PassQueen
a.k.a dumpsqueen.com

Nutanix Certified ProfessionalBusiness Continuity (NCP-BC) 7.5

Nutanix NCP-BC-7.5

Version Demo

Total Demo Questions: 11

Total Premium Questions: 116

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Business Continuity and Disaster Recovery Concepts	3
Topic 2, Data Protection	46
Topic 3, Disaster Recovery	26
Topic 4, Nutanix Leap	41
Total	116

QUESTION NO: 1

An organization defines a four-hour RPO and a one-hour RTO for a business application. A primary-site outage occurs immediately before the next scheduled recovery point completes. The administrator recovers the application at the recovery site in 45 minutes.

Which two statements are correct? (Choose two.)

A. The potential data loss can approach four hours.

B. The recovery operation exceeds the defined RTO.

C. The recovery operation guarantees zero data loss.

D. The application recovery completes within the defined RTO.

ANSWER: A D

Explanation:

The RPO defines the maximum acceptable amount of data that can be lost, while the RTO defines the maximum acceptable time to restore service. Because the outage occurred immediately before the next recovery point completed, the potential data loss can approach the four-hour RPO. Completing recovery in 45 minutes meets the one-hour RTO.

A successful recovery within the RTO does not reduce the possible data loss below the RPO. Likewise, the availability of a recent recovery point does not guarantee zero data loss when asynchronous replication is used.

QUESTION NO: 2

An administrator needs to perform a test failover on a Recovery Plan. What is the primary goal of this action?

A. Copy data from primary to recovery location before failover happens.

B. Perform a live production cutover and permanently switch all workloads to the recovery site.

C. Validate recoverability and network mappings without impacting production systems.

D. Disable replication before a planned maintenance window.

ANSWER: C

Explanation:

Validate recoverability and network mappings without impacting production systems. A test failover is a non-disruptive validation operation used to confirm that the Recovery Plan can recover protected workloads successfully at the recovery site. Nutanix Disaster Recovery uses the replicated recovery data to bring up the workload in a test context, enabling the administrator to verify that virtual machines boot, recovery-plan sequencing and dependencies operate as intended, and application services are available. It also provides an opportunity to verify the configured network mapping and to ensure recovered workloads can be tested without creating conflicts with the active production environment. This validation gives the organization practical evidence that its recovery configuration supports its business-continuity objectives before an actual outage requires a real failover. Test failover is therefore an essential periodic DR-readiness activity, rather than a production migration or an action that changes protection behavior. Nutanix describes recovery plans and their execution options in its Disaster Recovery documentation, including the use of test recovery operations to validate recoverability. See the Nutanix Disaster Recovery Guide and the Nutanix Disaster Recovery overview.

QUESTION NO: 3

An administrator has VMs in a container named **VMData** on Cluster A that need to be replicated to Cluster B. In order to configure new replications to Cluster B, the administrator creates a container named **VMData** on Cluster B with an advertised capacity of 200GB.

The VMs on Cluster A in the VMData container have these characteristics:

Large VMs taking up 4TB of space

Low data churn rate

100GB of recovery points

NGT is installed

Recovery points are application-consistent

After replications begin the administrator quickly notices replications are not progressing.

What best explains what happened?

A. Low data churn rate of source VMs slows the replication rate to the remote site.

B. Replications can't proceed if there are duplicate container names.

C. First replication is a full copy of the VMs, and the container isn't large enough.

D. Application-consistent snapshots take up double the space on remote clusters.

ANSWER: C

Explanation:

First replication is a full copy of the VMs, and the container isn't large enough. When Nutanix protection-domain replication is first established, the remote site must receive a baseline copy of the protected VM data before later recovery points can be transferred incrementally. The source workload occupies approximately 4 TB, but the target VMData container has only 200 GB of advertised capacity. Consequently, Cluster B cannot accommodate the initial replication baseline and the replication cannot make meaningful progress. The destination must be sized to hold the replicated VM data and provide additional headroom for recovery-point retention, metadata, and ongoing changes while replication is active. The low churn rate affects the amount of data sent after the baseline is established; it does not reduce the capacity needed for the initial copy. Likewise, application-consistent recovery points and NGT support enable suitable recovery-point behavior, but the immediate operational requirement is adequate destination capacity for the initial protected dataset. After increasing the destination container capacity appropriately, the administrator can allow the baseline replication to complete and then retain the required recovery points. Nutanix documents protection-domain replication and recovery-point operation in its Prism Web Console Guide and provides DR architecture guidance at Nutanix Disaster Recovery.

QUESTION NO: 4

Which two VM disk types does Self-Service restore support? (Choose two.)

A. SCSI

B. IDE

C. PCI

D. VGs

ANSWER: A B

Explanation:

Self-Service Restore supports VM disks presented through the SCSI and IDE controller types. This capability enables an authorized VM user to browse files from a selected recovery point and restore individual files without requiring a full VM recovery workflow. Nutanix Guest Tools (NGT) provides the in-guest components required for Self-Service Restore. When a user initiates a restore, Nutanix makes the relevant recovery-point disk available to the guest so its file system can be browsed and the required files can be copied back to the active VM.

