

- **Vendor:** Network Appliance
- **Exam Code:** NS0-403
- **Exam Name:** NetApp Certified Hybrid Cloud Implementation Engineer
- **New Updated Questions from** [Braindump2go](#)
- **(Updated in** [July/2022](#)**)**

[Visit Braindump2go and Download Full Version NS0-403 Exam Dumps](#)

Question: 21

You are using REST APIs to deploy and manage storage that is provided by NetApp Cloud Volumes ONTAP in Microsoft Azure. You are a federated user, and you want to communicate with an authenticated endpoint.

In this scenario, how should you acquire the JWT access token?

- A. Use the API key generator.
- B. Use a JWT access token generator.
- C. Use the refresh token generator.
- D. Use an e-mail address and password

Answer: A

Explanation:

Question: 22

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.)

- 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: AC

Explanation:

Question: 23

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

- A. traces
- B. errors
- C. metrics
- D. versions
- E. logs

Answer: ACE

Explanation:

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

Question: 24

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

[NS0-403 Exam Dumps](#) [NS0-403 Exam Questions](#) [NS0-403 PDF Dumps](#) [NS0-403 VCE Dumps](#)

<https://www.braindump2go.com/ns0-403.html>

machine learning platform.

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

- A. NetApp SnapMirror
- B. Apache Kafka
- C. NetApp Global File Cache
- D. NetApp Cloud Sync service

Answer: B

Explanation:

Question: 25

You are deploying persistent storage for Kubernetes with NetApp Trident.
In this scenario, what are the two objects that you must create? (Choose two.)

- A. Trident virtual storage pool
- B. Trident backend
- C. Kubernetes StorageClass
- D. Trident volume

Answer: BD

Explanation:

Question: 26

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 }}"
  vservers: "my_svm"
  aggregate_name: "aggr0"
  size: 10
  size_unit: gb
  policy: default
  junction_path: "/"
loop:
  - vol1
  - vol2
  - vol3
```

B)

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

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

Question: 27 _____

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

- A. Ansible
- B. Puppet
- C. Jenkins
- D. Terraform

Answer: A

Explanation:

Question: 28

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?

- A. Secret
- B. Pod
- C. Volume
- D. Config

Answer: D

Explanation:

Question: 29

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

- A. LIST
- B. SET

C. POST

D. DELETE

E. GET

Answer: CDE

Explanation:

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

Question: 30

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?

- A. Cloud Manager API
- B. AWSAPI
- C. Tenancy API
- D. ONTAPAPI

Answer: B

Explanation:

Question: 31

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?

- A. ontap-nas-flexgroup
- B. ontap-nas-economy
- C. ontap-nas
- D. aws—cvs

Answer: A

Explanation:

Question: 32

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.)

[NS0-403 Exam Dumps](#) [NS0-403 Exam Questions](#) [NS0-403 PDF Dumps](#) [NS0-403 VCE Dumps](#)

<https://www.braindump2go.com/ns0-403.html>

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

Explanation:

Question: 33

You are asked to provision storage for a development team to make use of NetApp Trident in their Kubernetes environment. You are advised that the development environment will be frequently performing large numbers of deployments. You want to ensure you do not exhaust the storage volume system limits.

In this scenario, which backend would you recommend to satisfy the requirements?

- A. ontap-nas
- B. ontap-san
- C. ontap-nas-economy
- D. ontap-nas- flexgroup

Answer: C

Explanation:

Question: 34

Which DevOps principle implies that your documentation and your configuration are the same?

- A. infrastructure as code
- B. continuous testing
- C. continuous deployment
- D. continuous monitoring

Answer: D

Explanation:

Question: 35

A customer wants their main data set to be stored in their primary location while providing a writable, persistent cache in another location closer to the consumer.

Which two solutions will satisfy the customer's requirements? (Choose two.)

- A. Deploy Cloud Volumes ONTAP with FlexCache in their primary location and ONTAP Select at the consumer site.
- B. Deploy an FAS8300 with FlashCache in their primary location and export the share to the VPN

[NS0-403 Exam Dumps](#) [NS0-403 Exam Questions](#) [NS0-403 PDF Dumps](#) [NS0-403 VCE Dumps](#)

<https://www.braindump2go.com/ns0-403.html>

IPspace.

C. Deploy Cloud Volumes ONTAP in their primary location and ONTAP Select with FlexCache at the consumer site.

D. Deploy Cloud Volumes ONTAP in their primary location and Cloud Volumes ONTAP with FlexCache on a public cloud region nearest to the consumer.

Ans

wer: AC

Explanation:
n: