

Nutanix Certified Master

Nutanix NCM-MCI

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, Deploy and Configure a Nutanix Environment	30
Topic 2, Manage a Nutanix Environment	33
Topic 3, Troubleshoot a Nutanix Environment	30
Topic 4, Optimize a Nutanix Environment	26
Total	119

QUESTION NO: 1

A customer needs a DR solution designed around the following requirements:

- There are 8 ms of latency between Sites A and B
- The administrator needs the shortest RPO/RTO possible
- Application consistent snapshots are required
- Use of Nutanix VSS

Which Data Protection Method should be used?

- A. Async
- B. Metro Availability
- C. Near Sync
- D. Cloud Connect

ANSWER: C

Explanation:

Near Sync is the appropriate protection method because it is designed for workloads requiring a very low recovery point objective while operating across sites with network latency suitable for near-synchronous replication. It can provide an RPO as low as one minute, substantially reducing potential data loss and enabling a fast recovery workflow at the remote site. An 8 ms intersite latency falls within the expected connectivity characteristics for this type of deployment.

Near Sync uses frequent replication of protection-domain snapshots to the remote site. For supported Windows workloads, Nutanix Guest Tools and the Nutanix VSS provider can quiesce applications before the snapshot is taken. This produces application-consistent recovery points rather than merely crash-consistent copies, allowing applications such as Microsoft workloads to recover from a transactionally consistent state. The combination of minute-level replication and VSS-based application consistency directly satisfies the stated RPO/RTO and recovery-consistency requirements. Nutanix describes its disaster-recovery capabilities at [Nutanix Disaster Recovery](#), and provides product documentation through the [Nutanix Documentation Portal](#).

QUESTION NO: 2

An administrator protected a DB VM running on a Nutanix cluster with NearSync Replication, which schedules to repeat every 15 minutes to a second site. The daily change rate is low. During nightly backups, the replication window is missed due to write volume. Which two options describe the expected behavior? (Choose two.)

- A. NearSync stops working and fails over to the remote site.
- B. Protection Domain transitions back to the hourly schedule.
- C. The Administrator needs to re-enable it manually.
- D. The system continuously tries to get to the NearSync schedule.

ANSWER: B D

Explanation:

NearSync Replication is designed to protect workloads with a short recovery-point objective while handling temporary periods in which the source workload produces more changed data than can be replicated within the NearSync interval. When the nightly backup causes write volume high enough to miss the 15-minute replication window, the protection domain automatically transitions back to its hourly asynchronous replication schedule. This provides continued protection without requiring an administrator to intervene or manually re-enable replication.

While operating on the fallback schedule, Nutanix continuously evaluates whether replication can again meet the configured NearSync objective. As the backup activity ends and the changed-data rate returns to a level that the available replication

bandwidth can accommodate, the system continuously tries to get to the NearSync schedule and resumes NearSync operation automatically when it can do so. This adaptive fallback and recovery behavior preserves data protection during a temporary replication backlog while avoiding an unnecessary failover. Nutanix documents NearSync as an asynchronous disaster-recovery capability with automatic handling for workloads that temporarily cannot sustain the configured short replication interval. See the [Nutanix Prism Web Console Guide: Data Protection](#) and [Nutanix Disaster Recovery](#).

QUESTION NO: 3

An administrator receives an error indicating that the CVMs in the cluster are not syncing to any NTP servers. An investigation of the issue finds:

- The NTP servers are configured in Prism

- The time on all CVMs is the same
- Both the CVMs and AHV hosts are configured for the UTC time zone

Which two steps can be taken to troubleshoot this issue? (Choose two.)

A. Confirm that the NTP servers are reachable from the CVMs.

B. Restart genesis on all CVMs.

C. On a CVM, run the command `allssh ntpq -pn`.

D. Restart the chronos service on all CVMs.

ANSWER: A C

Explanation:

Confirming that the NTP servers are reachable from the CVMs is an essential first diagnostic step because CVMs, rather than merely Prism configuration, must be able to communicate with the configured NTP sources over the network. Validate name resolution when hostnames are used, routing, and firewall policy for NTP traffic on UDP port 123. A successfully configured NTP source in Prism does not by itself prove that each CVM can reach or receive responses from that source.

