCNA

Kubernetes and Cloud Native Associate (KCNA)

NEW QUESTION 1
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 2
Which of the following are not the metrics for Site Reliability Engineering?

A. Service Level Objectives 'SLO'
B. Service Level Agreements 'SLA'
C. Service Level Indicators 'SLI'
D. Service Level Definition 'SLD'

Answer: D

Explanation:
SLI defined quantitative measure of some aspect of the level of service that is provided.
SLOs are key to making data-driven decisions about reliability, they're at the core of SRE practic-es.
SLAs an explicit or implicit contract with your users that includes consequences of meeting (or missing) the SLOs they contain.

NEW QUESTION 3
What command use to get documentation about kubernetes resource type

A. alias k='kubectl' k api-resources
B. alias k='kubectl' k api-list
C. alias k='kubectl' k explain
D. alias k='kubectl' k get resource

Answer: C

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

explain

List the fields for supported resources.

This command describes the fields associated with each supported API resource. Fields are identified via a simple JSONPath identifier:

<type>.<fieldName>[.<fieldName>]

Add the --recursive flag to display all of the fields at once without descriptions. Information about each field is retrieved from the server in OpenAPI format.

Use "kubectl api-resources" for a complete list of supported resources.

Usage

```
$ kubectl explain RESOURCE
```

Get the documentation of the resource and its fields

```
kubectl explain pods
```

Get the documentation of a specific field of a resource

```
kubectl explain pods.spec.containers
```

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

Pods

Pods are the smallest deployable units of computing that you can create and manage in Kubernetes.

A *Pod* (as in a pod of whales or pea pod) is a group of one or more containers, with shared storage and network resources, and a specification for how to run the containers. A Pod's contents are always co-located and co-scheduled, and run in a shared context. A Pod models an application-specific "logical host": it contains one or more application containers which are relatively tightly coupled. In non-cloud contexts, applications executed on the same physical or virtual machine are analogous to cloud applications executed on the same logical host.

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

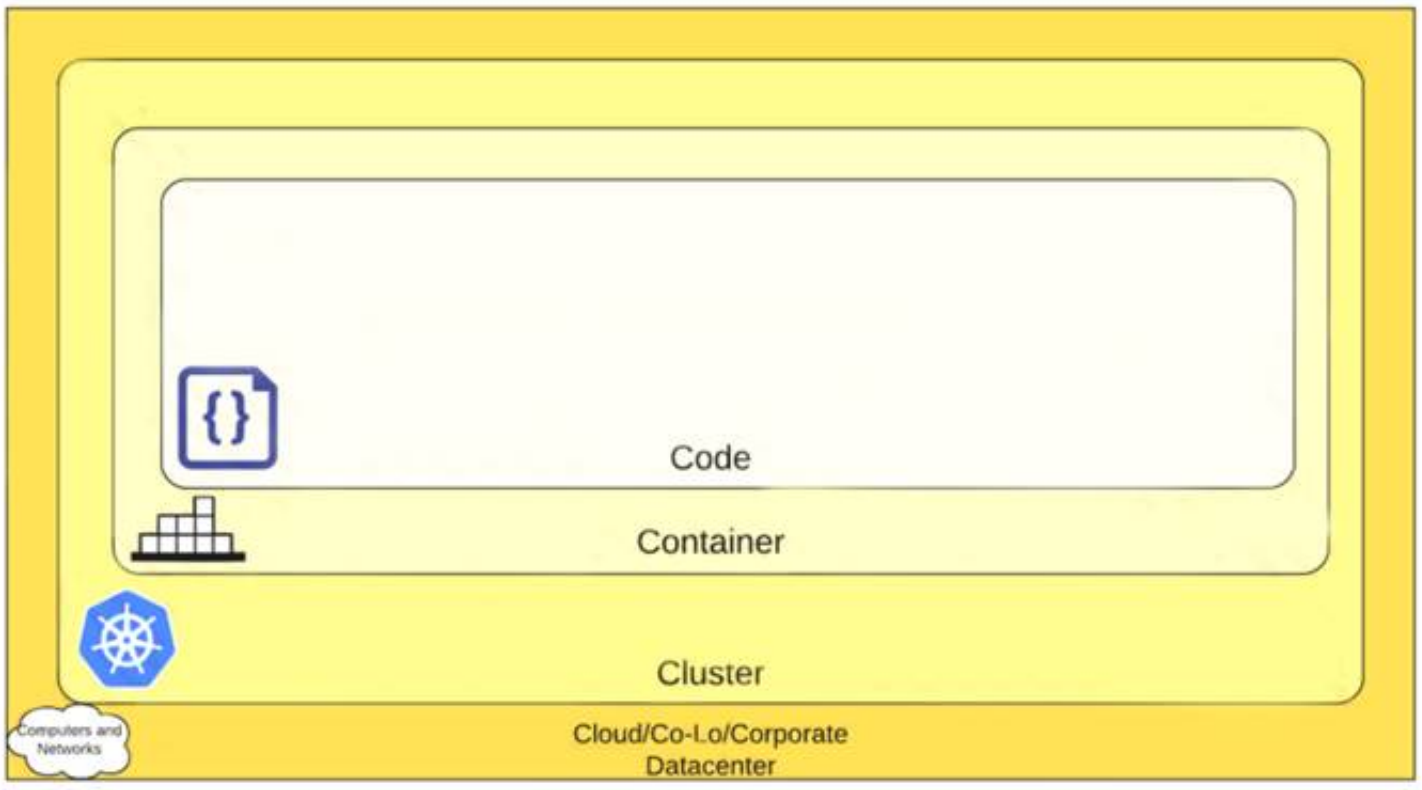

NEW QUESTION 6
What tool allows us to build useful visual representations of prometheus data?

