

Exam4Future

Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI) 5.20

Nutanix NCM-MCI-5.20

Version Demo

Total Demo Questions: 18

Total Premium Questions: 186

Buy Premium PDF

<https://exam4future.com>

support@exam4future.com

exam4future.com

support@exam4future.com

QUESTION NO: 1

An administrator is using Nutanix Move to migrate a Windows VM from ESXi to AHV. Automated guest preparation is failing with an error:

Account has UAC enabled error

The administrator is using the local built-in administrator account for the Windows VM Admin approval mode must remain enabled.

What should the administrator do to continue with the migration with Nutanix Move?

- A. Use a Domain Admin account for the Windows VM
- B. Reboot the Windows VM and try the migration again
- C. Follow manual VM preparation guidelines
- D. Place Windows VM in Maintenance Mode

ANSWER: C

Explanation:

Follow manual VM preparation guidelines is correct because Nutanix Move automated guest preparation needs to make required guest-level configuration changes before cutover. When the Windows VM must retain User Account Control (UAC) Admin Approval Mode, the automated preparation workflow cannot safely apply the changes it requires and reports the UAC-related account error. Manual preparation is the supported path for completing the necessary source-guest prerequisites without requiring Move to alter the mandated UAC policy. The administrator can prepare the guest according to the Move manual preparation procedure, including the applicable Windows configuration and migration prerequisites, then continue the migration workflow using the manually prepared VM. This approach preserves the organization's requirement that Admin Approval Mode remain enabled while allowing Move to proceed with replication and cutover. Nutanix Move documentation describes guest preparation requirements and the manual preparation workflow for cases where automated preparation cannot be used. For background on the Windows policy involved, Microsoft documents that Admin Approval Mode controls how administrator accounts operate under UAC. See the [Nutanix Move User Guide](#) and Microsoft's [User Account Control settings documentation](#).

QUESTION NO: 2

A customer has NearSync configured. When the administrator tries to restore a snapshot from 3 minutes ago, it is not available. The snapshots are happening at 15-minute intervals instead of at the 5-minute configured interval. When the protection domain was initially set up, the snapshots were happening at the expected 5-minute interval.

What should the customer do to enable the snapshots to happen at the expected interval?

- A. Utilize Metro Availability to meet this requirement
- B. Change the protection domain to use Async DR
- C. Fix a connectivity issue because the protection domain reverted to Async
- D. Configure the protection domain to take snapshots on 15-minute intervals

ANSWER: C

Explanation:

Fix a connectivity issue because the protection domain reverted to Async. NearSync replication depends on healthy, continuous communication between the source and target clusters so that the frequent replication cycles required by the configured recovery point objective can complete. If NearSync cannot maintain the required intercluster connectivity, Nutanix

AHV Turbo improves virtual-machine I/O efficiency by optimizing the AHV data path and reducing overhead associated with handling I/O. To realize the greatest benefit, the guest workload must be capable of generating concurrent work. Therefore, **Workloads are multi-threaded.** is required: a multithreaded application can issue and process I/O in parallel rather than being constrained to one execution thread. This parallelism gives the optimized Turbo path sufficient concurrency to improve throughput and responsiveness under load.

Likewise, **VMs have more than one vCPU.** is required because the guest needs multiple virtual CPUs on which its concurrent application threads and I/O-related processing can execute. A multithreaded workload paired with multiple vCPUs can use CPU scheduling capacity more effectively and avoid restricting the workload to a single virtual execution context. Together, workload-level concurrency and adequate VM CPU configuration allow AHV Turbo's I/O optimizations to deliver greater performance. Nutanix describes AHV as its integrated virtualization platform and documents AHV configuration and performance capabilities in its [AHV overview](#) and the [Nutanix documentation portal](#).

QUESTION NO: 4

An administrator needs to investigate an AHV VM that is intermittently unable to reach a service on another subnet. The administrator wants to determine whether the issue is caused by the guest configuration or by the network path.

Which two actions should the administrator perform? (Choose two.)