Running `allssh ntpq -pn` from a CVM provides cluster-wide visibility into the NTP daemon's peer associations. The command displays configured peers by numeric address and reports association state, reachability, polling information, delay, offset, and jitter. This makes it possible to determine whether an NTP server is being contacted and whether a CVM has selected a synchronization source. Reviewing results from every relevant node is especially useful when the reported issue is cluster-wide, while the observed clocks currently appear consistent. Nutanix documents NTP configuration and validation as part of Prism and cluster administration; see the [Prism Web Console Guide](#) and the [AHV Administration Guide](#).

QUESTION NO: 4

A systems administrator needs to add more VMs to their Nutanix cluster.

Which two actions should the administrator perform to determine if the current cluster could accommodate the new VMs? (Choose two.)

A. Perform an Inventory with Life Cycle Management.

B. Determine utilization with Cluster Runway.

C. Enable Deduplication and Erasure Coding.

D. Utilize Optimize Resources for VM efficiency.

ANSWER: B D

Explanation:

"Determine utilization with Cluster Runway" is correct because Cluster Runway provides a capacity-planning view based on the cluster's historical resource-consumption trends. It helps an administrator assess available and projected CPU, memory,

and storage capacity, identify when one of those resources is expected to become constrained, and determine whether the anticipated VM workload can fit without creating an unacceptable capacity risk. This is the appropriate first step when evaluating whether existing infrastructure has sufficient headroom for growth.

“Utilize Optimize Resources for VM efficiency.” is also correct because Prism’s resource optimization capabilities help identify virtual machines whose allocated resources are not aligned with actual demand. Reviewing these recommendations can reveal opportunities to rightsize VM CPU and memory allocations, reclaim resources from oversized workloads, and improve usable capacity before new VMs are deployed. Used together, Cluster Runway supplies the cluster-level forecast while Optimize Resources identifies VM-level efficiency improvements that can increase available headroom. Nutanix describes Prism as providing infrastructure monitoring, operational insights, and optimization capabilities at [Nutanix Prism](#). Additional Nutanix capacity-planning and management documentation is available through the [Nutanix Support Portal documentation library](#).

QUESTION NO: 5

An administrator is implementing a VDI solution. The workload will be a series of persistent desktops in a dedicated storage container within a four-node cluster. Storage optimizations should be set on the dedicated storage container to give optimal performance including during a node failure event.

Which storage optimizations should the administrator set to meet the requirements?

- A. Compression, Deduplication, and Erasure Coding
- B. Deduplication and Erasure Coding
- C. Compression and Deduplication
- D. Compression only

ANSWER: C

Explanation:

Compression and Deduplication is the appropriate storage-optimization combination for a dedicated container holding persistent VDI desktops. Persistent desktops commonly contain highly repetitive operating-system files, application binaries, patches, and similar user desktop images. Nutanix deduplication reduces the physical capacity consumed by those repeated blocks and can reduce unnecessary storage I/O. Compression further reduces the capacity footprint of the remaining data, allowing the container to accommodate more desktops while maintaining efficient use of the cluster’s storage resources.

For a VDI workload, these efficiency features are especially useful because many desktop instances have substantial commonality even when each user retains a persistent desktop. The cluster’s replication-based data protection continues to provide data availability through a node failure event, while the selected data-reduction settings avoid adding an erasure-coding data layout to this performance-sensitive desktop container. Storage optimization settings are configured at the container level, allowing the administrator to apply them specifically to the VDI workload without affecting containers used for other application profiles.

See Nutanix documentation resources at [Nutanix Documentation](#) and the Nutanix storage efficiency overview in the [Nutanix Bible](#).

QUESTION NO: 6

An administrator is notified that a bare metal database server is down. This database server is being served storage using a Nutanix Volume Group. Upon investigating, the administrator finds that the disks in the database server that map to the vDisks in the volume group have gone offline.

What is causing this issue?

- A. The Volume Group Load Balancer has been disabled.
- B. Port 9443 is blocked in the server firewall.
- C. Port 3260 has been blocked in the server firewall.

D. A CVM serving the Volume Group has gone offline.

ANSWER: C

Explanation:

Port 3260 has been blocked in the server firewall. Nutanix Volume Groups present storage to external, bare-metal clients through iSCSI. The client operating system connects its iSCSI initiator to the Volume Group target and maintains iSCSI sessions over TCP port 3260. If a firewall rule on the database server prevents this traffic, the initiator cannot maintain or re-establish its connection to the target. The operating system then loses access to the LUNs represented by the Volume Group vDisks, causing the mapped database disks to appear offline and potentially taking the database service down.

