

DUMPSBOSS.

Nutanix Certified Professional Multicloud Infrastructure (NCP-MCI) 7.5

Nutanix NCP-MCI-7.5

Version Demo

Total Demo Questions: 12

Total Premium Questions: 120

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Describe the Nutanix Cloud Platform	2
Topic 2, Plan a Nutanix Cloud Platform Deployment	8
Topic 3, Deploy a Nutanix Cloud Platform	9
Topic 4, Configure and Manage a Nutanix Cloud Platform	79
Topic 5, Troubleshoot a Nutanix Cloud Platform	22
Total	120

QUESTION NO: 1

While configuring a remote site for PD-Based DR, an administrator elects to skip vStore mapping. Assuming the same container names don't exist on the two sites, what is the expected behavior of data placement at the remote site for replications?

- A. Replications fail to complete.
- B. The default container is always chosen.
- C. Container with lowest ID is chosen.
- D. A container is randomly chosen.

ANSWER: A

Explanation:

Replications fail to complete. In Protection Domain-based disaster recovery, a vStore mapping defines the destination storage location for data replicated from a protected container at the source site. When explicit vStore mapping is skipped, Nutanix relies on a destination container with the same name as the source container. This name-based behavior allows replication to proceed without a manually configured mapping only when an equivalent container name is available at the remote site. In this scenario, no identically named container exists at the remote site, so Nutanix cannot resolve a valid destination vStore for the replicated data. The replication operation therefore cannot place the snapshots and fails rather than selecting an arbitrary, lowest-ID, or default container. This behavior protects data placement consistency and ensures that replicated Protection Domain data is written only to an intended destination container. Administrators should either create matching container names on both sites or configure vStore mappings during remote-site configuration to map each source container to the appropriate remote container. Nutanix documentation for data protection and replication is available through the [Nutanix Documentation portal](#) and the [Nutanix product documentation page](#).

QUESTION NO: 2

An administrator configured an IPAM network in Prism Central for a development project. New VM deployments now fail during creation with an address-allocation error.

Existing VMs on the same network continue to operate normally, and the AHV hosts have network connectivity.

What is the most likely corrective action?

- A. Expand the IPAM address pool or reclaim unused addresses from the network.
- B. Configure a Data Services IP address for the cluster.
- C. Change the AHV uplink bond from active-backup to active-active.
- D. Assign a second default gateway to the affected VM subnet.

ANSWER: A

Explanation:

The administrator should expand the available IP address pool or reclaim unused IP addresses in the configured IPAM network. Prism Central IPAM assigns addresses from the network pool during VM deployment. When no addresses remain available, a new VM cannot receive an IP address even though the underlying AHV network and existing VMs remain functional.

Changing the VLAN ID can disrupt existing workloads and does not create additional IP addresses. A Data Services IP is used for cluster data services and is not a replacement for VM IPAM capacity.

QUESTION NO: 3

An administrator is trying to migrate a VM to a new container. The Migration of vDisks stalls.

What could be the cause of this?

- A. Insufficient storage space to create the source snapshot.
- B. Insufficient storage space in the target storage container.
- C. Insufficient storage space in the local storage container.
- D. Insufficient storage space to create the target snapshot.

ANSWER: B

Explanation:

Insufficient storage space in the target storage container is the likely cause. A Nutanix storage-container migration relocates the VM's virtual disks to the selected destination container. The target container must therefore have adequate available capacity to accept the migrated vDisk data and the temporary overhead associated with the migration workflow. If capacity is exhausted or cannot accommodate the required allocation, the data-transfer operation cannot continue normally and may appear stalled in Prism rather than completing the vDisk move. Before retrying, an administrator should review the target container's usage and available capacity in Prism, account for the VM's provisioned and actual data footprint, and either free capacity or select a container with sufficient space. Nutanix recommends monitoring storage capacity at the container level because containers are the policy and placement boundary used for VM data. The Prism documentation describes using the storage views to monitor containers and capacity, while the administration documentation covers VM storage and container operations. See [Nutanix Prism documentation](#) and [Nutanix storage container management documentation](#).

