Free Questions for NS0-304

Shared by Torres on 16-04-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An administrator wants to automate the configuration of SnapMirror policies between cloud and on-premises deployments in AWS using Ansible. What must the administrator do first?

Options:

A- Set up AWS Control Tower for automation
B- Subscribe to Ansible Automation Platform
C- Install the ONTAP collection using Ansible Galaxy
D- Install the Ansible plugin for aws_ec2 inventory

Answer:

C

Explanation:

To automate the configuration of SnapMirror policies between cloud and on-premises deployments in AWS using Ansible, the administrator needs to begin by installing the NetApp ONTAP collection from Ansible Galaxy. This collection contains modules specifically designed to manage NetApp ONTAP storage systems, including the management of SnapMirror configurations. Here are the steps to do this:

Installation of ONTAP Collection: Open your command line interface and run the command ansible-galaxy collection install netapp.ontap. This command pulls the ONTAP collection from Ansible Galaxy, which includes all necessary modules for managing NetApp ONTAP, including SnapMirror.

Configuration of Ansible Environment: Ensure that your Ansible environment is set up to connect to both your AWS environment and the on-premises NetApp ONTAP systems. This typically involves configuring the appropriate credentials and network settings in your Ansible playbooks and inventory files.

Writing Ansible Playbooks: With the ONTAP collection installed, you can now write Ansible playbooks that utilize the SnapMirror modules to automate the configuration of SnapMirror policies as required.

For further information on using the NetApp ONTAP Ansible collection, please refer to the official documentation available at: NetApp ONTAP Ansible Collection Documentation.

Question 2

Question Type: MultipleChoice

An administrator configures the trident ontap-san driver and specifies useCHAP=true.

```
{
    "version": 1,
    "storageDriverName": "ontap-san",
    "backendName": "ontap_san_chap",
    "managementLIF": "192.168.0.135",
    "svm": "ontap_iscsi_svm",
    "useCHAP": true,
    "username": "vsadmin",
    "password": "FaKePaSsWoRd",
    "igroupName": "trident",
    "chapInitiatorSecret": "cl9qxIm36DKyawxy",
    "chapTargetInitiatorSecret": "rqxigXgkesIpwxyz",
    "chapTargetUsername": "iJF4heBRT0TCwxyz",
}
```

Which parameter is missing from the configuration?

Options:

A- dataUF
B- chapSecret
C- clientPrivateKey
D- chapUsername

Answer:

D

Explanation:

When configuring the Trident ONTAP-SAN driver with CHAP (Challenge Handshake Authentication Protocol) enabled (useCHAP: true), it is essential to specify both the initiator and target CHAP secrets and their corresponding usernames. In the configuration provided, while the CHAP secrets for both initiator and target are specified (chapInitiatorSecret and chapTargetInitiatorSecret), only the CHAP target username (chapTargetUsername) is listed. The missing parameter necessary for the complete CHAP configuration is the CHAP initiator username (chapUsername). This username is used along with the chapInitiatorSecret to authenticate the initiator to the storage system.

To correct this, add the chapUsername field to the configuration file, ensuring that the initiator's username matches the configured environment and that it is correctly paired with the chapInitiatorSecret. This inclusion ensures that both sides of the CHAP authentication process are

properly identified, thereby providing the necessary security for SAN communication.

For further guidance on configuring CHAP with the Trident ONTAP-SAN driver, refer to the NetApp Trident documentation: NetApp Trident Documentation.

Question 3

Question Type: MultipleChoice

An administrator is setting up NetApp Cloud Tiering. They are creating a new S3 object storage bucket that needs to be compliant with the default IAM policy for the cloud connector.

How must the bucket be configured to meet the policy?

Options:

A- It must have cross-region replication enabled.
B- It must be configured to support NetApp Cloud Sync.
C- The prefix must be set to cloud-tier.
D- The prefix must be set to fabric-pool.