SCSI is the normal controller type for many virtual data disks and is therefore supported for file-level recovery. IDE is also supported, which is particularly relevant for operating-system or boot disks in configurations where IDE is used. In both cases, the recovery operation is VM-centric: the protected VM disk from the snapshot is used as the source for restoring files into the running workload. This allows recovery of a specific file or directory while avoiding the impact and time associated with restoring an entire virtual machine. Nutanix documentation describes Self-Service Restore as an NGT-enabled, file-level restore capability for protected VMs. See the Nutanix Data Protection Guide: Self-Service Restore and the Prism User Guide: Self-Service Restore.

QUESTION NO: 5

An organization wants authorized application owners to perform file-level recovery from protected Windows VMs without requiring a full VM restore.

Which two conditions are required for Nutanix Self-Service Restore to support this workflow? (Choose two.)

A. NGT must be installed and operational on the protected VM.

B. The VM disk must use a supported SCSI or IDE controller.

C. The VM must be protected by Synchronous replication.

D. A full clone of the VM must be created before files can be restored.

ANSWER: A B

Explanation:

Self-Service Restore requires Nutanix Guest Tools (NGT) on the protected VM because the workflow uses guest-level components to present and browse data from the selected recovery point. The source disk must also use a supported VM disk controller type, such as SCSI or IDE.

A full VM clone is not required for file-level recovery, and synchronous replication is not a prerequisite for Self-Service Restore.

QUESTION NO: 6

An administrator performed an unplanned failover from AZ1 (primary) to AZ2 (recovery). Later, AZ1 is restored and the administrator wants to fail back the entities to AZ1. Which two statements are correct? (Choose two.)

A. Failback uses the same recovery plan, but the failover operations for failback must be performed at the primary AZ.

B. Unplanned failover to AZ1 can use recovery points that were created locally in the past at AZ1, even if no recovery points replicate back from AZ2.

C. After unplanned failover, replication begins in the reverse direction, and failback success depends on recovery points being replicated back from AZ2 to AZ1 after failover.

D. After unplanned failover, failback success depends on recovery points being created locally on AZ1 before outage.

ANSWER: B C

Explanation:

After an unplanned failover, the recovery location becomes the active production location. To protect the changes made while workloads run there, Nutanix Disaster Recovery reverses the replication direction: new recovery points are generated at AZ2 and replicated to the restored AZ1. Therefore, “After unplanned failover, replication begins in the reverse direction, and failback success depends on recovery points being replicated back from AZ2 to AZ1 after failover” correctly describes the normal failback workflow. Waiting for reverse replication ensures the restored site has a current, usable recovery point containing the post-failover changes before workloads are returned.

“Unplanned failover to AZ1 can use recovery points that were created locally in the past at AZ1, even if no recovery points replicate back from AZ2” is also correct as a recovery-of-last-resort capability. If the recovery location cannot be used or reverse replication cannot complete, AZ1 can be activated from recovery points retained locally before its original outage. This restores service from the most recent locally available point, although it does not include changes made at AZ2 after the original failover. Nutanix describes these failover and recovery workflows in its Disaster Recovery documentation and provides an overview of bidirectional DR capabilities on its Nutanix Disaster Recovery product page.

QUESTION NO: 7

After a storage failover in a synchronous replication environment, users report higher I/O latency and degraded performance compared to pre-failover levels. Cluster health is normal, and replication is functioning correctly. Which post-failover cleanup action could the administrator perform to correct the performance issue?

A. Increase the snapshot retention period of the affected VMs.

B. Manually realign the compute and storage locations of the affected VMs.

C. Reconfigure the replication schedule of the affected VMs.

D. Wait for Acropolis Dynamic Scheduler to rebalance the workloads of the affected VMs.

ANSWER: B

Explanation:

Manually realign the compute and storage locations of the affected VMs is the appropriate cleanup action because a synchronous-replication storage failover can leave VM compute running at one site while the active storage copy is being served from the other site. In that state, guest reads and writes require network traversal between the VM’s host and the remote active storage location. Although both clusters and replication may remain healthy, this loss of data locality increases I/O latency and can reduce workload performance.

After the failover is stable, the administrator should live-migrate the affected VMs’ compute to the cluster that currently owns and serves the active storage. This restores local data access between the VM and its storage path, avoiding unnecessary intersite I/O and returning performance closer to normal pre-failover behavior. The action is especially relevant when the event was a storage-only failover rather than a coordinated application or VM recovery operation. Nutanix Metro Availability is designed to provide synchronous replication and continuous availability, while operational recovery procedures should ensure that workload placement is aligned with the active site. See Nutanix’s Metro Availability overview and the Prism Central Metro Availability documentation.

QUESTION NO: 8

An administrator at a retail company is preparing to apply a critical OS security patch to a production SQL Server VM. Before starting, the administrator ensures the VM is part of a Protection Domain and manually creates a local snapshot (Recovery Point). Ten minutes after the patch is applied, the SQL service fails to start, and the database logs indicate corruption in the system registry. The administrator needs to return the VM to its functional state from ten minutes ago as quickly as possible. The administrator decides to perform a Restore from the local snapshot. Under which condition would an In-place Restore (Revert) fail, forcing the administrator to use the Clone (Out-of-place) option instead?