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 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

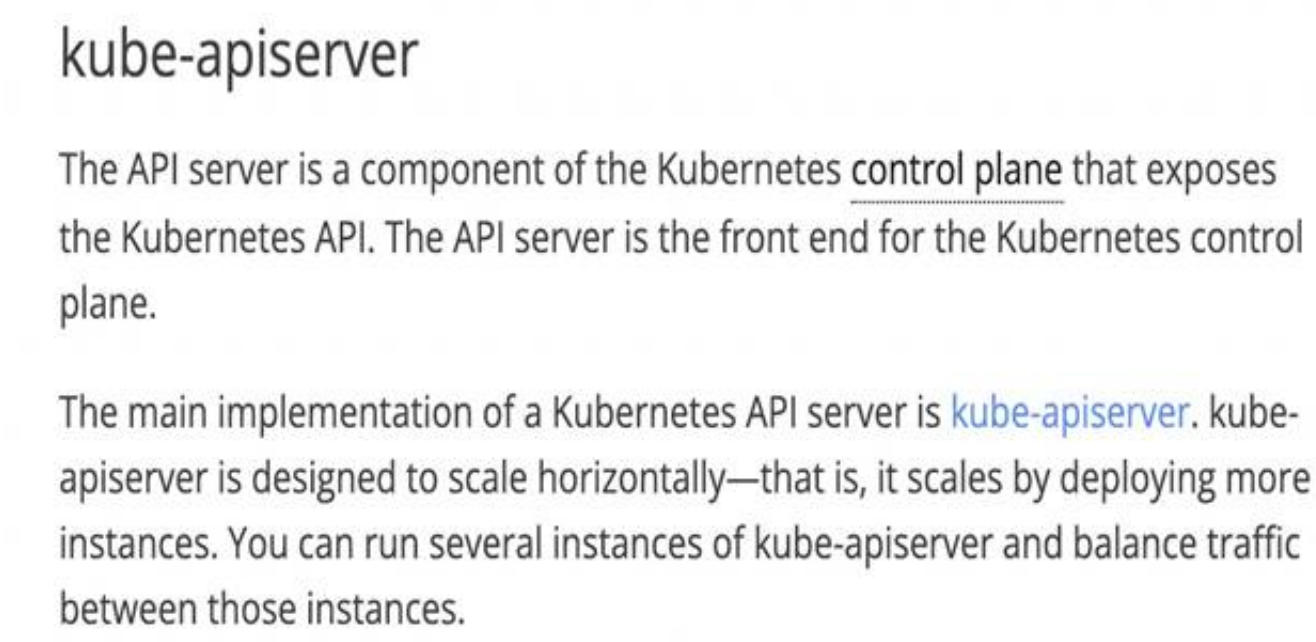

NEW QUESTION 8
Stateful set requires which service for the network identity of pods?

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

Answer: C

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

Using StatefulSets

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

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

NEW QUESTION 9
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 10
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
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 15
Which project in this list is a leading project in the observability space?