QUESTION NO: 4

An administrator is tasked with configuring a remote site for Data Protection. At the source site, the VMs are in a storage container named Marketing, but they need the data to be replicated to a storage container named Backups at the destination site.

What must be configured in order to accomplish this goal?

- A. Network mapping
- B. Network Compression
- C. vStore mapping
- D. vStore PD

ANSWER: C

Explanation:

vStore mapping must be configured because Nutanix Protection Domain replication uses a mapping to associate the source-site storage container with the intended target-site storage container. In this scenario, the source vStore/container is named *Marketing*, while the target vStore/container is named *Backups*. Since those names differ, Nutanix requires an explicit mapping so that replicated data is placed in the designated destination container rather than relying on an identically named remote container.


A vStore mapping is configured as part of the remote-site and protection-domain data-protection design. It defines the relationship between a local vStore and its remote counterpart, allowing administrators to use different container naming conventions at separate sites while maintaining predictable replication placement. This is particularly useful when a recovery site has containers organized by recovery purpose, retention policy, performance tier, or administrative ownership. After the mapping is established, protection-domain replication can send the protected VM data from *Marketing* to *Backups* as required.

For Nutanix documentation on Prism data protection and protection-domain replication configuration, see the [Nutanix Prism Web Console Guide](#) and the [Nutanix Support Portal](#).

QUESTION NO: 5

Refer to Exhibit:

Refer to the exhibit.



An administrator has been tasked with enabling LACP on an AHV cluster. In the exhibit, which option is the best to select for Bond Type when completing the task?

- A. Active-Backup
- B. No Uplink Bond
- C. Active-Active
- D. Active-Active with MAC Pinning

ANSWER: C

Explanation:

Active-Active is the appropriate AHV bond type for an uplink configuration that uses Link Aggregation Control Protocol (LACP). In AHV, this selection corresponds to the Open vSwitch *balance-tcp* bonding mode. It distributes network flows across the physical uplinks using a hash that can include layer-two, layer-three, and layer-four information, while allowing the AHV host and the connected physical switch ports to negotiate and maintain the aggregated link through LACP.

The physical switch ports connected to each AHV host must be configured as members of the same LACP port channel or link aggregation group, with matching VLAN and trunk settings. Selecting Active-Active ensures the host-side bond is configured for this negotiated, aggregated-uplink design and provides the intended load distribution and resiliency characteristics for the cluster's network traffic. Nutanix documents AHV virtual-switch and uplink-bond configuration, including the *balance-tcp* mode used for LACP, in the [AHV Administration Guide: Networking](#). Additional official guidance on AHV networking concepts is available in the [Nutanix AHV overview](#).

QUESTION NO: 6

The accounting department has been successfully utilizing Erasure Coding for a long time. During the weekend, there was a large increase in data usage. The administrator expected lower space usage on the storage container.

What could be the cause of this low space savings?

- A. Data Services IP is not configured
- B. Erasure coding auto deactivated
- C. Prism Central failure in the erasure coding cluster
- D. Erasure coding window default settings

ANSWER: D

Explanation:

Erasure coding window default settings is correct because Nutanix erasure coding is intended for write-cold data, rather than data that has just been created or recently changed. New writes are initially stored in the normal replicated layout so that the cluster can service active I/O efficiently. Nutanix waits for data to remain inactive for the configured erasure-coding delay before Curator can convert eligible extents to an erasure-coded layout. The default delay is seven days (604,800 seconds). Consequently, a substantial amount of data written during a weekend is still within that write-cold waiting period and continues to consume replicated capacity. The storage container therefore will not immediately display the space reduction normally associated with erasure coding, even though the feature remains enabled and has delivered savings for older data. Once the newly written data remains unchanged long enough to become eligible, Curator processing can apply erasure coding and the expected capacity savings can appear. Administrators can review data-reduction behavior and storage-container configuration in Prism and should account for this aging interval when evaluating capacity after a burst of writes. See Nutanix documentation on [storage container data reduction](#) and [erasure coding](#).

