



Exam Code: NS0-403

Exam Name: NetApp Certified Hybrid Cloud Implementation Engineer

Website: <https://VCEup.com/>

Team-Support: <https://VCEplus.io/>



Question No: 1

You want to manage your applications using NetApp Astra Control within your Azure account. Which two actions must be completed before accomplishing this task? (Choose two.)

- A. Manually provision Azure Blob storage
- B. Delegate a subnet to Azure NetApp Files.
- C. Set up a capacity pool.
- D. Export your Azure AD Administrator key and import it into Astra Control.

Answer: BD

Explanation:

Question No: 2

Your company uses all four of NetApp's storage offerings: ONTAP, StorageGRID, Element, and ESeries.

You are asked to provide persistent storage to the OpenShift group by using NFS.

In this scenario, which of the available storage deployments would provide access?

- A. StorageGRID
- B. ONTAP Select
- C. E-Series
- D. Element

Answer: C

Explanation:

Question No: 3

You want to add your NetApp StorageGRID embedded Prometheus instance as a data source to Grafana. What must be done to accomplish this task?

- A. Unblock port 9090 on the StorageGRID primary admin node.
- B. Deploy a gateway node in StorageGRID.
- C. Reboot the StorageGRID primary admin
- D. Create a dashboard for StorageGRID in Grafana.

Answer: C

Explanation:

Question No: 4

You have a single node Cloud Volumes ONTAP (CVO) system in AWS.

In this scenario, which two activities will cause system downtime? (Choose two.)

- A. Change the disk type of a volume from HDD (st1) to SSD (gp2).
- B. Change the Cloud Volumes ONTAP EC2 instance type.

C. Upgrade the Cloud Volumes ONTAP system software.

D. Apply a new Cloud Volumes ONTAP license.

Answer: CD

Explanation:

Question No: 5

As a DevOps engineer, you want a single tool that uses one automation language consistently across orchestration, application deployment, and configuration management.

In this scenario, which tool would you choose?

A. Docker

B. Ansible

C. Selenium

D. Octopus

Answer: D

Explanation:

Question No: 6

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

A. Subversion

B. version

C. Perforce

D. revision

E. Git

Answer: BCD

Explanation:

Question No: 7

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

A. plan

B. architect

C. operate

D. destroy

Answer: CD

Explanation:

Question No: 8

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: CD

Explanation:

Question No: 9

A customer made a mistake and deleted some important notes in their Jupyter notebooks. The customer wants to perform a near-instantaneous restore of a specific snapshot for a JupyterLab workspace.

Which command in the NetApp DataOps Toolkit will accomplish this task?

- A. `./ntap_dsutil.py restore snapshot`
- B. `./ntap__dsutil__k8s.py clone volume`
- C. `./ntap_ds\it±l_k8s.py list jupyterlab-snapshot`
- D. `./ntap_clsut±l_k8s .py restore jupyteclab-snapshot`

Answer: C

Explanation:

Question No: 10

You used a Terraform configuration to create a number of resources in Google Cloud for testing your applications. You have completed the tests and you no longer need the infrastructure. You want to delete all of the resources and save costs.

In this scenario, which command would you use to satisfy the requirements?

- A. `terraform untaint`
- B. `terraform apply`
- C. `terraform force—unlock`
- D. `terraform destroy`

Answer: D

Explanation:

Question No: 11

Click the Exhibit button.

```
resource "aws_iam_instance_profile" "instance_profile" {
  name = "connector_instance_profile"
  role = aws_iam_role.role.name
}

resource "netapp-cloudmanager_connector_aws" "cl-occm-aws" {
  name = "TF-ConnectorAWS"
  region = "us-east-1"
  key_name = "AWSKey"
  company = "NetApp"
  instance_type = "t3.xlarge"
  subnet_id = aws_subnet.PrivateSubnet.id
  security_group_id = aws_security_group.SGPrivateRedHat.id
  iam_instance_profile_name = "connector_instance_profile"
  account_id = "account-xxxxxxxxxx"
}
```

Referring to the exhibit, which two source resources? (Choose two.)