A. Jaeger
B. Vitess
C. Argo
D. Kubernetes

Answer: A

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

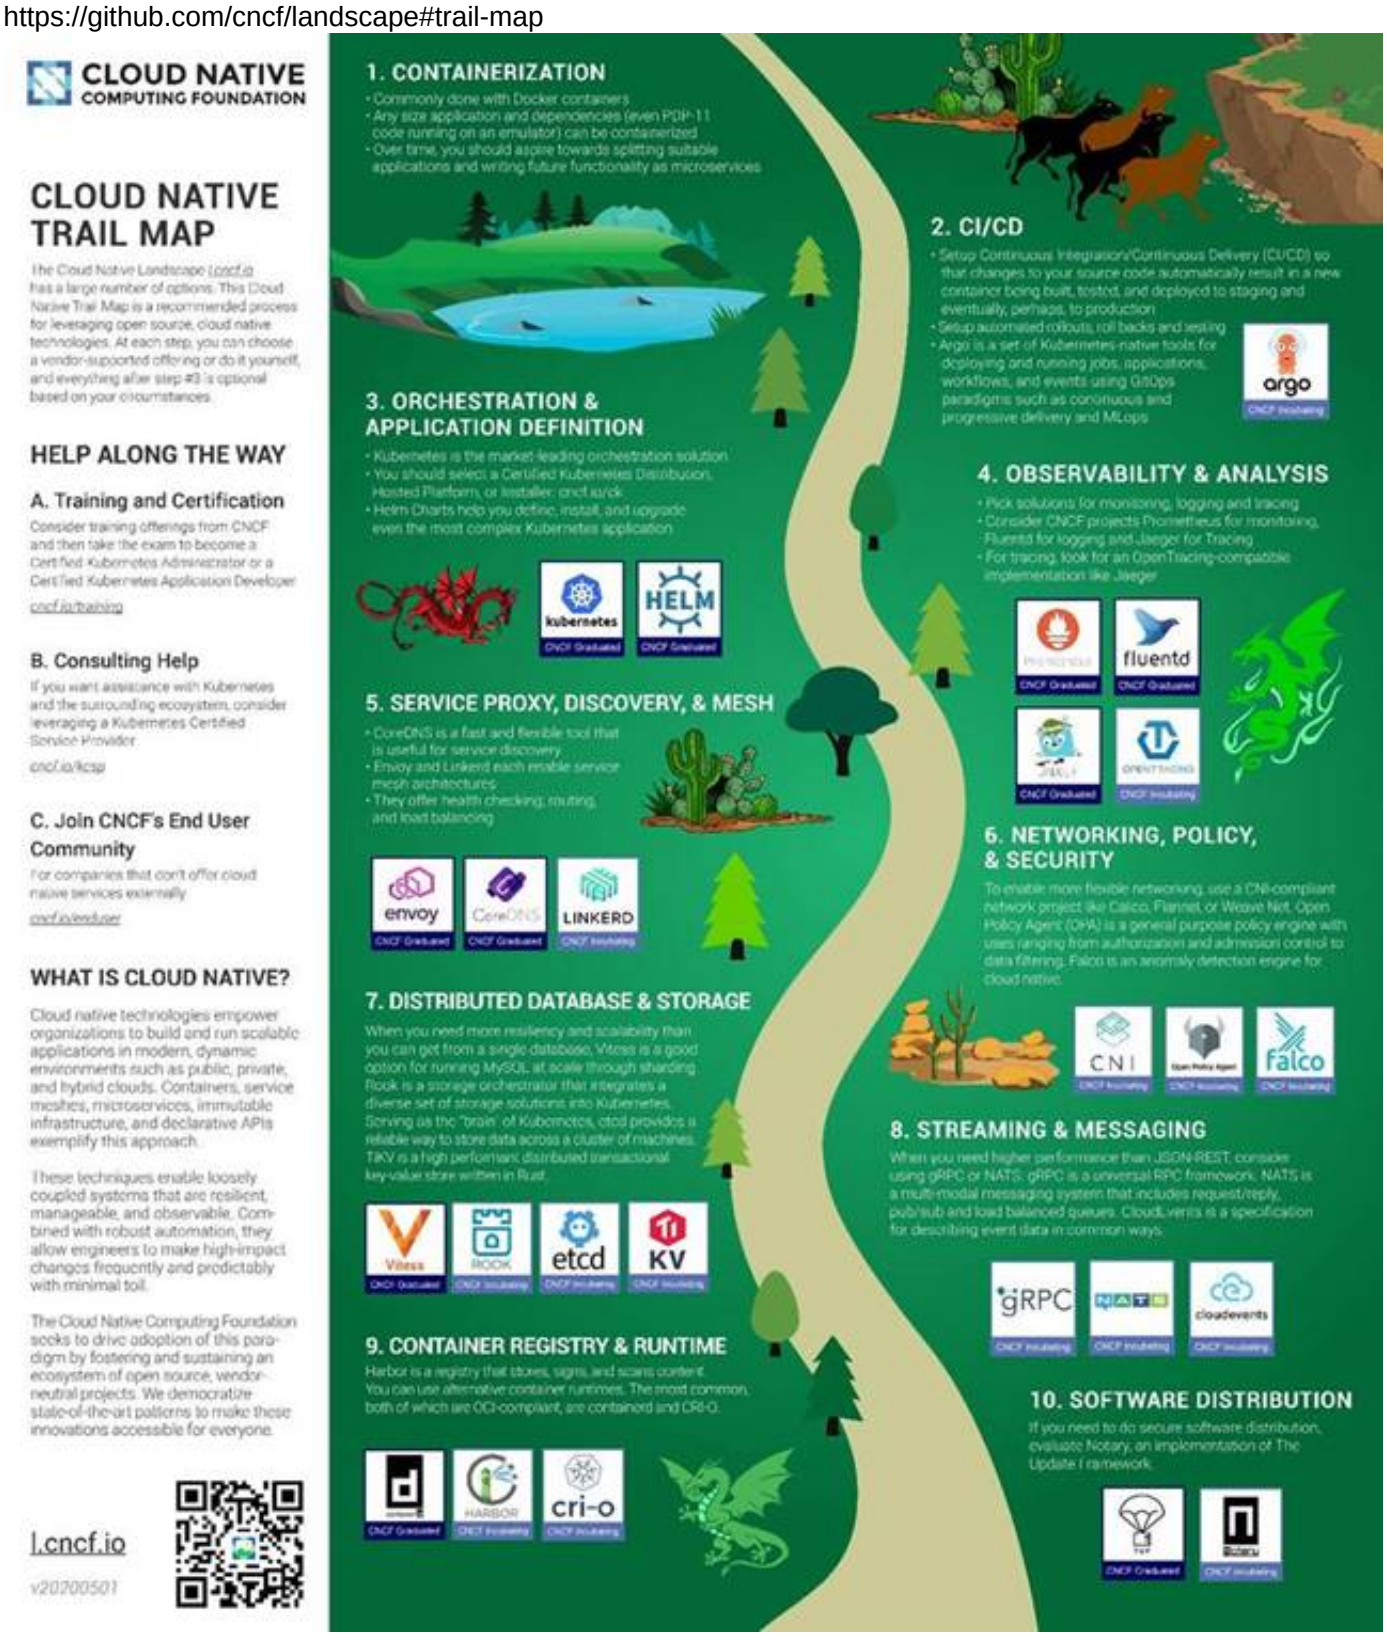

NEW QUESTION 20
What tool allows you to create self-managing, self-scaling, self-healing storage?

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 25
Which authentication method allows JWTs to authenticate?

