Free Questions for NS0-304 by vceexamstest

Shared by Rogers on 22-07-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A local ONTAP system has FabricPool configured to a public cloud. Storage use has grown exponentially due to end-of-year activities. After a few weeks, tiering to the cloud stops.

What should the administrator do?

Options:

A- Configure tiering to a second cloud provider

B- Increase the tiering license

C- Increase the cooling period

D- Adjust aggregate fullness

Answer:

D

Explanation:

When tiering to the cloud stops in an ONTAP system configured with FabricPool, especially after a rapid increase in storage use, it is likely due to reaching the capacity threshold of the aggregate. ONTAP systems with FabricPool will halt tiering if the aggregate becomes too full. The solution is to adjust the aggregate fullness, either by increasing the aggregate's capacity or by managing the existing data more effectively (e.g., deleting unneeded data or moving data to another aggregate).

Configure tiering to a second cloud provider: This might help in distributing data, but it does not address the issue if the problem is local aggregate capacity.

Increase the tiering license: Generally, tiering licenses are about the amount of data that can be tiered rather than a technical limitation affecting tiering functionality.

Increase the cooling period: This might delay data movement but does not resolve the issue of aggregate fullness halting tiering.

Adjusting the aggregate fullness directly addresses the root cause by ensuring there's sufficient capacity within the local system to continue tiering operations. Information about managing aggregate capacity in ONTAP systems can be found in the ONTAP management documentation or the FabricPool administration guide.

Question 2

Question Type: **MultipleChoice**

An administrator needs to monitor their storage and compute resources in their hyperscaler and their private data center. Which tool meets this requirement?

Options:

A- SPOT by NetApp

B- Active IQ Digital Advisor

C- BlueXP Observability

D- BlueXP Classification

Answer:

C

Explanation:

To monitor storage and compute resources across both a hyperscaler and a private data center, BlueXP Observability is the appropriate tool. This part of the BlueXP suite offers a unified view of infrastructure health, performance, and capacity. Here's the benefit of using BlueXP Observability:

Unified Monitoring: BlueXP Observability provides a single pane of glass for monitoring resources, regardless of their location---whether in the cloud or on-premises. This includes real-time data on performance, capacity utilization, and system health.

Cross-Environment Support: It supports various environments, making it suitable for hybrid deployments. This capability allows administrators to have a holistic view of their entire infrastructure.

Alerts and Metrics: The tool offers customizable alerts and detailed metrics that help in proactive management and troubleshooting of storage and compute resources.

BlueXP Observability's extensive capabilities in monitoring and managing diverse IT environments make it an ideal choice for enterprises that operate across multiple platforms.

For more information on how to utilize BlueXP Observability for infrastructure monitoring, refer to the NetApp BlueXP documentation: NetApp BlueXP Documentation.

Question 3

Question Type: **MultipleChoice**

An administrator configures FSx for ONTAP to use as storage in their cloud environment. The administrator cannot access their NFS file system on clients located in another VPC.

What should the administrator configure?

Options:

A- VPC peering between the two VPCs

B- Routing using an AWS Transit Gateway between the two VPCs

C- An additional instance of FSx for ONTAP in same VPC as the client

D- A Direct Connect Gateway between the two VPCs

Answer:

A

Explanation:

To address the issue of not being able to access an NFS file system hosted on FSx for ONTAP in one Virtual Private Cloud (VPC) from clients located in another VPC, the administrator should configure VPC peering between the two VPCs. Here's why and how:

VPC Peering Setup: VPC peering allows two VPCs to communicate with each other as though they are part of the same network. This is essential for enabling direct access to the NFS file system across different VPCs.

Configure Network Routes: Once VPC peering is established, configure the network routes to ensure that traffic destined for the NFS file system can traverse the peered VPC connection.

Verify Accessibility: Test the NFS file system access from the client VPC to ensure that the configuration is correct and that the file system is accessible.

VPC peering is a straightforward solution that does not require the complexity and additional cost associated with options like Transit Gateways or Direct Connect. It's well-suited for enabling direct network connectivity between VPCs within the same cloud provider.

For more details on setting up VPC peering, refer to AWS documentation: AWS VPC Peering Guide.

Question 4

Question Type: MultipleChoice

An administrator must configure SVM-DR between two instances of Cloud Volumes ONTAP (CVO); one is deployed in Azure, and the other in AWS.

What must be configured to enable replication traffic between the two CVO instances?

Options:

A- Internet Gateway

B- Direct Connect

C- ExpressRoute

D- Virtual Private Network

Answer:

D

Explanation:

To enable replication traffic between two instances of Cloud Volumes ONTAP (CVO) deployed in Azure and AWS, a Virtual Private Network (VPN) must be configured. This setup is crucial because it provides a secure and private communication channel over the internet, which is necessary for the replication of data between different cloud providers. Here's the process:

Setup VPN Connection: Establish a VPN connection between the Azure and AWS environments. This involves configuring VPN gateways in both clouds to enable encrypted traffic flow between the two instances of CVO.