- A. source = "hashicorp/terraform"
- B. source = "hashicorp/aws"
- C. source = "terraform/terraform"
- D. source = "NetApp/netapp-cloudman"

Answer: AC

Explanation:

Question No: 12

Your organization is adopting DevOps principles to deliver innovation rapidly. Which three methods support this approach? (Choose three.)

- A. Automate processes as much as possible.
- B. Plan to reduce the blast radius of failures.
- C. Implement version control for code changes.
- D. Implement daily stand-up meetings for all team members.
- E. Allow more people within the organization to approve changes.

Answer: ABE

Explanation:

Question No: 13

Your customer is writing a curl command to test an API call to deploy Cloud Volumes OWTAP for AWS. They want to know what the content-Type parameter should be set to in the header of the HTTP request.

In this scenario, what would you tell the customer?

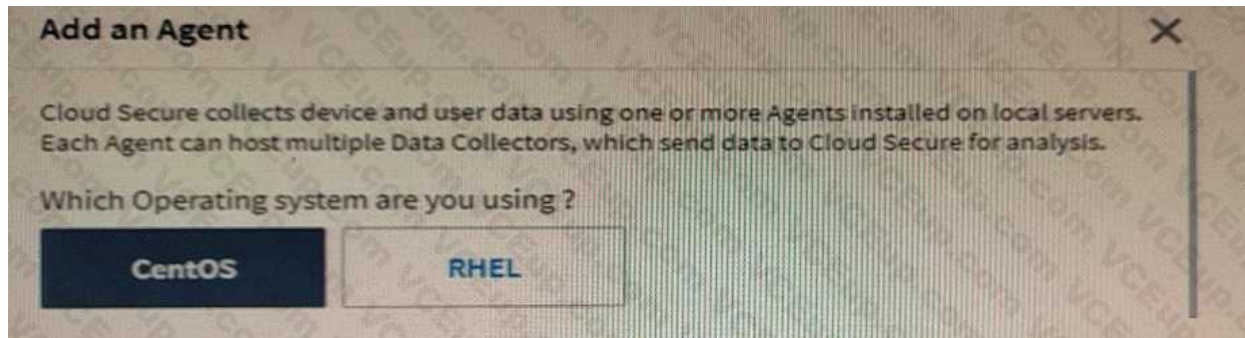
- A. Set the parameter to text/plain.
- B. Set the parameter to application/json
- C. Set the parameter to application/xml
- D. Set the parameter to text/html.

Answer: C

Explanation:

Question No: 14

As shown in the exhibit, you are installing NetApp Cloud Secure agents.



In this scenario, which component must be deployed for Cloud Secure to report findings?

- A. connector from Cloud Manager
- B. data broker
- C. container in Astra
- D. acquisition unit from Cloud Insight

Answer: C

Explanation:

Question No: 15

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

Explanation:

Question No: 16

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

Explanation:

Question No: 17

You need to find the documentation for an Ansible module without access to the Internet. In this scenario, which command will show you the documentation for na_ontap_volume?

- A. ansible—doc netapp.ontap.na_ontap_volume
- B. ansible—doc —lookup netapp.ontap.na_ontap_volume
- C. ansible—help netapp.ontap.na_ontap_volume
- D. ansible —doc netapp.ontap.na_ontap_volume

Answer: D

Explanation:

Question No: 18

You are asked to configure NetApp Trident to provide dynamic and persistent storage using Azure NetApp Files to your Azure Kubernetes Service (AKS) cluster.

In this scenario, which three elements do you need to configure your Trident back end? (Choose three.)

- A. an App Registration or service principal
- B. an ONTAP cluster credential
- C. an Azure portal credential
- D. an Azure subscription ID
- E. an Azure location/region with a subnet delegated to Microsoft .Netapp/volumes

Answer: ACD

Explanation:

Question No: 19

During software development, a team creates use cases, class diagrams, Unified Modeling Language models, software architecture documents, requirements documents, and design documents.

Collectively, in this scenario, what are these tangible by-products called?

