



KCNA^{Q&As}

Kubernetes and Cloud Native Associate (KCNA)

Pass Linux Foundation KCNA Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/kcna.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Linux Foundation Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

Which authentication method allows JWTs to authenticate?

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

Correct Answer: A

QUESTION 2

What is the command used to scale the application?

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

Correct Answer: C

Explanation: [https://kubernetes.io/docs/reference/generated/kubectl/kubectl- commands#scale](https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#scale)



scale

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

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

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

Usage

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

example

Scale a replica set named 'foo' to 3

```
kubectl scale --replicas=3 rs/foo
```

Scale a resource identified by type and name specified in "foo.yaml" to 3

```
kubectl scale --replicas=3 -f foo.yaml
```

If the deployment named mysql's current size is 2, scale mysql to 3

```
kubectl scale --current-replicas=2 --replicas=3 deployment/mysql
```

Scale multiple replication controllers

```
kubectl scale --replicas=5 rc/foo rc/bar rc/c
```

QUESTION 3

What Linux feature is used to provide isolation for containers?

- A. Processes
- B. Services
- C. NetworkPolicy
- D. Control groups

Correct Answer: D

Explanation: Control groups provide isolation for container processes, keeping them separate from other processes on the host.

QUESTION 4

What is FinOps?

- A. The first step in any cloud transformation
- B. Stage beyond DevOps or DevSecOps, where organization transition to serverless technologies
- C. Using data to make cost savings decisions about cloud usage



D. Specialized cloud features used by financial industries (example: banks, insurance, etc)

Correct Answer: C

Explanation: <https://www.servicenow.com/products/it-asset-management/what-is-finops.html>

What is the origin of FinOps?

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

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

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

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

QUESTION 5

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

Correct Answer: B

Explanation: <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>



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
```

[KCNA VCE Dumps](#)

[KCNA Study Guide](#)

[KCNA Exam Questions](#)