



# KCNA<sup>Q&As</sup>

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.pass4itsure.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

In distributed system tracing, is the term used to refer to a request as it passes through a single component of the distributed system?

- A. Log
- B. Span
- C. Trace
- D. Bucket

Correct Answer: B

Explanation: [https://www.splunk.com/en\\_us/data-insider/what-is-distributed-tracing.html](https://www.splunk.com/en_us/data-insider/what-is-distributed-tracing.html)



## 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.



Text, letter

**QUESTION 2**

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

Correct Answer: E

SaaS (Software as a Service)	FaaS (Functions as a Service)	PaaS (Platform as a Service)	CaaS (Container as a Service)	IaaS (Infrastructure as a Service)	On-Prem (private cloud)	
Functions	Functions	Functions	Functions	Functions	Functions	Cloud Service Provider Responsible
Applications	Applications	Applications	Applications	Applications	Applications	
Runtime	Runtime	Runtime	Runtime	Runtime	Runtime	Customer Responsible
Middleware or Containers	Middleware or Containers	Middleware or Containers	Middleware or Containers	Middleware or Containers	Middleware or Containers	
Operating System	Operating System	Operating System	Operating System	Operating System	Operating System	Customer and Cloud Service Provider have Shared Responsibility
Virtualization	Virtualization	Virtualization	Virtualization	Virtualization	Virtualization	
Servers	Servers	Servers	Servers	Servers	Servers	
Storage	Storage	Storage	Storage	Storage	Storage	
Networking	Networking	Networking	Networking	Networking	Networking	

**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



Correct Answer: C

Explanation: AgroCD: <https://argo-cd.readthedocs.io/en/stable/developer-guide/ci/#can-i-retrigger-thechecks-without-pushing-a-new-commit> FluxCD: <https://fluxcd.io/>

---

#### QUESTION 4

Continuous delivery is \_\_\_\_\_.

- A. Manually deploying the code
- B. Coding, Building and Testing the code
- C. Automatically deploying code to [container or server] environment

Correct Answer: C

---

#### QUESTION 5

To specify a Kubernetes object which language is used?

- A. JSON
- B. Go
- C. YAML
- D. Node
- E. Python

Correct Answer: C

Explanation: <https://kubernetes.io/docs/concepts/overview/working-with-objects/kubernetes-objects/>

# Understanding Kubernetes Objects

This page explains how Kubernetes objects are represented in the Kubernetes API, and how you can express them in `.yaml` format.