- A. application programming interface
- B. software development kits
- C. build tools
- D. artifacts

Answer: B

Explanation:

Question No: 20

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: ADE

Explanation:

Question No: 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 No: 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 No: 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-ofobservability/#:~:text=Observability%20is%20divided%20into%20three,called%20three%20pillars%20of%20observability.>

Question No: 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 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 No: 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. Kubemetes StorageClass
- D. Trident volume

Answer: BD

Explanation:

Question No: 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: "/{{ 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

Explanation:

Question No: 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 No: 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 No: 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 No: 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. AWS API
- C. Tenancy API
- D. ONTAP API

Answer: B

Explanation:

Question No: 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 No: 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.)

- 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 No: 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 No: 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 No: 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 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.

Answer: AC

Explanation:

Question No: 36

You are responsible for setting up a process to handle version control for developers' source code.

You need to decide which tools you will use.

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

- A. Docker Hub
- B. git
- C. Jfrog
- D. subversion

Answer: AD

Explanation:

Question No: 37

You recently started at a new company where the persistent storage for Kubernetes was not provided by NetApp. You want to make sure you take full advantage of the additional annotations available from your NetApp storage.

In this scenario, which three parameters would you use to accomplish this task? (Choose three.)

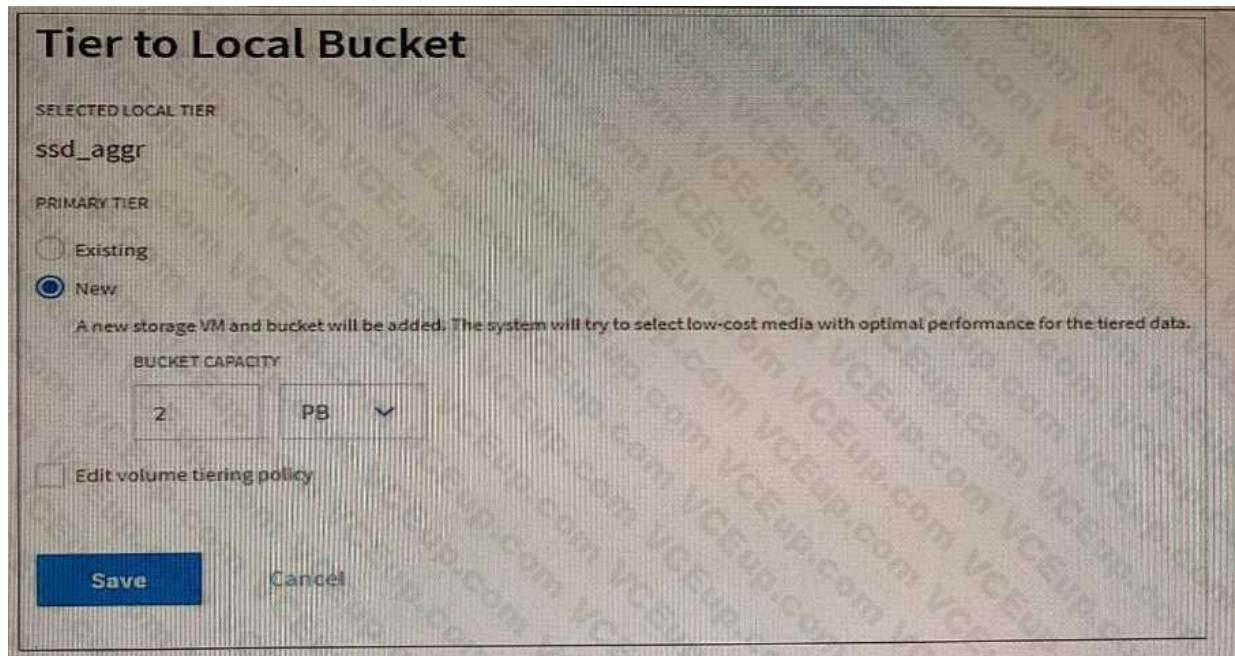
- A. accessModes,
- B. snapshotPolicy
- C. unixPermissions
- D. cloneFromPVC
- E. mountOptions

Answer: ABE

Explanation:

Question No: 38

Click the Exhibit button.



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

Explanation:

Question No: 39

Your developers each need their own unique copy of source code. Currently, you have 300 developers, and the source code that they are working on is 200 GBs in size.