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

Answer: A

NEW QUESTION 30
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 35
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

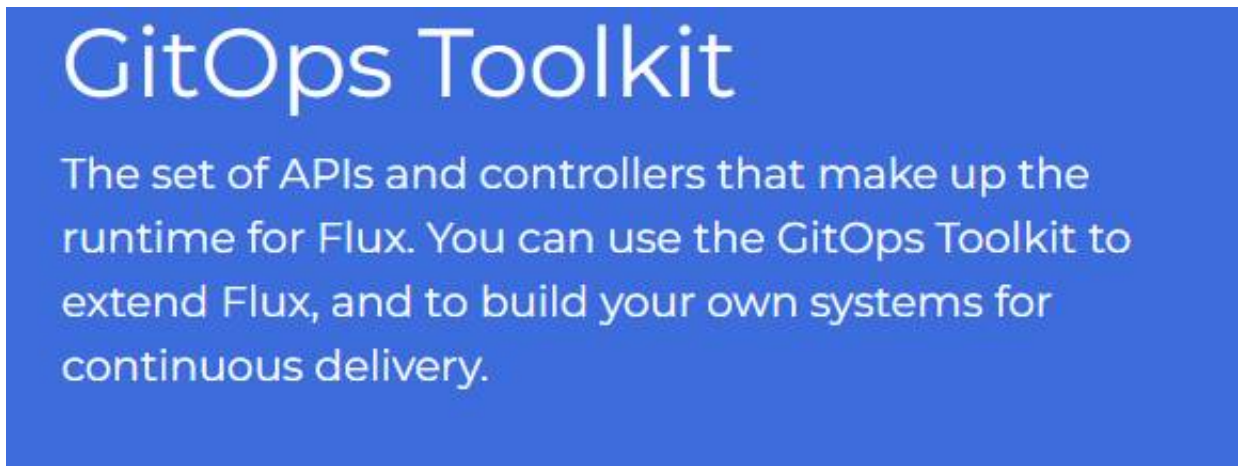

Note: Argo CD is a GitOps tool and not using GitOps toolkit

NEW QUESTION 39
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 41
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

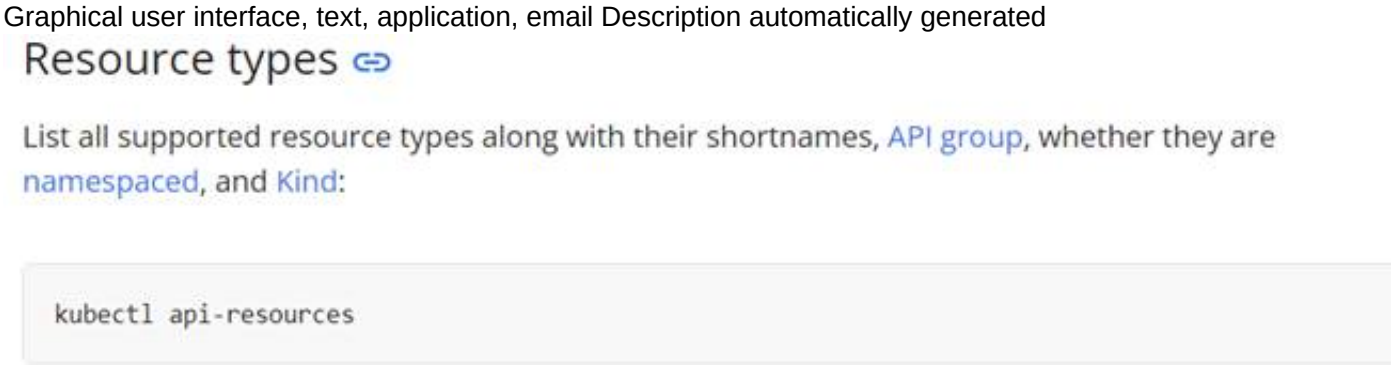

NEW QUESTION 43
Which of the following is not the Kubernetes AutoScaling Strategy?

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 44
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 47
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 48
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

