PassQueen
a.k.a dumpsqueen.com

NetApp Certified Hybrid Cloud Administrator

Netapp NS0-302

Version Demo

Total Demo Questions: 11

Total Premium Questions: 119

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Hybrid Cloud Fundamentals	38
Topic 2, ONTAP Administration	10
Topic 3, Cloud Volumes ONTAP Administration	46
Topic 4, Cloud Manager Administration	25
Total	119

QUESTION NO: 1

You want to reduce Amazon EBS capacity costs for a Cloud Volumes ONTAP single-node system. The system is used only during business hours and has no requirement to serve data overnight.

Which action reduces compute charges while retaining the deployed working environment?

A. Schedule the Cloud Volumes ONTAP system to shut down outside business hours.

B. Disable AutoSupport on the working environment.

C. Set the volume tiering policy to none.

D. Delete the Connector after deployment.

ANSWER: A

Explanation:

Schedule the Cloud Volumes ONTAP system to shut down outside business hours is correct. Stopping the Cloud Volumes ONTAP instance reduces charges associated with the cloud compute resources while the system is not required. BlueXP can schedule automatic shutdown and restart operations for supported Cloud Volumes ONTAP working environments.

Stopping the system does not eliminate storage capacity charges for resources such as EBS disks, but it reduces instance-runtime costs. See the BlueXP documentation for managing Cloud Volumes ONTAP state.

QUESTION NO: 2

Click the Exhibit button.

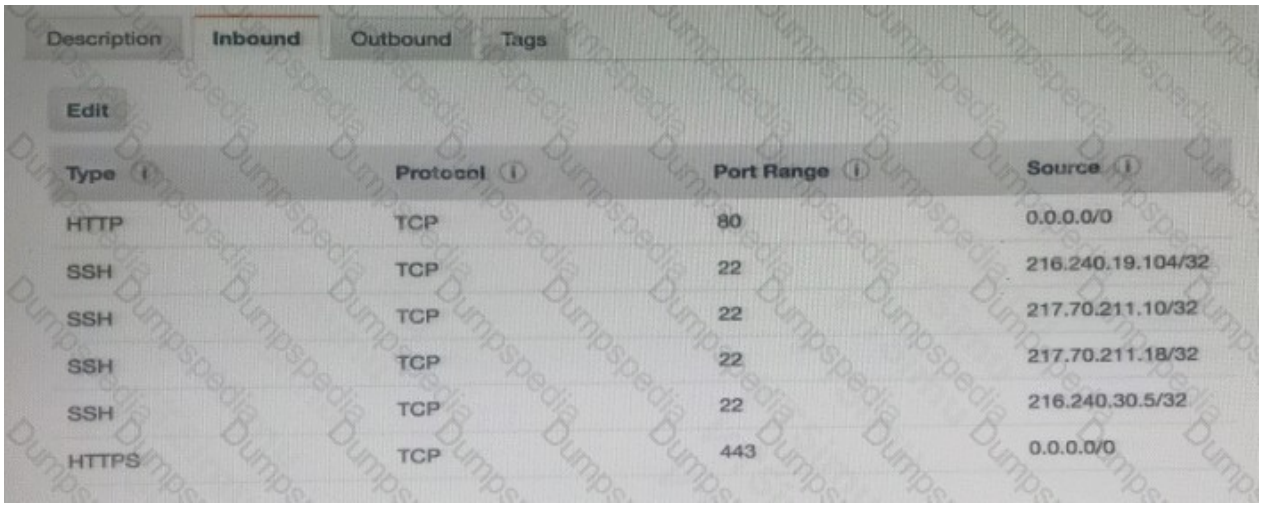

You have deployed Cloud Volumes ONTAP for AWS. During the deployment, you use an existing security group as shown in the exhibit. You are not able to connect NFS or SMB clients to the Cloud Volumes ONTAP Instance.

In this scenario, which two changes must be made to the security group to solve this problem? (Choose two.)

A. Change to a new security group with access controls according to the documentation.

B. Change the allowed ports to allow SMB and NFS outbound.

C. Change the allowed ports to allow SMB and NFS Inbound.

D. Change the source addresses to allow the 0.0.0.0/0 address.

ANSWER: A D

Explanation:

Change to a new security group with access controls according to the documentation. is correct because Cloud Volumes ONTAP requires a security group that includes the documented data-protocol access controls as well as the required management and service connectivity rules. When an existing security group does not provide the necessary Cloud Volumes ONTAP access model, assigning a security group built according to NetApp's documented AWS requirements provides the required rule set for client data access.