Which NetApp ONTAP feature would help to solve this issue with the least amount of overhead?

- A. FlexCache
- B. SnapMirror
- C. FlexGroups
- D. FlexClone

Answer: C

Explanation:

Question No: 40

You are asked to accelerate your company's release cycle by adopting DevOps practices. As part of this project, you want to use infrastructure as code to deploy your storage environments during the CI/CD pipeline.

In this scenario, which tool would you use to satisfy this requirement?

- A. Docker
- B. Ansible
- C. Selenium
- D. Maven

Answer: B

Explanation:

Question No: 41

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 YAML.
- 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: D

Explanation:

Question No: 42

You are developing microservices on a Kubernetes system. When developers push new code to a branch, you want to have your application deployed on Kubernetes using the CI/CD pipeline that automatically picks up new changes in the main branch-In this scenario, what are three steps to accomplish this deployment workflow? (Choose three.)

- A. Issue a pull request.
- B. Reset the Git branches.
- C. Commit changes to a Git branch

D. Configure Git hooks.

E. Merge code changes to the main branch.

Answer: ACD

Explanation:

Question No: 43

Click the Exhibit button.

```
{
  "version": 1,
  "storageDriverName": "solidfire-san",
  "Endpoint": "https://<user>:<password>@<mvip>/json-rpc/8.0",
  "SVIP": "<svip>:3260",
  "TenantName": "<tenant>",
  "UseCHAP": true,
  "Types": [
    {"Type": "Bronze", "Qos": {"minIOPS": 1000, "maxIOPS": 2000, "burstIOPS": 4000}},
    {"Type": "Silver", "Qos": {"minIOPS": 4000, "maxIOPS": 6000, "burstIOPS": 8000}},
    {"Type": "Gold", "Qos": {"minIOPS": 6000, "maxIOPS": 8000, "burstIOPS": 10000}},
  ],
  "defaults": {
    "type": "Silver"
  }
}
```

Refer to the exhibit and the Storage Class information shown below.

```
apiVersion:
storage.k8s.io/v1beta1 kind:
StorageClass
metadata:
  name: storage-class1
provisioner: cli.trident.netapp.io
parameters:
  backendType: "solidfire-san"
  IOPS: 50
```

What are the minimum IOPS, maximum IOPS, and burst IOPS assigned to the persistent volumes that were created by Trident?

A. "minIOPS": 6000, "maxIOPS": 8000, "burstIOPS": 10000

B. "minIOPS": 4000, "maxIOPS": 6000, "burstIOPS": 8000

C. "minIOPS": 50, "maxIOPS": 50, "burstIOPS": 50

D. "minIOPS": 1000, "maxIOPS": 2000, "burstIOPS": 4000

Answer: D

Explanation:

Question No: 44

You configured your NetApp Cloud Volumes ONTAP instance to protect your data against ransomware. A previously unknown ransomware file extension infects your data.

In this scenario, which two NetApp technologies will provide a writable volume of the dataset before the ransomware attack? (Choose two.)

A. FPolicy

B. SnapRestore

C. FlexCache

D. FlexClone

Answer: BC

Explanation:

Question No: 45

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

Explanation:

Question No: 46

You deployed a NetApp Cloud Volumes ONTAP system in your AWS environment and you need to activate the WORM feature to secure them.

In this scenario, which two statements are correct? (Choose two.)

A. You can activate WORM storage on individual volumes; you do not need to activate WORM at the system level.

B. When WORM storage is activated, data tiering to object storage cannot be enabled.

C. When WORM storage is activated, data tiering to object storage can be enabled.

D. You cannot activate WORM storage on individual volumes; WORM must be activated at the system level.

Answer: AC

Explanation:

Question No: 47

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

Explanation:

Question No: 48

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

Explanation:

Question No: 49

You have received an internal request for persistent storage where the developer would like to be able to specify the filesystem.

Which three NetApp Trident backend storage drivers would satisfy the requirement? (Choose three.)