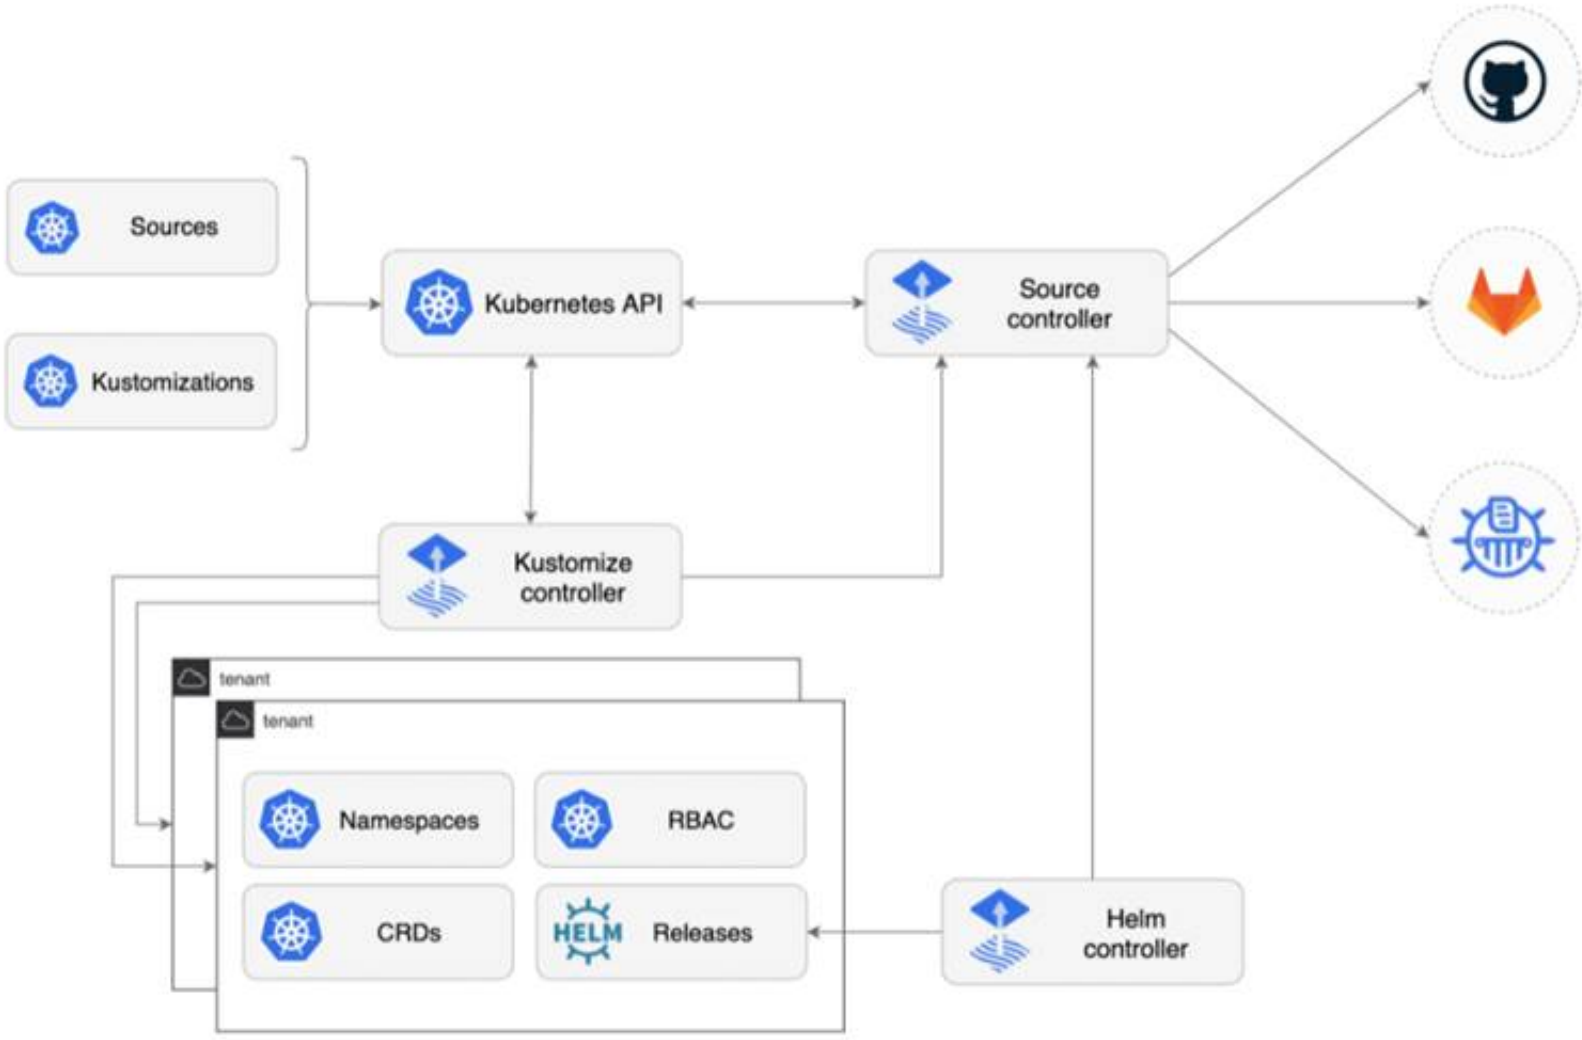

NEW QUESTION 49
Flux is built using which toolkit?

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 54
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

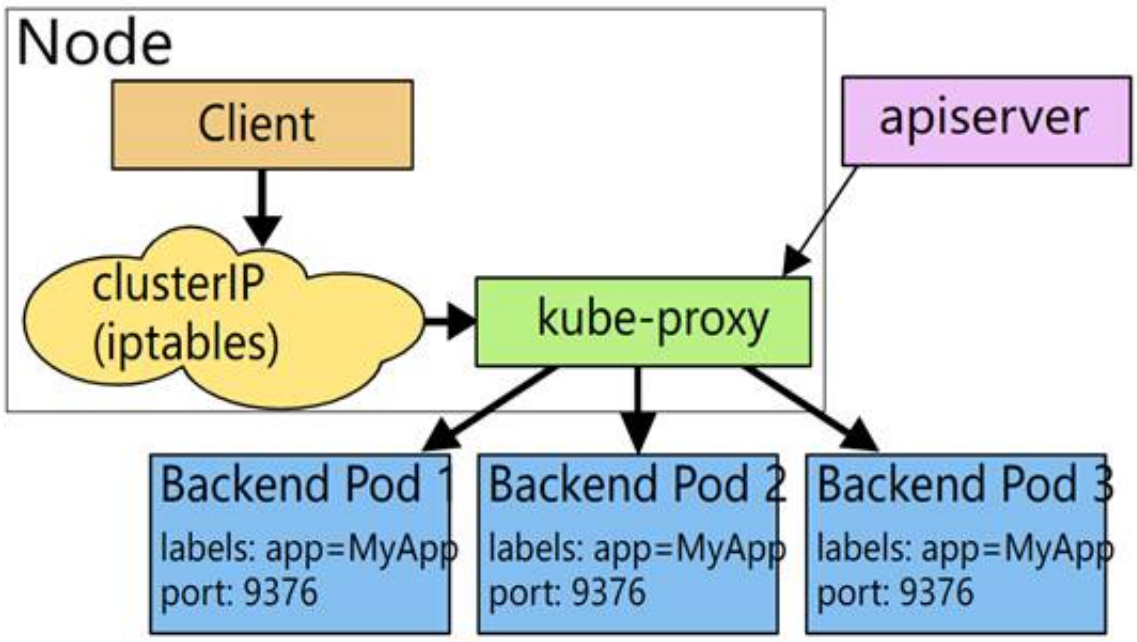

NEW QUESTION 59
What feature is used for selecting the container runtime configuration?