Configure Network Routing: Ensure that the routing rules are set to direct the replication traffic through the VPN connection. This might include setting up appropriate route tables that point to the VPN gateway.

Test and Verify Connectivity: After setting up the VPN, conduct tests to verify that the replication traffic is flowing correctly and securely between the two cloud environments.

Using a VPN is the most straightforward and typically the most cost-effective method to securely link AWS and Azure for the purpose of data replication, without the need for direct connectivity services like AWS Direct Connect or Azure ExpressRoute, which are more complex and costly solutions.

For guidance on setting up VPNs between AWS and Azure, refer to the respective cloud provider's documentation on VPN configuration.

Question 5

Question Type: MultipleChoice

An administrator is configuring Cloud Volumes ONTAP (CVO). The CVO instance does not have outbound network connectivity to send AutoSupport messages.

What will BlueXP automatically configure as the proxy server for AutoSupport?

Options:

A- Page blob

B- Mediator

C- Collector

D- Connector

Answer:

D

Explanation:

In a scenario where a Cloud Volumes ONTAP (CVO) instance lacks outbound network connectivity to send AutoSupport messages, BlueXP (formerly known as NetApp Cloud Manager) will automatically configure the Connector as the proxy server for AutoSupport. The Connector serves as a bridge between the customer's environment and NetApp cloud services, facilitating communication and data transfer, including AutoSupport messages, when direct connectivity is unavailable.

Page blob is a type of storage in Azure, not related to network functions.

Mediator and Collector are not standard terms used within NetApp for describing components involved in managing or proxying AutoSupport messages.

BlueXP's configuration to use the Connector as a proxy ensures that all monitoring and telemetry data crucial for the health and performance diagnostics of the CVO instance are relayed effectively, even in environments with restrictive outbound network policies. More details on this setup can be explored in the BlueXP or Cloud Volumes ONTAP documentation available on NetApp's website.

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.

Question 7

Question Type: MultipleChoice

An administrator wants to use BlueXP Observability to generate notifications whenever a volume in a FlexGroup on a GCP CVO system is nearing capacity.

Which option should the administrator use?

Options:

A- Data Collection API

B- Service Level Annotation

C- Log Monitor

D- AIQ ConfigAdvisor

Answer:

A

Explanation:

To generate notifications whenever a volume in a FlexGroup on a GCP CVO system is nearing capacity using BlueXP Observability, the Data Collection API should be utilized. This option allows for the configuration of customized monitoring and alerting based on specific data points and thresholds. Here's the process:

Setup Data Collection API: Configure the Data Collection API to monitor volume capacity metrics within your FlexGroup. This involves setting up the API to pull or receive data points related to storage utilization.

Define Alerts: Set thresholds for when capacity is considered nearing its limit (e.g., 80% full). Configure alerts to be triggered when these thresholds are approached, ensuring administrators are notified in advance to take necessary actions.

Implement Notification System: Integrate the alerting mechanism with your organization's notification system (e.g., email alerts, SMS, or a dashboard) to inform the relevant stakeholders or administrators promptly.

For detailed instructions on configuring the Data Collection API and setting up monitoring and alerting in BlueXP Observability, refer to the NetApp BlueXP documentation and API guides: NetApp BlueXP Documentation.

Question 8

Question Type: **MultipleChoice**

An administrator wants to add more disks to an aggregate to their existing Azure CVO instance. What is the supported method?

Options:

A- Use the BlueXP Advanced Allocation option

B- Assign more disks to the CVO VM in Azure Portal

C- Add more disks in ONTAP System Manager

D- Use the 'aggr add-disk' command on CLI

Answer:

C

Explanation:

To add more disks to an aggregate in an existing Azure CVO instance, the supported method is to use ONTAP System Manager. This tool provides a user-friendly graphical interface for managing ONTAP features, including disk and aggregate management. Here's how:

Access ONTAP System Manager: Log into ONTAP System Manager on your Azure CVO instance.

Manage Disks and Aggregates: Navigate to the Storage or Aggregates section, where you can view existing aggregates and the disks assigned to them.

Add Disks to Aggregate: Select the aggregate you wish to expand and follow the prompts to add additional disks. This process may involve selecting disks from a pool of unassigned disks or reconfiguring existing disk resources.

For more detailed guidance on managing aggregates and disks in Azure CVO, refer to the specific section in the ONTAP System Manager documentation on disk and aggregate management: NetApp ONTAP System Manager.

Question 9

Question Type: MultipleChoice

A company is setting up FlexCache in CVO to scale-out an on-premises system. What should the administrator do on the on-premises system?