Change the source addresses to allow the 0.0.0.0/0 address. permits client traffic from any IPv4 source to reach the required NFS and SMB listener ports, resolving the connectivity failure described. In a real production deployment, NetApp and AWS security best practice is to scope the source to the actual client subnet CIDR ranges rather than leave unrestricted public access. However, allowing `0.0.0.0/0` does make the security-group source range inclusive of the affected clients and therefore restores reachability in this scenario.

NetApp documents the AWS security-group requirements, including the required inbound protocol access, in its Cloud Volumes ONTAP security-group reference. AWS also describes how security-group rules control permitted inbound traffic in the Amazon VPC security groups documentation.

QUESTION NO: 3

You want to provide read-heavy workloads at a remote site with low-latency access to data stored on a Cloud Volumes ONTAP volume. The remote site requires a local cache while the original volume remains the authoritative source.

In this scenario, which ONTAP technology should you use?

A. FlexCache

B. SnapVault

C. CloudMirror

D. FabricPool

ANSWER: A

Explanation:

FlexCache is the correct technology. A FlexCache volume is a scalable, persistent cache of an origin volume. It allows clients near the cache to read data locally while ONTAP retrieves cache misses from the origin volume.

The origin volume remains the authoritative data source, while the cache improves read performance and reduces WAN traffic for repeated reads. NetApp documents this architecture in the ONTAP FlexCache overview.

QUESTION NO: 4

You have a cloud-connected storage system that is located in a collocation facility that uses AWS Direct Connect. The AWS Direct Connect circuit is already configured and you want to verify that your compute nodes in AWS are able to connect to the volumes created on the storage. You want to verify connectivity using the NFS protocol. In this scenario, what are three solutions? (Choose three.)

A. Write a text file and save it to the NFS export.

B. Use OnCommand Systems Manager to test NFS connectivity

C. Connect a remote desktop session to the export.

D. Create a NFS export and use the AWS VM Instance to mount the export.

E. Launch an AWS VM instance running Linux in the VPC that is connected to the AWS Direct Connect network.

ANSWER: A D E

Explanation:

Launching an AWS VM instance running Linux in the VPC that is connected to the AWS Direct Connect network provides an appropriate NFS client in the same AWS network context as the compute resources whose access must be validated. The instance can resolve and reach the storage virtual machine's NFS data LIF across the Direct Connect path.

Creating an NFS export and using the AWS VM instance to mount the export performs the essential end-to-end validation: the Linux NFS client contacts the export through the network, negotiates NFS access, and mounts the ONTAP volume using the applicable export policy permissions. A successful mount confirms that the required network path, NFS service, export configuration, and client authorization are functioning together.

Writing a text file and saving it to the NFS export validates actual client data access after the mount, rather than merely confirming that a connection can be established. This confirms that the mounted export is usable for file creation and write operations from the AWS-side compute environment. NetApp's NFS client-access verification guidance includes mounting the export from a client and testing access to it. See NetApp ONTAP: Verify NFS client access and AWS Direct Connect User Guide.

QUESTION NO: 5

As the site reliability engineer, you have successfully deployed Kubernetes clusters on AWS and Google Cloud Platform using NetApp Kubernetes Service.

In this scenario, what would enable you to manage these clusters through a single management Interface?

A. Federation

B. Helm Charts

C. Istlo

D. NetApp Trident

ANSWER: A

Explanation:

Federation is the capability that enables centralized management of Kubernetes clusters distributed across cloud providers such as AWS and Google Cloud Platform. In the NetApp Kubernetes Service context, federation provides a unified management layer for working with multiple independently deployed clusters rather than requiring the site reliability engineer to access and administer each cluster separately. This supports a consistent operational model across the multicloud estate, including the ability to view and manage federated cluster resources through one management interface while the clusters remain in their respective cloud environments.

More generally, Kubernetes federation is designed for multicluster use cases in which workloads and policies need coordinated administration across clusters. The Kubernetes project describes federation as an approach for managing multiple Kubernetes clusters, including clusters located in different regions or cloud providers. This aligns directly with the stated requirement to manage the AWS and Google Cloud Platform clusters from one interface. See the Kubernetes federation documentation and NetApp's Kubernetes integration documentation for related multicluster and Kubernetes-management context.

