

NATSP 200, Public Cloud Services Accreditation Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

- 1. Which types of data can be moved to the cloud using NetApp Cloud Tiering service?**
 - A. File metadata.**
 - B. Infrequently accessed data.**
 - C. Object data.**
 - D. Backup data.**
- 2. How can clients continue to access their NAS data when a Cloud Volumes ONTAP instance fails in Amazon Web Services?**
 - A. The client connects to a virtualized IP address by using the border gateway protocol in another Virtual Private Cloud (VPC).**
 - B. Clients connect to a floating IP address that fails over to the partner in another Availability Zone.**
 - C. The load balancer forwards client traffic to the surviving high-availability (HA) partner.**
 - D. Clients connect to a NAS LIF that fails over to the surviving node in the HA pair.**
- 3. What is a benefit of the NetApp Cloud Sync service?**
 - A. Automated data conversion between formats**
 - B. Flexible hourly pricing per sync relationship**
 - C. Simplified development of in-house migration tools**
 - D. Unlimited data replication options**
- 4. A customer has a large FC storage environment and wants to implement a chargeback system for internal users. Which NetApp product should be suggested?**
 - A. Active IQ**
 - B. Cloud Sync service**
 - C. Virtual Storage Console**
 - D. OnCommand Insight**

- 5. In NetApp Cloud Insights, what is the role of an acquisition unit?**
- A. Represents a physical or virtual resource that NetApp Cloud Insights monitors**
 - B. Displays inventory and performance data for public cloud resources**
 - C. Analyzes the metrics, events, and logs for all monitored systems**
 - D. Collects and sends data from infrastructure assets to NetApp Cloud Insights for analysis**
- 6. What is a primary function of the NetApp Cloud Tiering service?**
- A. Data migration from local environments to NAS**
 - B. Periodically syncing data between local storage and cloud**
 - C. Tiering inactive data objects to cloud storage**
 - D. Providing on-premises storage solutions for cloud data**
- 7. What is a key advantage of using Azure ExpressRoute compared to traditional internet connections?**
- A. Lower latency and higher reliability**
 - B. More routing options available**
 - C. Simultaneous connections to multiple clouds**
 - D. Automatic failover capabilities**
- 8. What is required for an infrastructure data collector in NetApp Cloud Insights to collect data from a resource?**
- A. Service path**
 - B. Acquisition unit**
 - C. Annotation**
 - D. Query**
- 9. In a large FC infrastructure, which statement about OnCommand Insight is correct?**
- A. It can monitor the fabric, switch, and ports**
 - B. It can monitor the initiator and target, but not the switch**
 - C. It can monitor the initiator and switch only**
 - D. It is unaware of hardware changes in the fabric**

10. What is retained in the default backup policy for NetApp Cloud Backup Service for AWS?

- A. Seven backup copies of each volume**
- B. Thirty backup copies of each volume**
- C. Daily backups only**
- D. Weekly backups only**

SAMPLE

Answers

SAMPLE

- 1. B**
- 2. B**
- 3. B**
- 4. D**
- 5. D**
- 6. C**
- 7. A**
- 8. B**
- 9. B**
- 10. B**

SAMPLE

Explanations

SAMPLE

1. Which types of data can be moved to the cloud using NetApp Cloud Tiering service?

- A. File metadata.
- B. Infrequently accessed data.**
- C. Object data.
- D. Backup data.

The NetApp Cloud Tiering service is designed specifically for optimizing the storage of data by intelligently moving infrequently accessed data to a cloud storage solution. This function is particularly valuable in managing costs, as it allows businesses to retain frequently accessed files locally while offloading larger volumes of less frequently used data to the cloud. Infrequently accessed data is data that doesn't need to be readily available for day-to-day operations and includes items like older project files, archives, or historical datasets that may only be accessed occasionally. By tiering this type of data to the cloud, organizations can take advantage of the cost-effectiveness and scalability of cloud storage, making efficient use of their on-premises storage resources. In contrast, file metadata, object data, and backup data may not be eligible for movement via the Cloud Tiering service in the same context as infrequently accessed data. File metadata refers to information about files, object data implies different data structures typically managed in cloud object storage, and backup data usually has specific requirements for integrity and locality. Consequently, the most suitable category that can be effectively managed and moved using the Cloud Tiering service is indeed infrequently accessed data. This tiering capability underlines the flexibility and efficiency of cloud strategies in contemporary data management.

2. How can clients continue to access their NAS data when a Cloud Volumes ONTAP instance fails in Amazon Web Services?

- A. The client connects to a virtualized IP address by using the border gateway protocol in another Virtual Private Cloud (VPC).
- B. Clients connect to a floating IP address that fails over to the partner in another Availability Zone.**
- C. The load balancer forwards client traffic to the surviving high-availability (HA) partner.
- D. Clients connect to a NAS LIF that fails over to the surviving node in the HA pair.