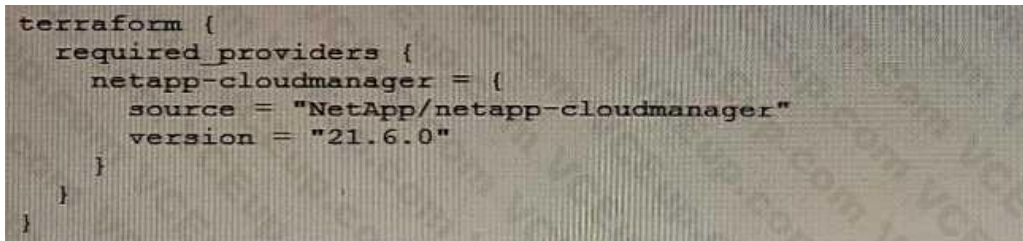
- A. ontap-nas
- B. ontap-nas-economy
- C. solidfire-san
- D. eseries-iscsi
- E. ontap-.san

Answer: BCD

Explanation:

Question No: 50

Click the Exhibit button.



```

terraform {
  required_providers {
    netapp-cloudmanager = {
      source = "NetApp/netapp-cloudmanager"
      version = "21.6.0"
    }
  }
}

```

Referring to the exhibit, what would happen if there is a newer version of netapp-cloudmanager available when you install the provider with the terraform init command?

- A. The newer version will be downloaded.
- B. The newer version will be downloaded only when it is a patch release.

C. Only the exact version 21.6.0 is downloaded.

D. Any version below version 21.6.0 is installed.

Answer: D

Explanation:

Question No: 51

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: CD

Explanation:

Question No: 52

Click the Exhibit button.



You are using NetApp Cloud Manager to set up replication, as shown in the exhibit.

Which two replication policies are supported in this scenario? (Choose two.)

A. Mirror

B. Tier

C. Clone

D. Backup

Answer: AD

Explanation:

Question No: 53

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: BC

Explanation:

Question No: 54

A team is implementing automated builds of Dev/Test environments. The team must ensure consistent and repeatable environments.

In this scenario, which two actions will satisfy this requirement? (Choose two.)

- A. Ensure that build scripts log to the same output file.
- B. Use variables in build scripts to identify environments.
- C. Use infrastructure as code to build the environments.
- D. Ensure that a service desk ticket is automatically logged when an environment is built.

Answer: AC

Explanation:

Question No: 55

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. ksqldb
- B. Jenkins
- C. Argo CD
- D. PostgreSQL

Answer: AB

Explanation:

Question No: 56

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

Explanation:

Question No: 57

Click the Exhibit button.

```
[
  "version": 1,
  "storageDriverName": "ontap-nas",
  "managementLIF": "10.0.0.1",
  "dataLIF": "10.0.0.2",
  "svm": "svm_nfs",
  "username": "vsadmin",
  "password": "secret",
  "defaults": {
    "spaceReserve": "volume",
    "exportPolicy": "myk8scluster"
  }
]
```

Referring to the exhibit, you are certain that the backend configuration information is correct, but you still cannot get the created PVCs to connect.

What are three reasons for this problem? (Choose three.)

- A. The Kubernetes nodes are configured to allow PVC over NFS.
- B. The NFS utilities are not installed on the Kubernetes nodes.
- C. The firewall settings between the Kubernetes nodes and the NetApp ONTAP cluster are not set up properly.
- D. The Kubernetes nodes are installed with NetApp drivers.
- E. The export policy rules allowing the Kubernetes nodes to connect are not set up properly.

Answer: ADE

Explanation:

Question No: 58

Which two definitions are true about data protection under Astra Control? (Choose two.)

- A. A backup is stored in an object store in the cloud.
- B. A snapshot is a point-in-time copy of an application that is stored on the same provisioned volume as the application.
- C. A snapshot is stored in an object store in the cloud.
- D. A backup is a point-in-time copy of an application that is stored on the same provisioned volume as the application.

Answer: A

Explanation:

Question No: 59

You have a hybrid storage environment with on-premises ONTAP 9.8 clusters and Cloud Volumes ONTAP for AWS 9.8 clusters. You require anomaly detection to automatically analyze data access patterns to identify risks from ransomware attacks.

In this scenario, which NetApp service satisfies this requirement?

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

Answer: A

Explanation:

Question No: 60

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: BCD

Explanation: