




Latest Version: 7.0

Question: 1

You are using NetApp Cloud Sync to replicate NFS to S3 as shown in the exhibit.

Schedule [ON] | No Tags

	Source nfs://217.70.21.18	✓ Sync Completed Initial Copy A few seconds ago Duration: < 1 min	
	Target s3://medium		Copied 692 Files 5.08 TB
	On-Prem Data Bro... on-prem1		Deleted 0 Files
		Failed 4 Files 241.28 MB	

Which network connection is required for the replication service?

Response:

- A. The data broker needs an inbound Internet connection.
- B. The NFS server needs a VPN connection to the S3 endpoint.
- C. The NFS server must be on the same subnet as the S3 endpoint.
- D. The data broker needs an outbound Internet connection.

Answer: D

Question: 2

A customer has deployed a NetApp HCI environment and has configured SnapMirror replication relationships to Cloud Volumes ONTAP in AWS. After migrating Cloud Manager from on-premises to AWS, the customer can no longer create new replication relationships between existing environments. The Cloud Manager can access the Cloud Volumes ONTAP working environment and the current SnapMirror relationships continue to make successful transfers of data.

In this scenario, where would you locate the source of the problem?

Response:

- A. The routing table for the subnet where Cloud Volumes ONTAP resides.
- B. The EC2 type of the Cloud Manager EC2 instance.
- C. The Security Group for the Cloud Volumes ONTAP EC2 instance.
- D. The IAM role for Cloud Manager EC2 instance.

Answer: C

Question: 3

You are implementing a backup scenario from an on-premises system to a NetApp Cloud Volumes ONTAP system on AWS. In this scenario, which two implementations ensure that you will recover from a single Amazon Elastic Block Storage (EBS) disk failure?

(Choose two.)

Response:

- A. Use SnapMirror to replicate data to a single-node NetApp Cloud Volumes ONTAP system. Then, from this secondary system, use SnapMirror to replicate data to a tertiary Cloud Volumes ONTAP system.
- B. Use SnapMirror to replicate data to a single-node NetApp Cloud Volumes ONTAP system.
- C. Use SnapMirror to replicate data to a single-node NetApp Cloud Volumes ONTAP system and perform periodic EBS snapshots of the Cloud Volumes ONTAP disks.
- D. Use SnapMirror to replicate data to an HA Cloud Volumes ONTAP system.

Answer: CD

Question: 4

You have FabricPool enabled for an on-premises AFF system. FabricPool is connected to Azure Blob storage. In this scenario, which configuration is required for high availability of data access in the hybrid aggregate?

Response:

- A. Duplicate the Azure Managed disk to another Availability Zone.
- B. Mirror the Blob storage container to another Azure region.
- C. Create multiple gateway subnets.
- D. Create redundant VPN or Express Route connections.

Answer: B

Question: 5

Which three volume workflows are supported by Trident for Kubernetes?

(Choose three.)

Response:

- A. volume import
- B. volume replication

- C. LUN resize
- D. volume clone
- E. volume snapshots

Answer: BDE

Question: 6

You are instructed to move towards an infrastructure as code process. You need to be able to track changes to the files that you are using. In this scenario, how would you satisfy the requirement?

Response:

- A. Enable WORM configurations for each file.
- B. Keep the files in a source-content versioning system.
- C. Enable encryption for each file.
- D. Keep the files in a block-storage device.

Answer: B

Question: 7

You are developing an Ansible playbook to create a volume within a NetApp ONTAP cluster. In this scenario, which Ansible module would be used to collect a list of existing aggregates?

Response:

- A. na_ontap_cluster
- B. netapp_e_facts
- C. na_ontap_gather_facts
- D. na_elementsw_cluster_config

Answer: C

Question: 8

Your software development team creates and deploys applications to Kubernetes. You want to provide the team with metrics from the applications resource usage. In this scenario, which two products would satisfy this requirement?

(Choose two.)

Response:

- A. Application Lifecycle Manager
- B. OnCommand System Manager

- C. Cloud Insights
- D. OnCommand Unified Manager

Answer: CD

Question: 9

You have deployed your application to an NKS cluster in your on-premises environment with ONTAP Select. You now want to move your application to an AWS environment. What are two ways to accomplish this task?

(Choose two.)

Response:

- A. Use CloudMirror.
- B. Use Cloud Sync.
- C. Use SnapSync.
- D. Use SnapMirror.

Answer: BD

Question: 10

You are setting up a containerized environment with persistent storage. Trident has been installed, and the Trident back end for NetApp SolidFire has been created by using the tridentctl command. Now, you need to deploy an application that requires an Extreme class of service.

```
{
  "version": 1,
  "storageDriverName": "solidfire-san",
  "Endpoint": "https://username:password@10.1.1.1/json-rpc/7.0",
  "SVIP": "10.1.2.50:3260",
  "BackendName": "sf",
  "TenantName": "tenant1",
  "InitiatorIFace": "default",
  "DefaultVolSz": 1,
  "UseCHAP": true,
  "Types": [{"Type": "Standard", "Qos": {"minIOPS": 1000, "maxIOPS": 2000, "burstIOPS": 4000}},
            {"Type": "Premium", "Qos": {"minIOPS": 4001, "maxIOPS": 6000, "burstIOPS": 8000}},
            {"Type": "Extreme", "Qos": {"minIOPS": 8001, "maxIOPS": 9000, "burstIOPS": 10000}}
]
```

Referring to the exhibit, which two configurations would accomplish this task?

(Choose two.)

Response:

- A.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: storage-class-hci-extreme
provisioner: csi.trident.netapp.io
parameters:
  snapshots: "true"
  BackendType: "solidfire-san"
  IOPS: "Extreme:8050"
```

B.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: storage-class-hci-extreme
provisioner: csi.trident.netapp.io
parameters:
  snapshots: "true"
  storagePools: "tenant1:Extreme"
```

C.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: storage-class-hci-extreme-tenant1
provisioner: csi.trident.netapp.io
parameters:
  snapshots: "true"
  BackendType: "solidfire-san"
  IOPS: "8050"
```

D.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: storage-class-hci-extreme
provisioner: csi.trident.netapp.io
parameters:
  snapshots: "true"
  storagePools: "sf:Extreme"
```

Answer: BD

Question: 11

You deployed a NetApp HCI private cloud system using a six-cable configuration. You are asked to make changes to the network configuration post-deployment. In this scenario, which configuration change is supported?

Response:

- A. Remove virtual interfaces.
- B. Change vSwitch names.
- C. Add virtual interfaces.
- D. Remove default port groups.

Answer: B

Question: 12

You are creating a new meme generation application and are using containers for service chaining. The end user uploads the original image to one container, the image resolution modifier to another container, and the text overlay to a third container.

In this scenario, knowing that you have NetApp ONTAP storage available to you, which solution will allow each container to access the same data?

Response:

- A. Re-encode the binary data into MIME64, and pass it as a string when instantiating new containers.

- B. Presenting the same LUN to all of your nodes' underlying OS will make that storage available to the container as well.
- C. Use the same container image for all three tasks and the data will exist because data in containers is persistent.
- D. Provision a persistent volume using Trident, and then have Kubernetes claim this volume, as necessary, to provide the same storage to each container.

Answer: B