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 64
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 67
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 71
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 76
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 81
Observability and monitoring are not the same?

A. True
B. False

Answer: A

NEW QUESTION 82
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 86
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 88
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

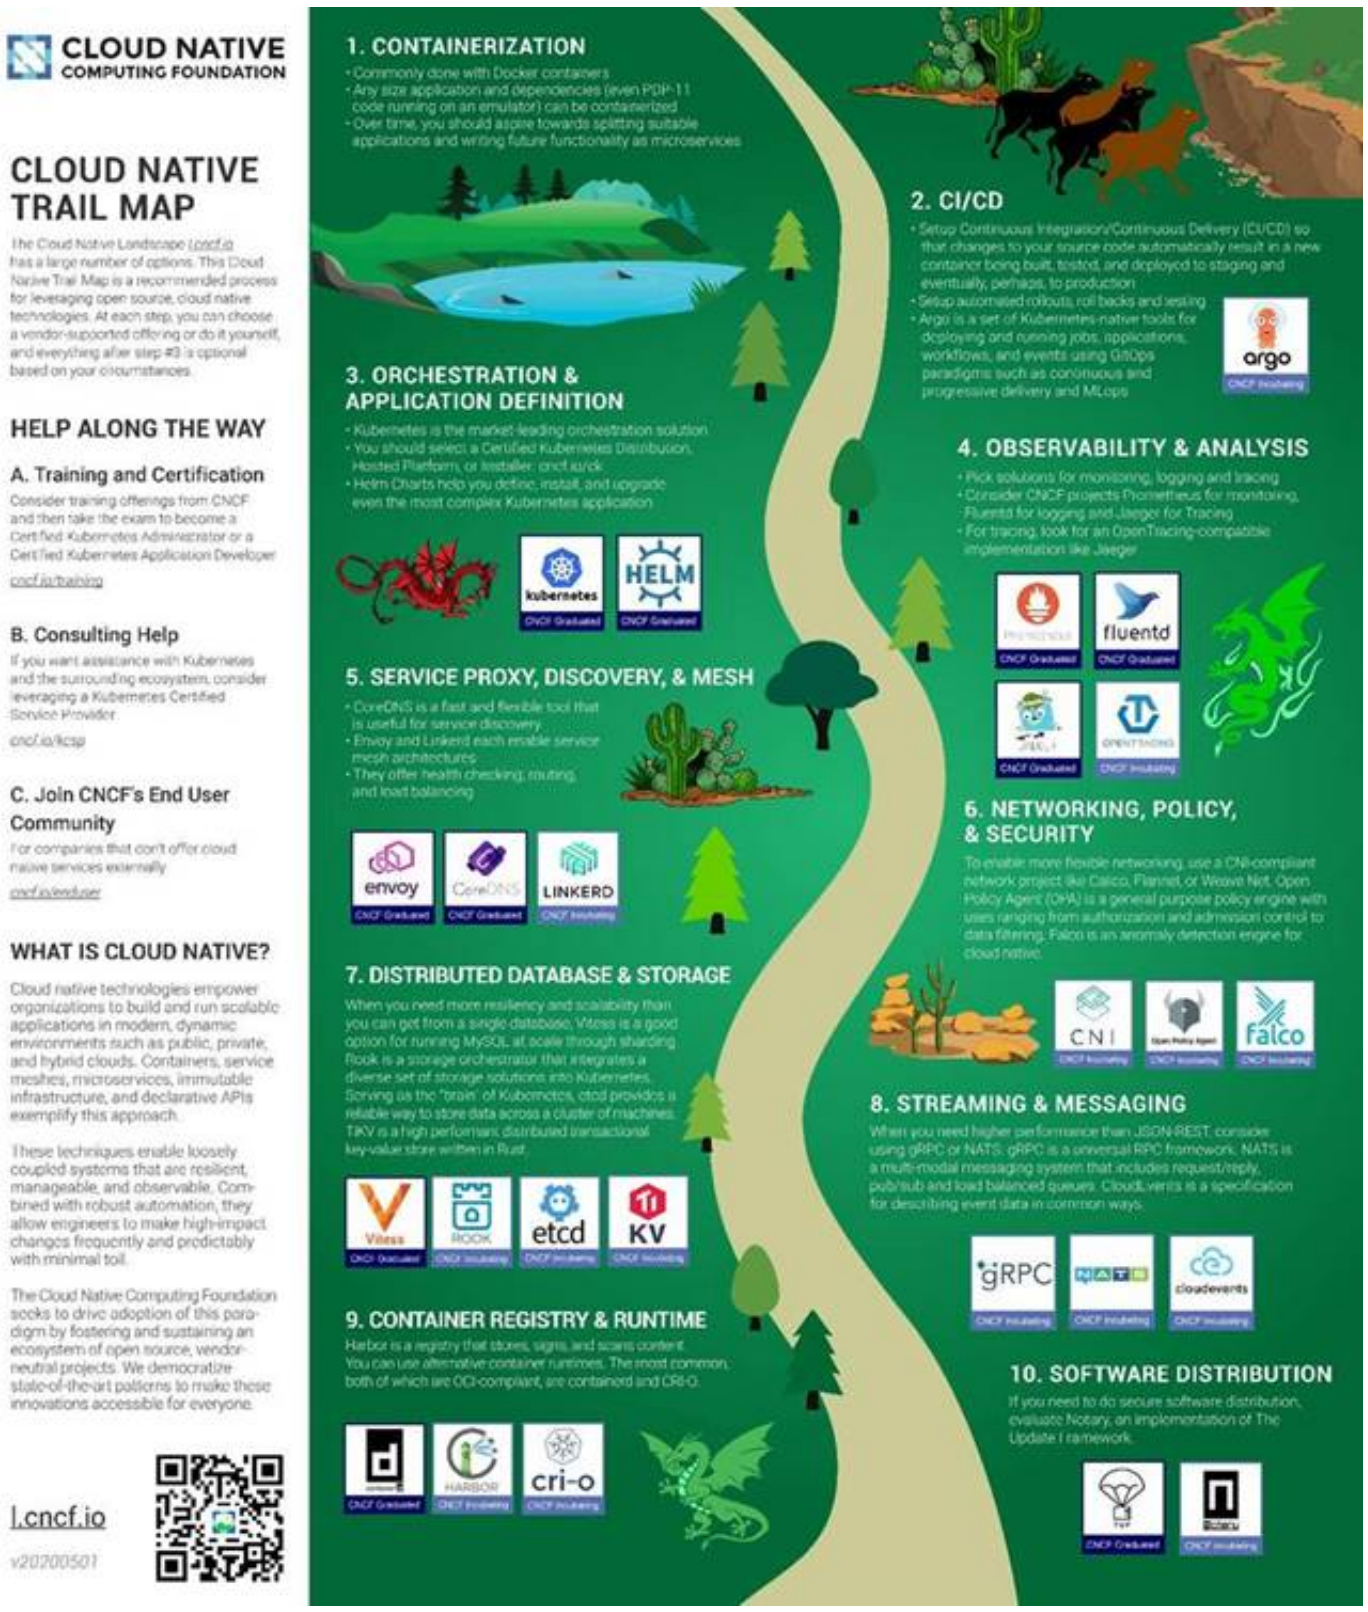

NEW QUESTION 93
Which is NOT a use case for the Kubernetes dashboard?

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 95
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 96
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 98
In distributed system tracing, is the term used to refer to a request as it passes through a single com-ponent 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 100
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 103
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 108
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