Clients can continue to access their NAS data when a Cloud Volumes ONTAP instance fails in Amazon Web Services through connection to a floating IP address that is designed to fail over to the partner in another Availability Zone. This mechanism allows for high availability, as the floating IP acts as a single point of access for clients, seamlessly redirecting their connections to the backup partner when a failure occurs. By leveraging this floating IP address, clients are not required to change their connection settings or addresses; instead, they experience uninterrupted access to their NAS data due to the automatic redirection facilitated by the failover process. This method enhances resilience, as it is specifically designed to handle failover scenarios while ensuring minimal disruption for the clients who are actively accessing the data. It also simplifies the management of network configurations and failover procedures, as users are shielded from mundane backend processes and can maintain consistent access regardless of the state of any individual node within the high-availability setup.

3. What is a benefit of the NetApp Cloud Sync service?

- A. Automated data conversion between formats
- B. Flexible hourly pricing per sync relationship**
- C. Simplified development of in-house migration tools
- D. Unlimited data replication options

The benefit of the NetApp Cloud Sync service lies in its flexible hourly pricing per sync relationship. This model allows organizations to scale their data synchronization needs based on actual usage rather than being locked into a fixed or rigid pricing structure. Such flexibility can lead to cost savings, especially for businesses that have variable data transfer needs. Customers can establish multiple sync relationships with different pricing, making it easier to manage costs effectively as data demands change over time. This pricing strategy is particularly advantageous for enterprises that need to adapt to fluctuating workload requirements or for those running projects that may not require continuous sync, minimizing unnecessary expenses. With this approach, users can optimize their expenditures while ensuring they have the resources needed for their data syncing tasks when they arise.

4. A customer has a large FC storage environment and wants to implement a chargeback system for internal users. Which NetApp product should be suggested?

- A. Active IQ
- B. Cloud Sync service
- C. Virtual Storage Console
- D. OnCommand Insight**

The suggested product for implementing a chargeback system in a large FC storage environment is OnCommand Insight. This tool is designed specifically for managing and monitoring storage environments, providing detailed analytics and reporting features that can help organizations understand their storage usage patterns. With OnCommand Insight, customers can gain insights into resource consumption and usage trends across their storage systems. This allows for accurate allocation of costs to various internal users or departments based on their actual storage utilization. It helps administrators create a chargeback or showback model, facilitating more efficient budget management and ensuring accountability for storage resource consumption. Additionally, OnCommand Insight comes with a comprehensive dashboard and reporting capabilities, enabling users to track metrics such as performance, capacity, and efficiency. This visibility is crucial for developing a robust chargeback system. By quantifying storage expenditures and linking them to specific users or departments, organizations can make informed decisions about future storage investments and resource allocation. Other options, while useful in their own capacities, do not provide the specific chargeback capabilities needed in this scenario.

5. In NetApp Cloud Insights, what is the role of an acquisition unit?

- A. Represents a physical or virtual resource that NetApp Cloud Insights monitors**
- B. Displays inventory and performance data for public cloud resources**
- C. Analyzes the metrics, events, and logs for all monitored systems**
- D. Collects and sends data from infrastructure assets to NetApp Cloud Insights for analysis**

The acquisition unit in NetApp Cloud Insights functions as a critical component responsible for collecting and transmitting data from various infrastructure assets to the Cloud Insights platform for analysis. This means that acquisition units are tasked with gathering metrics, events, and logs from monitored systems, which includes both physical and virtual resources. By facilitating this data collection process, acquisition units enable Cloud Insights to provide comprehensive visibility into resource performance, health, and usage patterns. Without acquisition units, Cloud Insights would lack the necessary input to perform effective monitoring and analysis. They ensure that data flows smoothly from the monitored systems to the analytics engine of Cloud Insights, which is essential for maintaining operational awareness and optimizing resource management in cloud environments. Thus, the specific role of the acquisition unit in this context is fundamental to the functionality of NetApp Cloud Insights as a whole.

6. What is a primary function of the NetApp Cloud Tiering service?

- A. Data migration from local environments to NAS**
- B. Periodically syncing data between local storage and cloud**
- C. Tiering inactive data objects to cloud storage**
- D. Providing on-premises storage solutions for cloud data**

The primary function of the NetApp Cloud Tiering service is to tier inactive data objects to cloud storage. This means that the service intelligently moves data that is not actively being used from high-cost, on-premises storage to lower-cost cloud storage options. This process helps organizations optimize their storage costs and manage their data more efficiently by ensuring that only actively used data remains on the more expensive local storage systems. By automating the tiering process, Cloud Tiering allows businesses to take advantage of the scalability and cost-effectiveness of cloud storage solutions while maintaining easy access to their data when needed. This functionality is particularly useful in scenarios where data access patterns are unpredictable, as it helps in freeing up space for active workloads and reduces the overall storage footprint. The other options do not capture the essence of what Cloud Tiering does. While data migration and syncing might involve aspects of data movement, they do not specifically focus on the tiering of inactive data to optimize storage costs. Similarly, providing on-premises storage solutions focuses on local infrastructure rather than leveraging the benefits of cloud storage for tiered data management.

