



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



NS0-403 MCQs
NS0-403 TestPrep
NS0-403 Study Guide
NS0-403 Practice Test
NS0-403 Exam Questions



Network-Appliance

NS0-403

NetApp Certified Hybrid Cloud Implementation Engineer certification



<https://killexams.com/pass4sure/exam-detail/NS0-403>

Question: 264

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

- 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: 265

What are three version control processes? (Choose three.)

- A. Subversion
- B. version
- C. Perforce
- D. revision
- E. Git

Answer: A,B,C,D

Question: 266

A customer was running a traditional application with the data being provided by an NetApp ONTAP NFSv3 export. The customer is now migrating that same application to run with a container in a Kubernetes environment using NetApp Trident.

In this scenario, how would the customer ensure that the new container-based application has access to the same data?

- A. Use the rsync command to copy the data to a worker node in the Kubernetes cluster and allow the kubelet to auto-discover the storage for the container.
- B. Import the volume using the tridentctl command and provide the corresponding persistent volume claim YAM
- C. Define a backend YAML in Trident and use the kubectl command to create the backend in the Trident namespace.
- D. Define the connection information for the NFSv3 export in the storage class object and then reference the storage class in the persistent volume claim.

Answer: C

Question: 267

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: A,B,D

Question: 269

Using Astra Control, which three steps must be completed to protect your Kubernetes-based applications? (Choose three.)

- A. Set the snapshot retention.
- B. Set the snapshot schedule.
- C. Copy the protection. YAML file to the/etc/ directory.
- D. Choose the application to protect.
- E. Install the appropriate application agent

Answer: A,D,E

Question: 270

Your customer is running their Kafka and Spark applications inside a Kubernetes cluster. They want to create a single backup policy within Astra to automate the backup of both applications and all their related Kubernetes objects at the same time.

Which method in Kubernetes should be used to accomplish this task?

- A. Create a Helm chart that deploys Kafka and Spark and their related objects to multiple namespaces.
- B. Put the applications and their objects in a single namespace, or label all the objects with a single label that Astra can recognize.
- C. Create a Kubernetes custom resource definition that includes all of the objects that Astra needs to treat as a single entity.
- D. Use a single Trident-based StorageClass to provision all storage for Kafka and Spark.

Answer: C

Question: 271

Your organization is adopting DevOps techniques to accelerate the release times of your product. The organization wants to have more communication between teams, quick iterations, and a better method for training work in progress.

In this scenario, which three methodologies would you recommend? (Choose three.)

- A. waterfall
- B. agile
- C. kanban
- D. scrum
- E. iterative

Answer: A,B,C,D

Question: 272

You are using REST APIs to deploy and manage storage that is provided by NetApp Cloud Volumes ONTAP in Microsoft Azure. You have written a Python script that contains a POST request to the following URL: `http://localhost/occm/api/azure/vsa/working-environments`. However, you receive a 401 response failure message that indicates that the user is not authenticated.

In this scenario, which two steps must you perform to communicate with an endpoint? (Choose two.)

- A. Call an API endpoint with the JWT access token.
- B. Acquire a JWT access token from the OAuth token endpoint.
- C. Call an API endpoint with the refresh token.
- D. Acquire a refresh token from the OAuth token endpoint.

Answer: C

Question: 273

You have an AFF A400 system with NetApp ONTAP 9.8 installed. You want to configure the NetApp Cloud Tiering service to tier data to AWS object storage. The communication between your cluster and the AWS cloud needs to be encrypted. You have installed the FabricPool license, but your ONTAP cluster cannot reach the S3 bucket.

Which two configurations are causing the error? (Choose two.)

- A. You have not configured the TLS encryption.
- B. You have not configured your firewall to allow outbound traffic on port 443.
- C. You have not configured a LIF with a public IP address.
- D. You have not configured a LIF with the intercluster role.

Answer: A,C,D

Question: 274

What are two stages in the DevOps lifecycle? (Choose two.)

- A. plan
- B. architect
- C. operate
- D. destroy

Answer: A,C,D

Question: 275

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.)

- A. Limit the maximum volume count.
- B. Limit the available data LIFs.
- C. Use a shared SV
- D. Limit the maximum volume size.

Answer: A

Question: 276

You are researching different automation frameworks. One of the main features that you are seeking is the ability to provide an end state of what a particular environment should look like versus needing to define each step that is required to get to the end state.

In this scenario, which keyword should you be looking for as you research various automation frameworks?

- A. RESTful
- B. declarative
- C. interrogative
- D. imperative

Answer: D

Question: 277

You are working with your development team to automate the deployment of Cloud Volumes ONTAP instances across various hyperscalers. The developers are asking you for specific RESTful API calls and other details to facilitate the automation.

In this scenario, which two links would you provide to them? (Choose two.)

- A. the link to the API documentation on cloud.netapp.com
- B. the link to the Swagger interface found on the connector instance
- C. the link to the Swagger interface found on cloud.netapp.com
- D. the link to the API documentation on readthedocs.io

Answer: A,B,C

Question: 278

Click the Exhibit button.

Tier to Local Bucket

SELECTED LOCAL TIER

ssd_aggr

PRIMARY TIER

☐ Existing

☒ New

A new storage VM and bucket will be added. The system will try to select low-cost media with optimal performance for the tiered data.

BUCKET CAPACITY

2 PB ✓

☐ Edit volume tiering policy

Save Cancel

You are configuring a FabricPool aggregate as shown in the exhibit.

In this scenario, where is the S3 bucket created?

- A. AWS
- B. StorageGRID
- C. NetApp ONTAP
- D. Azure

Answer: C

Question: 279

You have a StorageGRID solution with 1 PB of object data. All data is geographically distributed and erasure coded across three sites. You are asked to create a new information lifecycle management (ILM) policy that will keep a full copy of the grid in Amazon S3.

In this scenario, which component must be configured for the ILM policy?

- A. archive node
- B. additional 1 PB license
- C. Cloud Storage Pool
- D. Cassandra database

Answer: A

Question: 280

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: "/{{ item }}"
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

Question: 281

Your customer is using Astra Control Service to back up applications. They require a consistent copy of the data that also captures data in flight.

Which two applications will Astra support in this scenario? (Choose two.)

- A. ksqIDB
- B. Jenkins
- C. Argo CD
- D. PostgreSQL

Answer: A,B

Question: 282

You have a hybrid storage environment with on-premises ONTAP 9.8 clusters and Cloud Volumes ONTAP for AWS 9-8 clusters. You have a requirement to automatically discover, map, and classify your data.

In this scenario, which NetApp service satisfies this requirement?

- A. Cloud Data Sense
- B. Cloud Secure
- C. Cloud Volume Service
- D. Cloud Sync

Answer: C

Question: 283

A team compiles a group of artifacts locally. The team decides to allow an automatic process to deploy the compiled artifacts to a production server.

In this scenario, which step was skipped in the DevOps process?

- A. continuous builds
- B. continuous delivery
- C. continuous integration
- D. continuous deployment

Answer: C

Question: 284

You have a 4-node AFF A800 cluster that is serving file data. You are required to replicate copies of the data volumes directly to Amazon S3.

In this scenario, which two NetApp tools satisfy this requirement? (Choose two.)

- A. Cloud Sync
- B. XCP
- C. FabricPool
- D. SnapMirror Cloud

Answer: A,C,D

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.