QUESTION NO: 7

An administrator uploads an operating-system ISO to the Prism Central Image Configuration service. The image is visible in Prism Central, but it is not available when creating a VM on a newly registered AHV cluster.

What action should the administrator take?

- A. Create a new storage container on the target cluster.
- B. Configure a Data Services IP address on the target cluster.
- C. Convert the ISO image to a volume group.
- D. Add the newly registered cluster to the image placement configuration.

ANSWER: D

Explanation:

Add the newly registered cluster to the image placement configuration. Prism Central images must be placed on the clusters where they will be used. Image placement makes the image available to the selected cluster for VM creation and other supported image operations.

Registering a cluster with Prism Central does not automatically place every existing image on that cluster. Creating a new storage container or configuring a Data Services IP does not resolve image availability for VM deployment.

QUESTION NO: 8

An organization protects a workload between two sites and requires a recovery point objective of 15 minutes.

Asynchronous replication is configured hourly and does not meet the business requirement. The sites have sufficient bandwidth and latency for a lower-RPO replication design.

Which replication type should the administrator configure?

- A. NearSync replication
- B. Asynchronous replication
- C. Metro Availability
- D. Protection-domain archival replication

ANSWER: A

Explanation:

NearSync replication is appropriate for a 15-minute RPO requirement. NearSync provides lower recovery point objectives than hourly asynchronous replication while avoiding the zero-data-loss and latency requirements associated with synchronous replication or Metro Availability.

Standard asynchronous replication is suitable for longer scheduled RPO values, such as hourly protection. Metro Availability is intended for synchronous active-active availability use cases and is not required merely to meet a 15-minute RPO.

QUESTION NO: 9

A database server has a corrupt database after a VM was restored from a recovery point. The database VM is part of the VIP_VMS category that is protected by a Protection Policy with asynchronous schedule.

After checking the state of NGT tools, an administrator is investigating the Protection Policy looking for the root cause of the database corruption.

What is the most probably cause?

- A. Set RPO is too high for protecting database VMs.
- B. Roll-up schedule is used in Protection Policy instead of Linear.
- C. The VM should have been added to the Applications category.
- D. Take App-Consistent Recovery Point option is not selected.

ANSWER: D

Explanation:

Take App-Consistent Recovery Point option is not selected. is the most probable cause because database workloads require a recovery point that represents a transactionally usable state, rather than merely a storage-level point-in-time image. When application-consistent recovery points are enabled, Nutanix uses Nutanix Guest Tools (NGT) to coordinate with the guest operating system and supported applications before creating the recovery point. This quiescing process helps flush pending writes and place the application into a consistent state, allowing the database to recover correctly after restoration.

If this setting is not enabled, asynchronous protection can create a crash-consistent recovery point. Such a point can be appropriate for many general workloads, but a database may have in-flight writes, uncommitted transactions, or data held in memory when the snapshot is taken. Restoring that state can leave the database unable to complete its normal crash recovery, resulting in corruption or an unusable database. Since the administrator has already verified that NGT is operational, the Protection Policy's omission of application-consistent recovery points directly identifies the relevant configuration issue. Nutanix describes protection-policy recovery-point behavior in its [Nutanix Support Portal documentation](#) and provides NGT deployment and operation resources through [Nutanix Support](#).

QUESTION NO: 10

An administrator has been tasked to remove a node from a Nutanix cluster running SQL VMs to help with a hardware refresh. After the node removal started, the administrator received the following error:

Node node_uuid cannot be removed: Cluster needs at least 5 usable nodes.

What action does the administrator need complete to remove the node?