7. What is a key advantage of using Azure ExpressRoute compared to traditional internet connections?

- A. Lower latency and higher reliability**
- B. More routing options available**
- C. Simultaneous connections to multiple clouds**
- D. Automatic failover capabilities**

A key advantage of using Azure ExpressRoute compared to traditional internet connections is the provision of lower latency and higher reliability. ExpressRoute enables a private connection between your on-premises infrastructure and Azure data centers, bypassing the public Internet. This private pathway significantly reduces the variability and interruptions that can occur with internet-based connections, resulting in a consistent and dependable performance for various workloads. Lower latency in ExpressRoute connections positively impacts application performance, especially for data-intensive applications that require rapid responses. Additionally, the enhanced reliability stems from reduced exposure to common internet issues such as congestion, outages, and security vulnerabilities. This is essential for businesses that require stable, always-on access to their cloud services, making Azure ExpressRoute an ideal choice for performance-sensitive applications and environments that rely on high availability.

8. What is required for an infrastructure data collector in NetApp Cloud Insights to collect data from a resource?

- A. Service path**
- B. Acquisition unit**
- C. Annotation**
- D. Query**

An infrastructure data collector in NetApp Cloud Insights needs an acquisition unit to effectively collect data from a resource. An acquisition unit functions as the component responsible for connecting with the data source, retrieving necessary metrics, and transmitting them to Cloud Insights for analysis. This unit acts as the intermediary between the data source and the Cloud Insights platform, ensuring that data is accurately gathered and made available for monitoring and visualization. Other aspects, such as a service path, annotations, or queries, play different roles in the context of Cloud Insights. A service path may relate to the routing of data or service interactions, while annotations provide helpful metadata for entities but do not facilitate the direct collection of data. Queries are often used to retrieve or filter data post-collection but do not have an active role in the data gathering process itself, which is solely managed by the acquisition unit. Hence, the acquisition unit is essential and specifically designed for the data collection task in this scenario.

9. In a large FC infrastructure, which statement about OnCommand Insight is correct?

- A. It can monitor the fabric, switch, and ports**
- B. It can monitor the initiator and target, but not the switch**
- C. It can monitor the initiator and switch only**
- D. It is unaware of hardware changes in the fabric**

In a large Fibre Channel (FC) infrastructure, OnCommand Insight plays a pivotal role in monitoring the performance and health of various components. The correct statement notes that OnCommand Insight can monitor the initiator and target, emphasizing its capability to provide insights into the endpoints of the storage network. Understanding the roles of initiators and targets is crucial in FC environments. The initiator typically refers to the host system that sends requests for data, while the target refers to the storage devices that respond to those requests. By monitoring these components, OnCommand Insight can help in diagnosing performance issues, optimizing resource utilization, and ensuring that traffic between the initiators and targets flows smoothly. This focus on the data path makes OnCommand Insight an essential tool for those managing storage networks, as it allows for the identification of bottlenecks or failures at critical points within the storage infrastructure. While the tool has broad capabilities in terms of monitoring, the statement in the correct answer appropriately highlights its specific strengths in overseeing the relationship between initiators and targets in the system. In contrast, the other options either misrepresent its monitoring capabilities or suggest limitations that are not characteristic of OnCommand Insight's functionality. Understanding these capabilities is key for effective management and optimization of a Fibre Channel infrastructure in

10. What is retained in the default backup policy for NetApp Cloud Backup Service for AWS?

- A. Seven backup copies of each volume**
- B. Thirty backup copies of each volume**
- C. Daily backups only**
- D. Weekly backups only**

The default backup policy for NetApp Cloud Backup Service for AWS retains thirty backup copies of each volume. This retention policy allows for a significant number of restore points, which can be crucial for data recovery purposes. With a thirty-copy retention, users have the flexibility to recover data from a range of time, which can be advantageous in situations where data corruption or loss may have occurred at various points. This feature is particularly important for maintaining data integrity and availability in cloud environments where users need to safeguard against potential issues that may affect their volumes. The higher number of backup copies within the default policy supports better recovery options, ensuring that users can restore their data to a specific point in time, thereby minimizing data loss and enhancing operational continuity.