Restoring TCP connectivity on port 3260 between the bare-metal host and the applicable Nutanix Volume Group target addresses is required. The administrator should validate both outbound client traffic and any relevant return-path filtering, then confirm that iSCSI sessions reconnect and that the host rescans its storage devices. TCP port 3260 is the registered iSCSI target service port, as shown in the [IANA service-name and port-number registry](#). Nutanix Volume Groups use iSCSI connectivity for external client access; consult the applicable [Nutanix documentation](#) for the platform release in use.

QUESTION NO: 7

An organization is deploying multiple isolated tenant environments on an AHV cluster using Flow Virtual Networking.

Which two constructs are used to provide logical network isolation for the tenant environments? (Choose two.)

- A. Virtual Private Cloud (VPC)
- B. Subnet
- C. Storage container
- D. Volume Group

ANSWER: A B

Explanation:

A Virtual Private Cloud (VPC) provides a logically isolated networking boundary in Flow Virtual Networking. Each tenant can use a separate VPC to isolate its virtual network topology, routing, and address space from other tenant environments.

Subnets are created within a VPC and provide the IP network segments to which VM interfaces are connected. A subnet can define the associated address range and gateway configuration for a particular application tier or tenant segment. Storage containers and Volume Groups are storage constructs and do not provide logical tenant network isolation. See the [Nutanix Flow Virtual Networking overview](#) and the [Nutanix Documentation portal](#).

QUESTION NO: 8

An administrator needs to make sure an RF2 cluster can survive a complete rack failure without negatively effecting workload performance. The current cluster configuration consists of the following:

- 30 nodes: distributed 10 nodes per rack across three 42U rack

- Each node is configured with 40TB usable storage all flash (Cluster Total 1.2 PB Usable)
- Current cluster utilization is 900TB storage

Which configuration changes should be made to make sure that the cluster meets the requirements?

- A. Expand the cluster by adding 3 additional nodes per rack, 9 nodes total.
- B. Expand the cluster by adding 2 additional nodes per rack, 6 nodes total.
- C. Expand the cluster by adding 9 nodes to a new 42U rack

D. Expand the cluster by adding 8 nodes to a new 42U rack

ANSWER: A

Explanation:

Expand the cluster by adding 3 additional nodes per rack, 9 nodes total is correct because the added capacity must be distributed evenly across the three existing racks. This preserves rack-aware fault-domain symmetry while ensuring that, after any one rack is lost, the surviving two racks retain sufficient storage headroom for the existing 900 TB workload and the RF2 resilience operations that follow the failure. With 13 nodes in each rack, a rack failure leaves 26 nodes. At 40 TB usable per node, the surviving cluster provides 1,040 TB usable capacity. That leaves 140 TB free, or approximately 13.5% headroom, after hosting the current 900 TB consumption. This free capacity is necessary to avoid running the surviving nodes at an excessively high utilization level, which can adversely affect performance and operational flexibility during data re-protection and normal workload activity. The design also keeps node count, storage capacity, and workload distribution balanced across all rack fault domains before a failure occurs. Nutanix capacity planning considers failure-domain tolerance and available capacity rather than only total installed capacity; see the [Nutanix Capacity Planning Guide](#) and [Nutanix Enterprise Storage overview](#).

QUESTION NO: 9

A new Nutanix cluster is ready for production on a remote site with no IT staff. Changes to the infrastructure will cost money and probably cause a cluster stop.

The administrator needs to perform a health check on the environment that should include:

- Network connectivity checks
- Network performance status ▪ Storage performance status

Which tool should be used to perform this test?

- A. X-Ray
- B. curator_cli
- C. NCC
- D. diagnostic, py

ANSWER: A

Explanation:

X-Ray is the appropriate tool because it is Nutanix's automated test and validation framework for assessing a cluster's operational readiness and performance. It can execute a defined test workload from a controller VM and test VMs, then collect results that validate the environment before it is handed over for production use. Its available test coverage includes network connectivity validation, network-performance measurement, and storage-performance testing—the specific areas required by this health check.

This is particularly suitable for a remote site because the administrator can run a repeatable test plan and review the generated results remotely, rather than relying on manual, ad hoc verification. X-Ray is designed to validate infrastructure behavior under controlled test conditions, helping establish a baseline for the newly deployed cluster and identify issues before production workloads are introduced. Nutanix describes X-Ray as a tool for testing and validating the performance and functionality of Nutanix environments; see the [Nutanix X-Ray product page](#) and the [Nutanix X-Ray Guide](#).