- A. Add the replacement node to the cluster first and then remove the original node.
- B. Change the cluster Replication Factor (RF) to RF2 and then remove the node.
- C. Change the storage containers Replication Factor (RF) to RF2 and remove the node.
- D. Shutdown the CVM and use ncli to force the removal of the node.

ANSWER: A

Explanation:

Add the replacement node to the cluster first and then remove the original node. The cluster message indicates that the current configuration requires a minimum of five usable nodes to retain its configured resiliency while keeping data available. This commonly applies to a cluster using Replication Factor 3, which maintains three copies of data and is designed to tolerate simultaneous node failures. Nutanix prevents a removal that would leave insufficient eligible nodes because the cluster must preserve the required number of data replicas and continue to meet its fault-tolerance commitments during the operation.

For a hardware-refresh workflow, adding the replacement node first supplies the additional usable capacity needed for the cluster to satisfy that requirement. After the node has joined the cluster and its CVM and storage resources are healthy, Nutanix can redistribute data as necessary and maintain the required replication level. The administrator can then safely initiate removal of the old node without violating the minimum-node condition. This approach preserves availability and resilience for the SQL virtual machines throughout the refresh. Nutanix documents node expansion and removal procedures in the [Prism Web Console Guide](#) and describes replication-factor resiliency in the [Data Resiliency documentation](#).

QUESTION NO: 11

An administrator needs to protect a volume group using a protection policy that supports an RPO of 1 hour. The environment consists of two clusters in different geographic regions. Which replication type should be used?

- A. NearSync
- B. Synchronous
- C. Metro Availability
- D. Asynchronous

ANSWER: D

Explanation:

Asynchronous replication is the appropriate replication type because it supports recovery point objectives beginning at one hour and is designed for disaster-recovery use cases between geographically separated clusters. For a volume group, Nutanix Protection Domains and protection policies can use asynchronous replication to periodically transfer changed data from the source cluster to a remote cluster. This approach allows the recovery point to be based on the configured replication schedule, making a one-hour RPO an appropriate and supported target.

Geographic separation also makes asynchronous replication a practical fit because it is intended to tolerate the network latency and distance commonly associated with inter-region connectivity. The replication process sends snapshots according to the protection policy rather than requiring each write to be acknowledged by the remote site before it completes. Administrators should configure the protection policy schedule to meet the required one-hour RPO and ensure adequate bandwidth for the expected changed-data rate. Nutanix documentation describes asynchronous replication as the solution for longer RPO targets, including hourly schedules, while its lower-RPO replication technologies address much shorter recovery windows. See the [Nutanix Data Protection Guide](#) and [Nutanix Disaster Recovery](#) for replication and DR design guidance.

QUESTION NO: 12

An administrator wants to adjust how often the platform checks for security drift without modifying controls that run at fixed intervals.

Which setting should the administrator modify?

- A. The AIDE schedule
- B. The SCMA schedule

C. The list of security entities inspected by SCMA

D. The SCMA baseline enforcement behavior

ANSWER: B

Explanation:

The SCMA schedule is the appropriate setting because Security Configuration Management Automation is Nutanix's mechanism for periodically assessing the cluster against its security configuration baseline. Its schedule controls the cadence at which those assessments are performed, enabling an administrator to select how frequently the environment is checked for configuration drift while leaving the defined security controls and their individual fixed execution behavior intact. This directly satisfies a requirement to change only the frequency of the broader security-drift evaluation.

SCMA separates scheduling from the definition of what is evaluated and from remediation or enforcement choices. Therefore, adjusting its schedule changes when the platform performs its automated compliance assessment, rather than changing the baseline, the inspection scope, or the behavior of the controls themselves. Nutanix documents security configuration and compliance capabilities in its Security Guide and provides product documentation through the Nutanix Support Portal: [Nutanix Security Guide](#) and [Nutanix Support Portal documentation](#).