Free Questions for NS0-403 by certscare

Shared by Walton on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are asked to provision storage for a development team to make use of NetApp Trident in their Kubernetes environment. You want to ensure that Trident will not be able to exhaust the storage system resources.

In this scenario, which two actions would you take to satisfy the requirements? (Choose two.)

Options:

A- Limit the maximum volume count.

B- Limit the available data LIFs.

C- Use a shared SVM.

D- Limit the maximum volume size.

Answer:

A

Question 2

Question Type: MultipleChoice

Your customer wants to use NetApp Trident to automate data lake provisioning on premises using a NetApp AFF array in AWS using NetApp Cloud Volumes ONTAP. Their unstructured datasets will require 200 TB of storage, will be accessed by hundreds of worker nodes simultaneously, and will have millions of small files in a complex directory hierarchy.

In this scenario, which Trident driver is appropriate for this task?

Options:

A- ontap-nas-flexgroup

B- ontap-nas-economy

C- ontap-nas

D- aws---cvs

Answer:

A

Question 3

Question Type: MultipleChoice

You are asked to administer your NetApp Cloud Central account settings to manage users, service accounts, workspaces, connectors, and subscriptions.

In this scenario, which REST API would you access?

Options:

A- Cloud Manager API

B- AWS API

C- Tenancy API

D- ONTAP API

Answer:

B

Question 4

Question Type: **MultipleChoice**

What are three REST API methods? (Choose three.)

Options:

A- LIST

B- SET

C- POST

D- DELETE

E- GET

Answer:

C, D, E

Explanation:

https://www.techtarget.com/searchapparchitecture/tip/The-5-essential-HTTP-methods-in-RESTful-API-development

Question 5

Question Type: MultipleChoice

Click the Exhibit button.

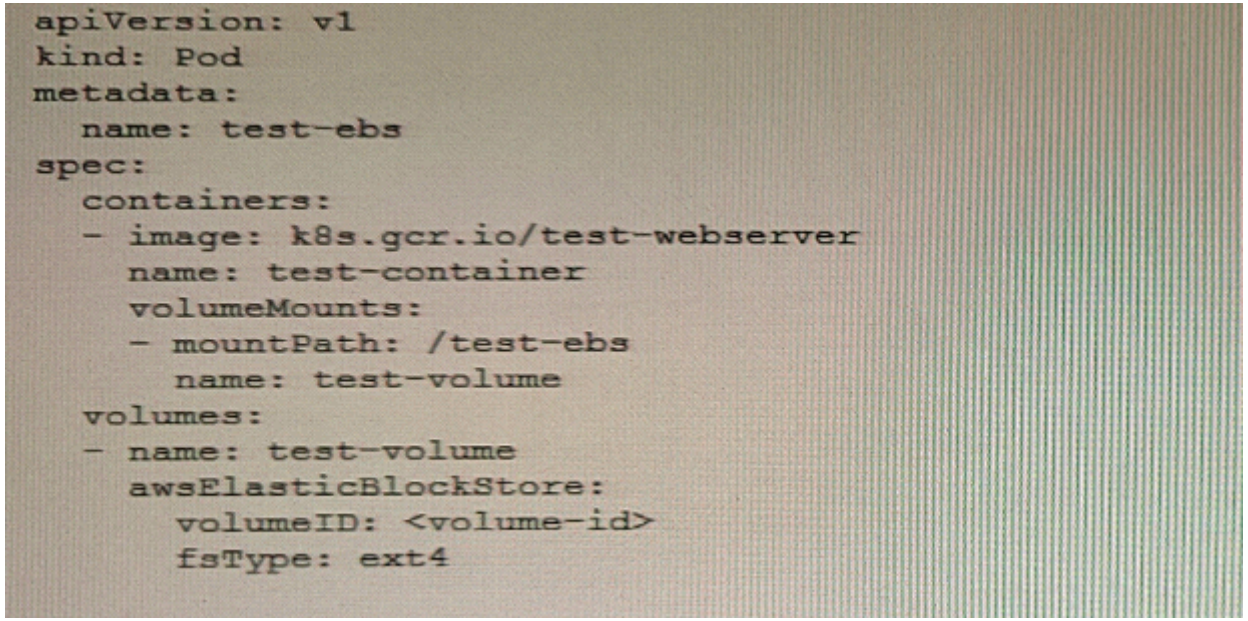

```
apiVersion: v1
kind: Pod
metadata:
  name: test-ebs
spec:
  containers:
  - image: k8s.gcr.io/test-webserver
    name: test-container
    volumeMounts:
    - mountPath: /test-ebs
      name: test-volume
  volumes:
  - name: test-volume
    awsElasticBlockStore:
      volumeID: <volume-id>
      fsType: ext4
```

You are reviewing a Kubernetes manifest to deploy an Amazon Elastic Block Store (AWS EBS).

Which two Kubernetes API object types are shown in the exhibit?

Options:

A- Secret

B- Pod

C- Volume

D- Config

Answer:

D

Question 6

Question Type: **MultipleChoice**

You are using Nvidia DeepOps to deploy Kubernetes clusters on top of NetApp ONTAP AI systems. In this scenario, which automation engine is required?

