

Question: 1

A developer has created an application based on customer requirements. The customer needs to run the application with the minimum downtime. Which design approach regarding high-availability applications, Recovery Time Objective, and Recovery Point Objective must be taken?

- A. Active/passive results in lower RTO and RPO. For RPO, data synchronization between the two data centers must be timely to allow seamless request flow.
- B. Active/passive results in lower RTO and RPO. For RPO, data synchronization between the two data centers does not need to be timely to allow seamless request flow.
- C. Active/active results in lower RTO and RPO. For RPO, data synchronization between the two data centers does not need to be timely to allow seamless request flow.
- D. Active/active results in lower RTO and RPO. For RPO, data synchronization between the two data centers must be timely to allow seamless request flow.

Answer: A

Question: 2

DRAG DROP

An application is being built to collect and display telemetry streaming data.

Drag and drop the elements of this stack from the left onto the correct element functions on the right.

IOS-XE Device: IOS-XE Device	visualization platform
Elasticsearch: Elasticsearch	data collector
Kibana: Kibana	data generator
Python Application: Python Application	datastore

Answer:

Kibana: Kibana

Python Application: Python Application

IOS-XE Device: IOS-XE Device

Elasticsearch: Elasticsearch

Question: 3

A cloud native project is being worked on in which all source code and dependencies are written in Python, Ruby, and/or JavaScript. A change in code triggers a notification to the CI/CD tool to run the CI/CD pipeline.

Which step should be omitted from the pipeline?

- A. Deploy the code to one or more environments, such as staging and/or production.
- B. Build one or more containers that package up code and all its dependencies.
- C. Compile code.
- D. Run automated tests to validate the correctness.

Answer: C

Question: 4

Which two statements are considered best practices according to the 12-factor app methodology for application design? (Choose two.)

- A. Application code writes its event stream to stdout.
- B. Application log streams are archived in multiple replicated databases.
- C. Application log streams are sent to log indexing and analysis systems.
- D. Application code writes its event stream to specific log files.
- E. Log files are aggregated into a single file on individual nodes.

Answer: AC

Question: 5

An organization manages a large cloud-deployed application that employs a microservices architecture. No notable issues occur with downtime because the services of this application are redundantly deployed over three or more data center regions. However, several times a week reports are received about application slowness. The container orchestration logs show faults in a variety of containers that cause them to fail and then spin up brand new.

Which action must be taken to improve the resiliency design of the application while maintaining current scale?

- A. Update the base image of the containers.
- B. Test the execution of the application with another cloud services platform.
- C. Increase the number of containers running per service.
- D. Add consistent “try/catch(exception)” clauses to the code.

Answer: A

Question: 6

How should a web application be designed to work on a platform where up to 1000 requests per second can be served?

- A. Use algorithms like random early detection to deny excessive requests.
- B. Set a per-user limit (for example, 5 requests/minute/user) and deny the requests from the users who have reached the limit.
- C. Only 1000 user connections are allowed; further connections are denied so that all connected users can be served.
- D. All requests are saved and processed one by one so that all users can be served eventually.

Answer: B

Question: 7

An organization manages a large cloud-deployed application that employs a microservices architecture across multiple data centers. Reports have received about application slowness. The container orchestration logs show that faults have been raised in a variety of containers that caused them to fail and then spin up brand new instances.

Which two actions can improve the design of the application to identify the faults? (Choose two.)

- A. Automatically pull out the container that fails the most over a time period.
- B. Implement a tagging methodology that follows the application execution from service to service.
- C. Add logging on exception and provide immediate notification.
- D. Do a write to the datastore every time there is an application failure.
- E. Implement an SNMP logging system with alerts in case a network link is slow.

Answer: BC

Question: 8

Which two situations are flagged by software tools designed for dependency checking in continuous integration environments, such as OWASP? (Choose two.)

- A. publicly disclosed vulnerabilities related to the included dependencies
- B. mismatches in coding styles and conventions in the included dependencies
- C. incompatible licenses in the included dependencies
- D. test case failures introduced by bugs in the included dependencies
- E. buffer overflows to occur as the result of a combination of the included dependencies

Answer: AE

Question: 9

A network operations team is using the cloud to automate some of their managed customer and branch locations. They require that all of their tooling be ephemeral by design and that the entire automation environment can be recreated without manual commands. Automation code and configuration state will be stored in git for change control and versioning. The engineering high-level plan is to use VMs in a cloud- provider environment then configure open source tooling onto these VMs to poll, test, and configure the remote devices, as well as deploy the tooling itself.

Which configuration management and/or automation tooling is needed for this solution?

- A. Ansible
- B. Ansible and Terraform
- C. NSO
- D. Terraform
- E. Ansible and NSO

Answer: B

Question: 10

DRAG DROP

Drag and drop the git commands from the left into the correct order on the right to create a feature branch from the master and then incorporate that feature branch into the master.

git branch -d feature	step 1
git checkout -b feature master	step 2
git checkout master	step 3
git push origin master	step 4
git merge --no-ff feature	step 5

Answer:

git checkout -b feature master
git checkout master
git merge --no-ff feature
git branch -d feature
git push origin master