QUESTION NO: 6

You have designed and written applications as containers and want to deploy them on a Kubernetes cluster. You want to use Cloud Volumes ONTAP to provide dynamic persistent storage for your Kubernetes cluster.

In this scenario, what will automate this task?

A. Trident

B. Terraform

C. Istio

D. kubectl

ANSWER: A

Explanation:

Trident is correct because it is NetApp's Kubernetes storage orchestrator, designed to automate the consumption of NetApp storage by containerized applications. When Trident is installed in a Kubernetes cluster and configured with a Cloud Volumes ONTAP storage backend, Kubernetes users can request persistent storage through PersistentVolumeClaims. Trident dynamically provisions the appropriate PersistentVolumes and creates the backing storage on Cloud Volumes ONTAP according to the selected Kubernetes StorageClass and its policy settings. This provides application teams with self-service, dynamically provisioned persistent storage while allowing storage administrators to retain control of storage configuration, performance, and data-protection behavior. Trident supports Cloud Volumes ONTAP as a backend and integrates with standard Kubernetes CSI-based storage workflows, so applications use Kubernetes-native objects rather than manually creating ONTAP volumes for every workload. NetApp documentation describes Trident as a fully supported open-source storage orchestrator for containerized applications and documents Cloud Volumes ONTAP configuration as a supported backend. See the Astra Trident documentation and the ONTAP backend configuration documentation.

QUESTION NO: 7

Click the Exhibit button.

Referring to the exhibit, which compute server instance is able to natively modify file permissions in the "media-plex" share?

A. Ubuntu Server

B. SUSE Linux Enterprise Server

C. Microsoft Windows Server

D. Red Hat Enterprise Linux

ANSWER: C

Explanation:

Microsoft Windows Server is able to natively modify file permissions in the *media-plex* share because the share is accessed through the SMB/CIFS file-serving protocol and uses Windows-style security semantics. A Windows Server instance can connect to an SMB share and use its built-in File Explorer or Windows security tools to view and change NTFS discretionary access control lists (DACLs), subject to the permissions granted to the connected user. This includes managing permissions for users and groups directly from the file or folder Properties dialog, making it the native administrative environment for SMB share and NTFS ACL management.

NetApp ONTAP SMB shares support Windows clients and integrate with Microsoft security concepts such as Active Directory identities, share permissions, and NTFS file permissions. The effective access to data is determined through this Windows-compatible authorization model. As a result, a Windows Server compute instance is the appropriate system for natively administering the permissions associated with the displayed share. NetApp documents SMB support and Windows security-style access control in its ONTAP SMB administration guidance: NetApp ONTAP SMB administration documentation. Microsoft also documents how Windows manages file and folder permissions using NTFS ACLs: Microsoft NTFS overview.

QUESTION NO: 8

Click the Exhibit button.

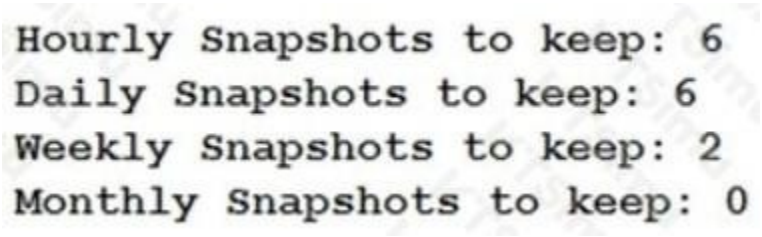

You have set the Snapshot policy shown in the exhibit on your Cloud Volumes Service for AWS.

On the 30th of the month at 10:00 p.m., how many Snapshot copies do you have?

A. 14

B. 12

C. 2

D. 6

ANSWER: B

Explanation:

12 is correct. At 10:00 p.m. on the 30th, the schedules configured in the displayed Snapshot policy that apply at that time have generated and retained their configured Snapshot copies. The policy's retention settings result in a combined total of twelve copies: the retained copies from each applicable schedule are counted together. Snapshot retention is governed by the number of copies to keep for each schedule, rather than by simply counting every time a snapshot operation has run historically. As new scheduled copies are created after the retention limit is reached, older copies for that schedule are automatically removed, maintaining the configured total. This behavior lets administrators retain short-term recovery points while controlling capacity consumption on the cloud volume. NetApp describes Snapshot policies as collections of schedules with associated retention counts, and explains that Snapshot copies are managed according to those policy rules. See the NetApp ONTAP documentation for creating Snapshot policies and the NetApp documentation on managing local Snapshot copies.