QUESTION NO: 10

An administrator needs to replace the self-signed certificate on a cluster. Which two requirements must be met as part of the process? (Choose two.)

- A. The cluster administrator must restart the interface gateway.

- B. The signed, intermediate and root certificates are chained.
- C. The existing certificate must be deleted prior to replacement.
- D. The imported files for the custom certificate must be PEM encoded.

ANSWER: B D

Explanation:

Replacing the cluster's self-signed TLS certificate requires a complete trust chain and certificate material in the format Prism expects. The signed, intermediate and root certificates are chained because clients must be able to validate the certificate presented by the cluster back to a trusted root certificate authority. The server certificate should be followed by the applicable intermediate CA certificate or certificates, in the proper order, so browsers, APIs, and other clients can establish the full chain of trust when connecting to Prism.

The imported files for the custom certificate must be PEM encoded. PEM is the text-based encoding convention used for the certificate and associated private-key material during the Prism custom-certificate workflow. PEM content is identifiable by its BEGIN/END header and footer lines, such as -----BEGIN CERTIFICATE----- . Using a correctly encoded PEM certificate and a complete CA chain helps ensure that Prism can import the custom certificate successfully and that connecting clients can validate it without certificate-chain errors. Nutanix documents certificate configuration as part of Prism security administration; see the [Prism Web Console Guide](#) and the [Prism Central Guide](#).

QUESTION NO: 11 - (HOTSPOT)

HOTSPOT

An administrator is configuring an All Flash cluster to support the following workloads: ■ Full clone virtual desktops

- Hadoop
- Write heavy databases

Which correct configuration options should the administrator use to meet the item requirements? For each procedure, indicate whether it is a correct or incorrect procedure that must be configured to meet the item requirements using the drop-down option.

Hot Area:

Answer Area

Procedure

Correct/Incorrect

One container with compression and deduplication enabled for full clone virtual desktops

	▼
Correct	
Incorrect	

One container with compression and Erasure Coding enabled for Hadoop

	▼
Correct	
Incorrect	

One container with compression enabled for heavy databases

	▼
Correct	
Incorrect	

One container with deduplication enabled for full clone virtual desktops

	▼
Correct	
Incorrect	

One container with Erasure Coding enabled for Hadoop

	▼
Correct	
Incorrect	

One container with deduplication enabled for Hadoop

	▼
Correct	
Incorrect	

ANSWER:

Answer Area

Procedure

One container with compression and deduplication enabled for full clone virtual desktops

Correct/Incorrect

☐

Correct

Incorrect

One container with compression and Erasure Coding enabled for Hadoop

☐

Correct

Incorrect

One container with compression enabled for heavy databases

☐

Correct

Incorrect

One container with deduplication enabled for full clone virtual desktops

☐

Correct

Incorrect

One container with Erasure Coding enabled for Hadoop

☐

Correct

Incorrect

One container with deduplication enabled for Hadoop

☐

Correct

Incorrect

Explanation:

Nutanix storage-efficiency features should be enabled according to the data pattern and operational behavior of each workload. Full-clone virtual desktops contain a high degree of repeated operating-system, application, and user-profile data across many separate virtual machines. Deduplication is therefore appropriate because it identifies duplicate data blocks and retains a single physical copy. Compression is also suitable in this container and further reduces the physical capacity consumed by data that cannot be eliminated as duplicates. Enabling both features provides the expected space-efficiency configuration for a full-clone VDI workload on an all-flash cluster.

Hadoop commonly stores large data sets that are capacity intensive and generally fit a data-protection approach optimized for usable capacity. Erasure Coding provides fault tolerance with materially lower capacity overhead than maintaining full replicated copies, making it an appropriate configuration for the Hadoop container. Compression is compatible with this objective because reducing the size of eligible data before storage further improves capacity efficiency. Nutanix explains Erasure Coding as a capacity-saving resiliency technology in its [Prism Web Console Guide](#), and documents the available data-reduction controls, including compression and deduplication, in the [data reduction documentation](#).

Write-heavy database data is an appropriate candidate for compression, which lowers consumed capacity without requiring the workload characteristics that make deduplication especially valuable. Conversely, Hadoop does not receive the same expected benefit from deduplication as identical full-clone desktop images do. The correct selections therefore enable compression and deduplication for full-clone desktops, compression and Erasure Coding for Hadoop, compression for write-heavy databases, deduplication for full-clone desktops, and Erasure Coding for Hadoop, while leaving deduplication for Hadoop marked Incorrect.