Options:

A- Ansible

B- Puppet

C- Jenkins

D- Terraform

Answer:

A

Question 7

Question Type: MultipleChoice

You are creating an Ansible playbook to deploy a NetApp ONTAP cluster. You want to take advantage of a loop to automate the creation of several volumes.

What would accomplish this task?

A)

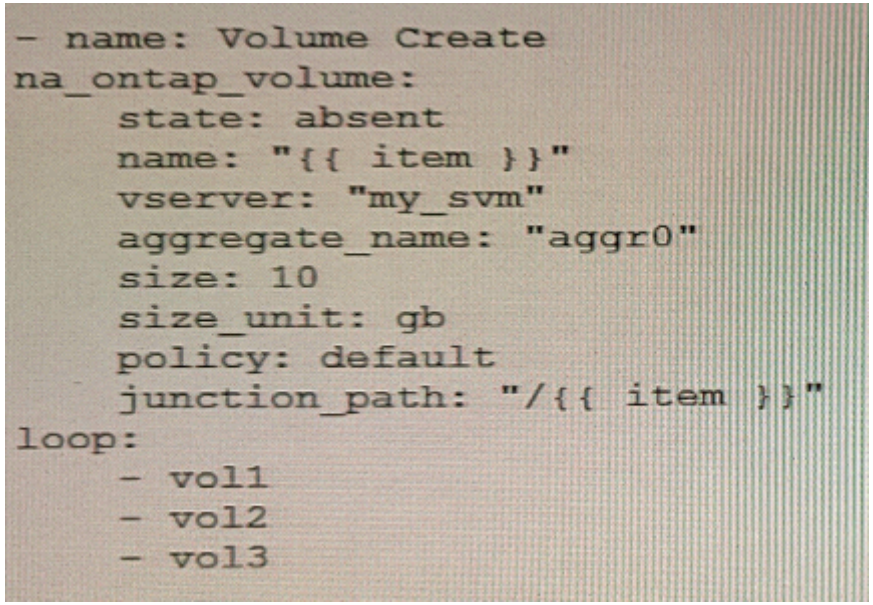

```
- name: Volume Create
na_ontap_volume:
    state: absent
    name: "{{ item }}"
    vserver: "my_svm"
    aggregate_name: "aggr0"
    size: 10
    size_unit: gb
    policy: default
    junction_path: "/{{ item }}"
loop:
    - vol1
    - vol2
    - vol3
```

B)

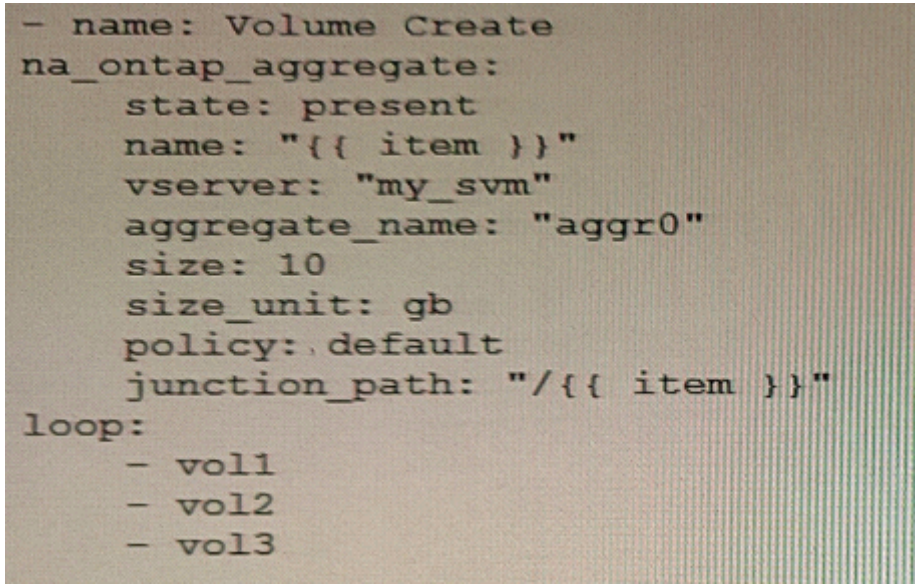

```
- name: Volume Create
na_ontap_aggregate:
    state: present
    name: "{{ item }}"
    vserver: "my_svm"
    aggregate_name: "aggr0"
    size: 10
    size_unit: gb
    policy: default
    junction_path: "/{{ item }}"
loop:
    - vol1
    - vol2
    - vol3
```

C)