QUESTION NO: 9

You need to troubleshoot a failed BlueXP copy and sync relationship between an NFS server and an Amazon S3 bucket.

Which two items should you verify on the data broker? (Choose two.)

A. Network connectivity to the source and target

B. Appropriate permissions for the source and target

C. An ONTAP cluster-management LIF on the data broker

D. A FabricPool license on the data broker

ANSWER: A B

Explanation:

The data broker must have **network connectivity to the source and target**. It must be able to reach the NFS server and the Amazon S3 endpoint over the applicable network paths, including DNS, routing, proxy, and firewall configuration.

The data broker also requires **appropriate permissions for the source and target**. The source credentials must permit access to the exported data, and the AWS credentials or assigned IAM role must permit the required S3 bucket operations. See the BlueXP copy and sync data broker documentation and the supported sync relationships reference.

QUESTION NO: 10

Your customer has an E-Series and wants to bock up the volumes to the cloud.

Which two components are required to accomplish this task using SANtricity Cloud Connector? (Choose two.)

A. AWS S3 bucket

B. Windows 2016 Datacenter Edition as a backup host

C. Cloud Connector License key for SANtricity

D. Red Hat or SUSE Enterprise Linux Server as a backup host

ANSWER: A C

Explanation:

AWS S3 bucket is required because SANtricity Cloud Connector is designed to back up E-Series volume data to an object-store target, including Amazon Simple Storage Service. The bucket provides the cloud destination in which Cloud Connector creates and manages the backup objects. Before configuring backup jobs, the administrator needs the AWS environment, bucket, and credentials or access permissions that permit the connector to access that bucket.

Cloud Connector License key for SANtricity is also required to enable backup and restore functionality in the SANtricity Cloud Connector application. The license is installed and validated through the Cloud Connector management interface, allowing the application to run protected backup and recovery operations for E-Series volumes. This is separate from configuring the cloud target: the object-store destination determines where backup data resides, while the Cloud Connector license authorizes the application capability used to create and restore those backups.

NetApp documents SANtricity Cloud Connector as a host-based application for backing up E-Series volumes to cloud object storage and describes the licensing and configuration workflow in its Cloud Connector documentation. See the SANtricity Cloud Connector documentation and NetApp's Amazon AWS storage-location configuration guide.

QUESTION NO: 11

You are the administrator for your company's Cloud Volumes ONTAP instances in AWS. You are experiencing problems with OnCommand Cloud Manager. NetApp support asks you to provide log files.

In this situation, which two methods would accomplish this task? (Choose two.)

A. Set up a remote session with amazon Web Services (AWS) support to demonstrate the problem.

B. Go to the Support Dashboard in the Cloud Manager UI and click the Send Now button.

C. Use SSH to log in to the Cloud Manager instance and collect the logs from the "/opt/application/netapp/cloudmanager/log" directory.

D. Go to the AWS Management Console and collect the logs from AWS CloudWatch.

ANSWER: B C

Explanation:

Go to the Support Dashboard in the Cloud Manager UI and click the Send Now button is correct because Cloud Manager can generate and transmit an AutoSupport package directly to NetApp Support. This provides NetApp with diagnostic information collected by Cloud Manager, allowing the support case to be investigated without requiring the administrator to

manually transfer files. The on-demand send operation is useful when Support requests current diagnostic data after an issue has occurred.

Use SSH to log in to the Cloud Manager instance and collect the logs from the “/opt/application/netapp/cloudmanager/log” directory is also correct because this is the Cloud Manager host’s log location. An administrator can securely connect to the Connector/Cloud Manager instance, gather the relevant log files from that directory, and attach them to a NetApp Support case or provide them through the requested secure transfer method. This approach is particularly useful when Support needs specific files or when logs must be reviewed before submission.

NetApp documents AutoSupport and support-data collection procedures for BlueXP/Cloud Manager in its BlueXP setup and administration documentation and maintains Cloud Manager documentation at NetApp Cloud Manager documentation.