A. NGT (Nutanix Guest Tools) version was updated immediately after the snapshot.

B. The VM ' s memory was increased from 16GB to 32GB to accommodate the patch installation.

C. The VM is currently powered on and cannot be shut down.

D. A new 50GB data disk was added to the VM configuration after the snapshot was taken.

ANSWER: D

Explanation:

A new 50GB data disk was added to the VM configuration after the snapshot was taken. An in-place restore returns the existing VM's protected virtual disks to the selected recovery point, so the VM's disk topology must remain compatible with the disk set represented by that recovery point. Adding a vDisk after the recovery point creates a configuration mismatch: the current VM has a disk for which the selected snapshot has no corresponding point-in-time disk state. In this situation, the administrator cannot safely revert the existing VM in place to that recovery point. The out-of-place restore method is appropriate because it creates a separate VM from the recovery point, using the VM and vDisk layout captured at that time. This provides a clean recovery target with the pre-patch operating-system and registry state, without attempting to reconcile the later-added disk with the older recovery point. Nutanix documents recovery-point restoration and cloning workflows in the Prism Web Console documentation; see the Protection Domains documentation and the Prism data-protection guide.

QUESTION NO: 9

What must an administrator verify before starting a migration of entities from a protection domain to a protection policy?

A. The entities must be added to a category in Prism Central.

B. The recovery plan must be active.

C. The entities must have no ongoing replication tasks.

D. A full backup of the protection domain must be completed.

ANSWER: C

Explanation:

The entities must have no ongoing replication tasks. Migration from a legacy protection domain to a protection policy changes the protection configuration and transfers management of the selected entities to the Prism Central policy-based workflow. Before initiating that operation, the administrator must confirm that replication activity for the relevant entities is idle and no replication task is in progress. This provides a stable protection state for the migration workflow and prevents the migration from overlapping an active data-transfer operation. Nutanix protection-domain replication is snapshot-based and can involve asynchronous or synchronous replication activity; the migration should be started only after the current replication work has completed and task status is clear. Administrators can use the Prism Element or Prism Central task and data-protection views to verify that no replication is running, then proceed with the entity migration and subsequent policy assignment. This check is specifically a pre-migration operational requirement, ensuring the system is not simultaneously modifying replication state while the protection-management model is being changed. Refer to the Nutanix documentation portal for the applicable Prism Central and Data Protection guides: Nutanix Documentation Portal. Nutanix also provides product documentation and support resources at Nutanix Product Support.

QUESTION NO: 10

An administrator is preparing to run a Recovery Plan at a recovery AZ for the first time. The organization requires the plan to start workloads successfully without relying on last-minute infrastructure changes.

Which two items should the administrator validate before executing the Recovery Plan? (Choose two.)

A. Validate that the recovery AZ has sufficient compute and storage capacity.

B. Validate that all protected VMs have NGT installed.

C. Validate that the primary and recovery subnets are stretched.

D. Validate that the Recovery Plan network mappings target available recovery-site networks.

ANSWER: A D

Explanation:

The recovery AZ must have sufficient compute and storage capacity to host the recovered workload, and the Recovery Plan must contain valid mappings to recovery-site networks. These prerequisites ensure that recovered VMs can be placed and connected according to the recovery design.

NGT is useful for supported guest-level functions, such as application-consistent snapshots and Self-Service Restore, but it is not a universal prerequisite for every Recovery Plan failover. A stretched subnet is required only when the design requires a VM to retain its original IP address; it is not required for all recovery plans.

QUESTION NO: 11

An administrator wants to protect the snapshots created on the cluster. Only authorized users should be allowed to modify or delete the snapshots on the cluster. How can the administrator harden the security of the snapshots?

A. Configure CVM and AHV hardening.

B. Configure a backup user with admin privileges.

C. Configure cluster lockdown and use SSH keys.

D. Configure Approval policy for snapshots.

ANSWER: D

Explanation:

Configure Approval policy for snapshots. An approval policy introduces an authorization workflow for sensitive snapshot operations, including modification or deletion. Rather than allowing a request to take effect immediately when it is initiated, the system can require review and approval by an authorized approver. This establishes separation of duties: a user who operates workloads or manages protection can initiate a request, while a separately authorized person validates it before the recovery data is changed or removed.

This control is particularly valuable for business-continuity recovery points. Snapshots are often the fastest recovery mechanism following accidental changes, operational mistakes, or a security incident. Requiring approval for destructive operations reduces the risk that a compromised administrative account, malicious insider, or inadvertent action can remove the snapshots required for recovery. It also provides an auditable governance process around actions that affect protected data. Approval policies should be combined with appropriately scoped Prism roles and trusted approver accounts so that only designated personnel can approve requests. Nutanix documents approval workflows and governance capabilities in the Prism Central and NCM documentation: Nutanix Prism Central Approvals and Nutanix Cloud Manager Self-Service Approvals.