Options:

A- Create a new volume as a cache

B- Generate cluster peering passphrase

C- Configure NFS export policies

D- Apply a FlexCache license

Answer:

B

Explanation:

When setting up FlexCache in Cloud Volumes ONTAP (CVO) to scale out an on-premises system, the critical first step on the on-premises system is to generate a cluster peering passphrase. This passphrase is used to establish a secure cluster peering relationship between the on-premises ONTAP system and the CVO in the cloud. Here's the process:

Cluster Peering Setup: Cluster peering is essential for FlexCache because it allows the on-premises system to communicate and share data with the CVO instance. The cluster peering passphrase is used to authenticate the peering session, ensuring security.

Generate the Passphrase: In the ONTAP system manager on the on-premises cluster, navigate to the cluster peering settings and generate or configure the passphrase that will be used for peering with the CVO.

Establish Peering: Once the passphrase is set, use it to create the cluster peer relationship from the on-premises ONTAP to the CVO, following the guided steps in ONTAP System Manager or using CLI commands.

For detailed instructions on setting up cluster peering for FlexCache, refer to the NetApp documentation on FlexCache and cluster peering: NetApp FlexCache Documentation.

Question 10

Question Type: MultipleChoice

An administrator wants to protect Kubernetes-based applications across both on-premises and the cloud. The backup must be application aware and protect all components and data for the application. The administrator wants to use SnapMirror for disaster recovery.

Which product should the administrator use?

Options:

A- NetApp SnapCenter

B- Astra Control Service

C- Cloud Backup Service

D- Astra Control Center

Answer:

B

Explanation:

Astra Control Service is the appropriate NetApp product for protecting Kubernetes-based applications across both on-premises and cloud environments. Astra Control Service is designed to provide application-aware data management, which means it understands the structure and dependencies of Kubernetes applications and can manage them holistically. This includes backup and recovery, application cloning, and dynamic scaling.

While SnapMirror could be used for disaster recovery by replicating data at the storage layer, it does not inherently understand or manage the Kubernetes application layer directly. SnapCenter is primarily focused on traditional data management for enterprise applications on NetApp storage and does not cater specifically to Kubernetes environments. Cloud Backup Service is for backup to the cloud and also does not provide the Kubernetes application awareness required in this scenario.

Thus, Astra Control Service, which integrates deeply with Kubernetes, allows administrators to manage, protect, and move containerized applications and their data across multiple environments, making it the best fit for the described requirements. For detailed information on Astra Control Service's capabilities with Kubernetes applications, refer to the official NetApp Astra Control Service documentation.

Question 11

Question Type: **MultipleChoice**

Refer to the exhibit.

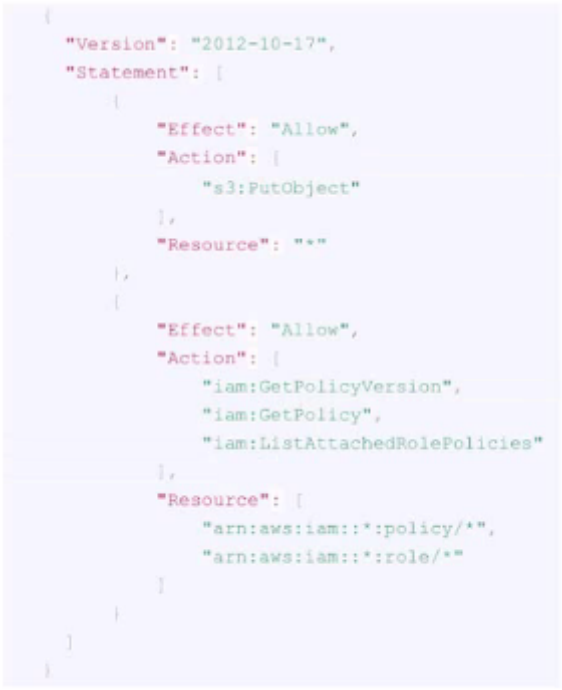

An administrator needs to review the IAM role being provisioned for Cloud Data Sense in order to scan S3 buckets. Which two permissions are missing? (Choose two.)

Options:

A- s3:DeleteObject

B- s3:Put0bjectAcl

C- s3:List*

D- s3:GetObjectAcl

E- s3:Get*

Answer:

C, E

Explanation:

For Cloud Data Sense to effectively scan S3 buckets, it requires permissions to list and get objects within the buckets. From the IAM policy provided in the exhibit, the permissions currently include s3:PutObject for object creation and a series of IAM-related permissions such as iam:GetPolicyVersion, iam:GetPolicy, and iam:ListAttachedRolePolicies. However, for scanning purposes, Data Sense needs to read and list the objects in the buckets. Therefore, the missing permissions are:

s3:List*: This permission allows the listing of all objects within the S3 buckets, which is necessary to scan and index the contents.

s3:Get*: This grants the ability to retrieve or read the content of the objects within the S3 buckets, which is essential for scanning the data within them.

These permissions ensure that Cloud Data Sense can access the metadata and contents of objects within S3 to perform its functionality.