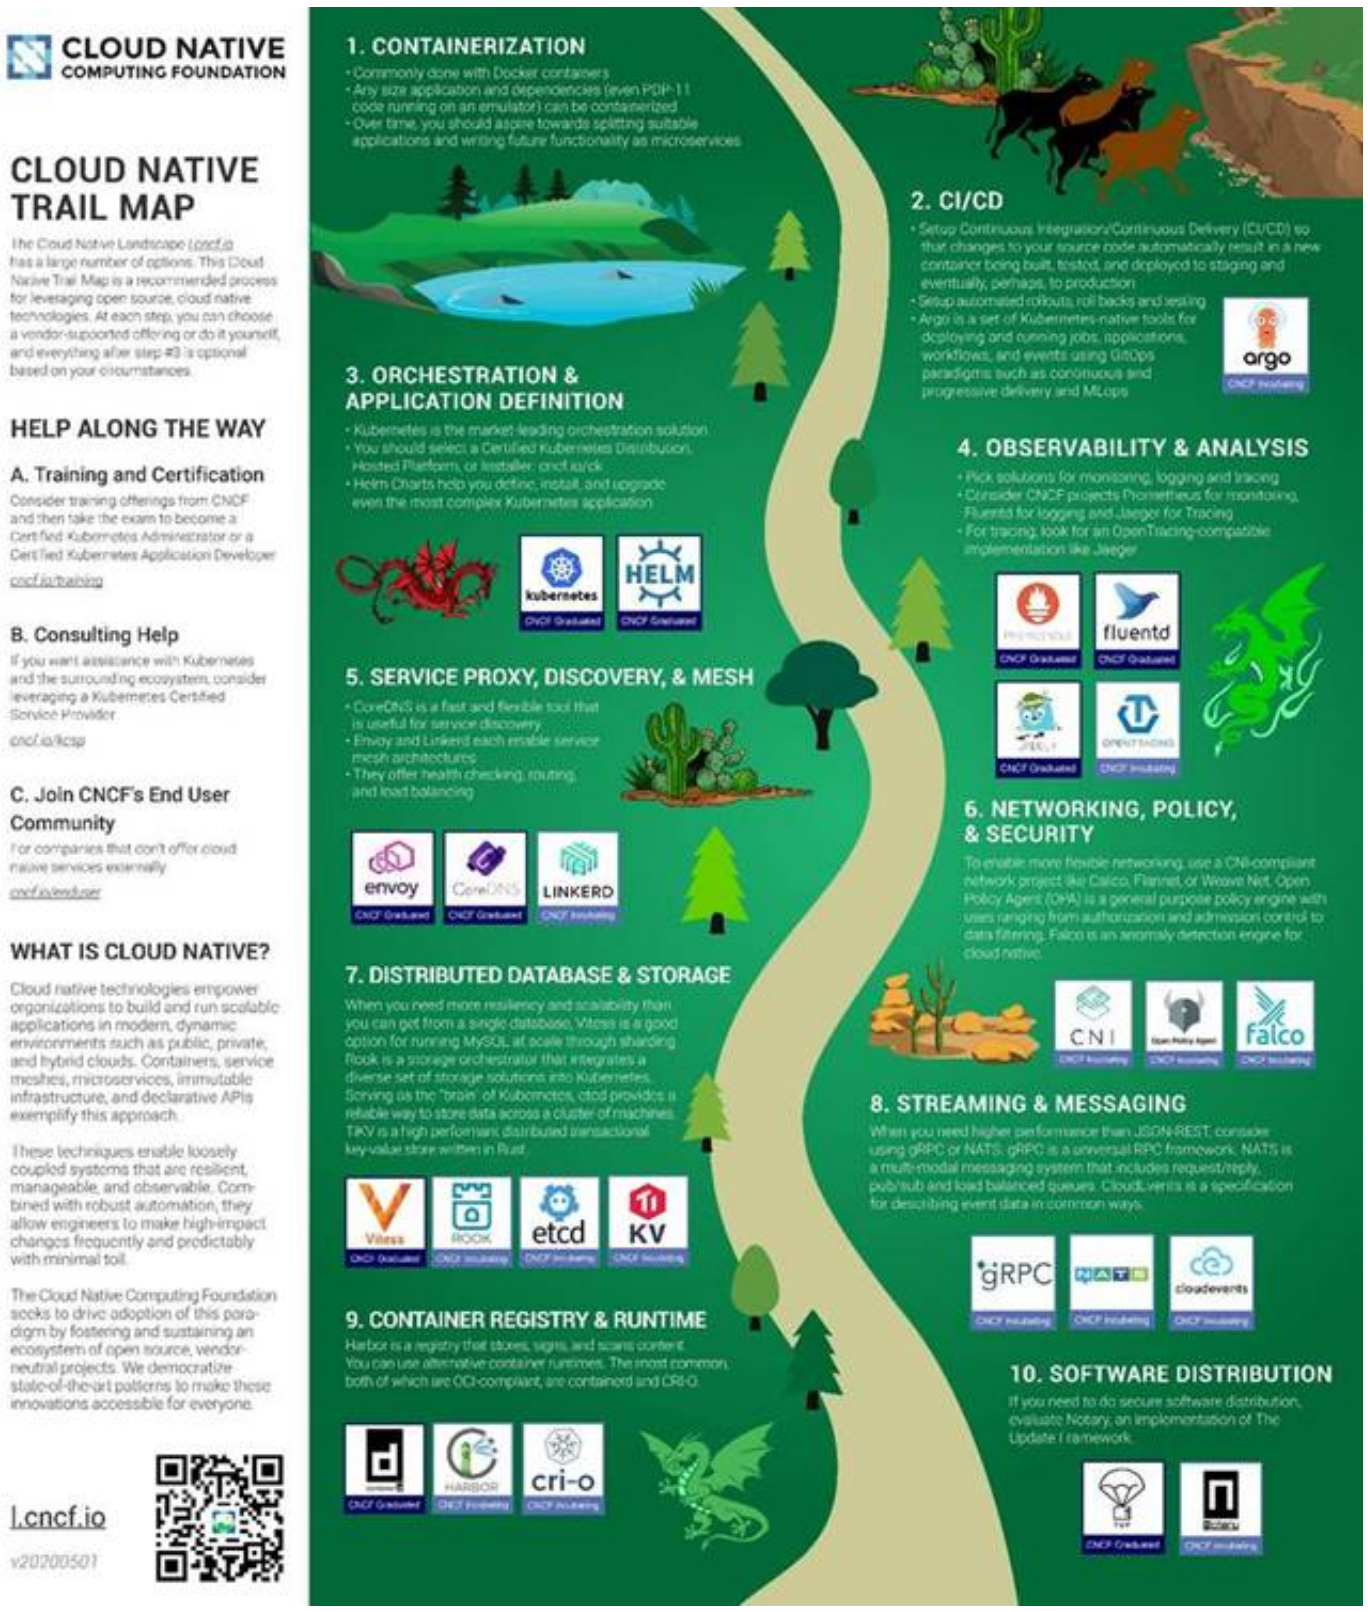

NEW QUESTION 111
Which access control component of Kubernetes is responsible for authorization and decides what requestor is allowed to do?

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 116
Various Container Orchestrator Systems (COS)?

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

Answer: ACD

NEW QUESTION 118
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 120
Which of the following are characteristics of Statefulsets?

A. Ordered, graceful deployment and scaling
B. Creates replica sets
C. Uses headless services

Answer: A

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

Using StatefulSets

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

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

NEW QUESTION 122
What is the functionality of the daemon set?

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

Answer: A

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

DaemonSet

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

Some typical uses of a DaemonSet are:

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

NEW QUESTION 126
How to get the logs of the previously terminated nginx container from the web pod?

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

Answer: A

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

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

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

NEW QUESTION 127
'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

```
cat pod.json | kubectl delete -f -
```

Delete pods and services with same names "baz" and "foo"

```
kubectl delete pod,service baz foo
```

Delete pods and services with label name=myLabel

```
kubectl delete pods,services -l name=myLabel
```

Delete a pod with minimal delay

```
kubectl delete pod foo --now
```

Force delete a pod on a dead node

```
kubectl delete pod foo --force
```

Delete all pods

```
kubectl delete pods --all
```

NEW QUESTION 131
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 136
Which command-line tool is used to interact with the Kubernetes cluster?

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

Answer: B

Explanation:
https://kubernetes.io/docs/reference/kubectl/
Graphical user interface, text, application, email Description automatically generated

Command line tool (kubectl)

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

This tool is named `kubectl`.

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

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

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

NEW QUESTION 137
.....

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