Answer:

D

Explanation:

When setting up NetApp Cloud Tiering with an S3 object storage bucket, it is crucial that the bucket configuration adheres to the default IAM policy for the cloud connector. Here's the configuration requirement:

Bucket Configuration with Specific Prefix: The IAM policy often specifies access permissions based on resource names or prefixes. For Cloud Tiering, particularly when integrating with FabricPool technology, the bucket should have a prefix set to fabric-pool. This allows the Cloud Tiering service to correctly identify and interact with the bucket, ensuring compliance with security policies and access controls.

Verify IAM Policy Configuration: Ensure that the IAM policy for the cloud connector includes permissions for operations on the S3 bucket with the fabric-pool prefix. This typically includes permissions to put, get, list, and delete objects within the bucket.

For further information on configuring S3 buckets for NetApp Cloud Tiering and detailed IAM

policy settings, please consult the NetApp Cloud Tiering documentation available on the NetApp website: NetApp Cloud Tiering Documentation.

Question 4

Question Type: MultipleChoice

An administrator has iSCSI LUNs on an AWS FSxN instance. The administrator is unable to mount the LUNs from a Linux host in the same AWS region. The Linux host is in a different VPC than FSxN.

What must the administrator configure to resolve this issue?

Options:

A- BGP peering
B- SVM peering
C- Cluster peering
D- VPC peering

Answer:

D

Explanation:

If an administrator has iSCSI LUNs on an AWS FSxN instance and is unable to mount these LUNs from a Linux host in the same AWS region due to the host being in a different Virtual Private Cloud (VPC), the solution is to configure VPC peering. Here's the process:

VPC Peering Setup: VPC peering allows two VPCs to communicate with each other as if they are in the same network. This enables the Linux host to connect to the AWS FSxN instance across different VPCs.

Configuration Steps: To set up VPC peering, the administrator must create a peering connection between the two VPCs in the AWS Management Console, and then update the route tables in each VPC to allow traffic to and from each other.

Mounting iSCSI LUNs: Once VPC peering is configured, the network route will be established, allowing the Linux host to successfully mount the iSCSI LUNs located on the FSxN instance.

For guidance on setting up VPC peering in AWS, consult the AWS documentation: AWS VPC

Peering Guide.

Question 5

Question Type: MultipleChoice

When deploying CVO, which two network types are used in HA configurations? (Choose two.)

Options:

A- Intracluster
B- iWARP
C- Transit Gateway
D- Intercluster

Answer:

A, D

Explanation:

When deploying Cloud Volumes ONTAP (CVO) in a High Availability (HA) configuration, two critical network types are used:

Intracluster: This network type is used for communication within the same cluster, particularly between nodes within the same HA pair. It is crucial for the synchronization and coordination of operations that support the cluster's internal processes and data management.

Intercluster: This network type facilitates communication between different clusters or between nodes across different data centers or geographical locations. It is typically used for data replication and disaster recovery purposes, ensuring data continuity and availability across diverse environments.

Understanding and configuring these network types correctly is essential for maintaining high availability and ensuring robust disaster recovery in CVO deployments.

For more details on network configuration in CVO HA setups, refer to the NetApp documentation on network management for Cloud Volumes ONTAP: NetApp CVO Network Documentation.

Question 6

Question Type: MultipleChoice

A customer requires unlimited backups be included for their CVO instance. Which two subscription models should the customer use? (Choose two.)

Options:

A- Professional
B- Essentials
C- Premium
D- Optimized
E- Edge Cache

Answer:

B, C

Explanation:

For a customer requiring unlimited backups in their Cloud Volumes ONTAP (CVO) instance, the Essentials and Premium subscription models are the appropriate choices. Both these subscription models offer unlimited backups as part of their service package, which is ideal for customers who prioritize extensive backup capabilities without the concern of hitting limits.

The Professional, Optimized, and Edge Cache plans typically have different focuses or limitations concerning backup capabilities:

Professional: Geared more towards smaller or less critical deployments without the breadth of features found in Premium or Essentials.

Optimized: Often focuses on performance optimization rather than extensive backup functionalities.

Edge Cache: Is used for caching services at the edge rather than core data management and backup functionalities.

Detailed information on these subscription models and their backup capabilities can be found in the NetApp Cloud Volumes ONTAP documentation or through consultation with NetApp sales representatives.