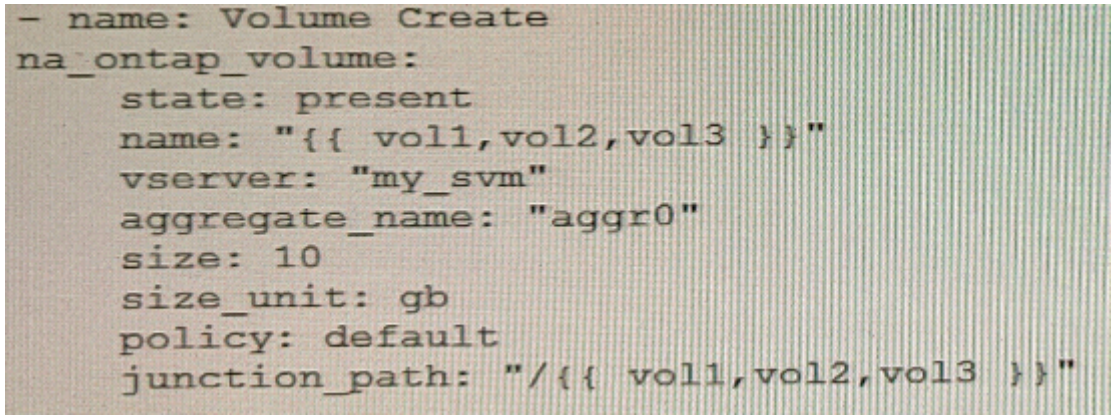

```
- name: Volume Create
na_ontap_volume:
    state: present
    name: "{{ vol1,vol2,vol3 }}"
    vserver: "my_svm"
    aggregate_name: "aggr0"
    size: 10
    size_unit: gb
    policy: default
    junction_path: "/{{ vol1,vol2,vol3 }}"
```

D)

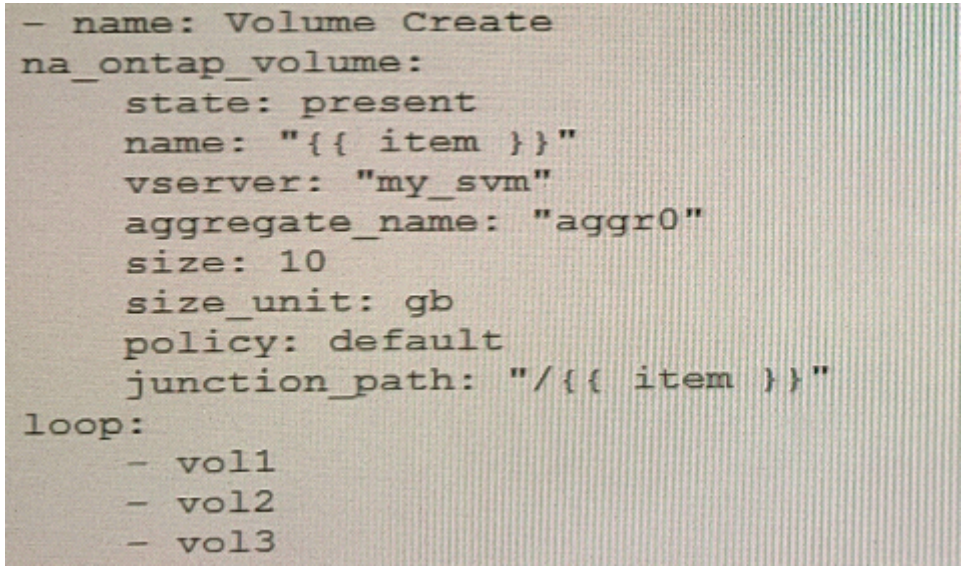

```
- name: Volume Create
na_ontap_volume:
    state: present
    name: "{{ item }}"
    vserver: "my_svm"
    aggregate_name: "aggr0"
    size: 10
    size_unit: gb
    policy: default
    junction_path: "/{{ item }}"
loop:
    - vol1
    - vol2
    - vol3
```

Options:

A- Option A

B- Option B

C- Option C

D- Option D

Answer:

C

Question 8

Question Type: MultipleChoice

You are deploying persistent storage for Kubernetes with NetApp Trident.

In this scenario, what are the two objects that you must create? (Choose two.)

Options:

A- Trident virtual storage pool

B- Trident backend

C- Kubemetes StorageClass

D- Trident volume

Answer:

B, D

Question 9

Question Type: MultipleChoice

Your customer wants to stream large amounts of real-time data from multiple sources at the application level into their Apache Spark platform for transformation before it goes into their machine learning platform.

In this scenario, which technology should be used for this task?

Options:

A- NetApp SnapMirror

B- Apache Kafka

C- NetApp Global File Cache

D- NetApp Cloud Sync service

Answer:

B

Question 10

Question Type: **MultipleChoice**

What are the three pillars of observability in DevOps? (Choose three.)

Options:

A- traces

B- errors

C- metrics

D- versions

E- logs

Answer:

A, C, E

Explanation:

https://www.dynatrace.com/resources/ebooks/three-pillars-of-observability/#:~:text=Observability%20is%20divided%20into%20three,called%20three%20pillars%20of%20observability.

Question 11

Question Type: MultipleChoice

An organization is implementing infrastructure as code to deploy NetApp Cloud Volumes ONTAP instances.

In this scenario, what are two ways in which this implementation improves operation efficiencies? (Choose two.)

Options:

A- Defining the system as code makes it more secure.

B- Best practices can be implemented directly into configurations.

C- Deployment failures are automatically resolved by configuration management.

D- System configurations can be version controlled.

Answer:

A, C