- A. Verify the VM guest default gateway and route table.
- B. Review the AHV network and VLAN configuration.
- C. Restart the Controller VM on the VM host.
- D. Move the VM disks to a different storage container.

ANSWER: A B

Explanation:

Verify the VM guest default gateway and route table and **review the AHV network and VLAN configuration** are correct. The guest route table determines the interface and gateway used for traffic destined to remote subnets. An incorrect default gateway, missing route, or route with an unsuitable metric can cause intermittent or asymmetric communication.

The administrator must also confirm that the AHV network assigned to the VM has the correct VLAN and that the VLAN is supported through the host uplinks and physical network. Reviewing both the guest routing configuration and the AHV/physical network path separates guest-level routing faults from Layer 2 or Layer 3 infrastructure issues. Restarting the CVM or changing the VM storage container does not diagnose intersubnet connectivity. See the [Nutanix Documentation portal](#) for AHV networking administration guidance.

QUESTION NO: 5

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 `ntpq -pn`.
- D. Restart the chronos service on all CVMs.

ANSWER: A C

Explanation:

Confirm that the NTP servers are reachable from the CVMs because NTP synchronization requires each Controller VM to have network connectivity to the configured time sources. Connectivity validation should include confirming the appropriate network route, DNS resolution when hostnames are used, and that UDP port 123 is not blocked between the CVMs and the NTP servers. A shared UTC time zone and matching CVM clocks do not by themselves prove that the CVMs have an active, usable NTP peer relationship.

On a CVM, run the command `ntpq -pn` to inspect the NTP daemon's peer associations directly. This command displays the configured peers, their reachability status, delay, offset, jitter, and the currently selected synchronization source. It is therefore the key diagnostic for determining whether a server is being contacted and whether it is eligible to synchronize time. Nutanix's NTP troubleshooting guidance specifically uses `ntpq -pn` to review CVM NTP peer status. See the [Nutanix Knowledge Base guidance for reviewing ntpq output](#) and the [NTP ntpq command documentation](#).

QUESTION NO: 6

An administrator is troubleshooting an AHV VM that cannot communicate with workloads on a newly created VLAN. The VM has the correct IP address and subnet mask.

Which two items should the administrator verify? (Choose two.)

- A. The VLAN is allowed on the physical switch uplinks
- B. The AHV network is mapped to the correct VLAN ID
- C. The CVM IP address matches the VM IP address
- D. Deduplication is enabled on the VM storage container

ANSWER: A B

Explanation:

The VLAN is allowed on the physical switch uplinks and **the AHV network is mapped to the correct VLAN ID** are correct. An AHV network definition associates VM traffic with a VLAN ID, but the physical switch ports connected to the AHV host uplinks must also be configured to carry that VLAN. If the VLAN is missing from the switch trunk or the network definition uses an incorrect VLAN ID, the VM can have valid guest IP settings while its traffic cannot reach the intended Layer 2 network.

Changing the CVM IP address is unrelated to VM VLAN connectivity. Creating a protection domain affects data protection rather than networking. Enabling deduplication changes storage optimization and does not alter VLAN tagging or switch trunk configuration. Nutanix documents AHV VLAN and uplink configuration in the [Nutanix AHV Administration documentation](#).

QUESTION NO: 7 - (HOTSPOT)

An administrator is configuring an All Flash cluster to support the following workloads:

- Full clone virtual desktops
- Hadoop
- Write heavy databases

Which three correct configuration options should the administrator use to meet the item requirements? (Choose three.)

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.

Procedure

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

Correct/Incorrect

Select

Select

Correct

Incorrect

One container with compression and Erasure Coding enabled for Hadoop

Select

Select

Correct

Incorrect

One container with compression enabled for heavy databases

Select

Select

Correct

Incorrect

One container with deduplication enabled for full clone virtual desktops

Select

Select

Correct

Incorrect

One container with Erasure Coding enabled for Hadoop

Select

Select

Correct

Incorrect

One container with deduplication enabled for Hadoop

Select

Select

Correct

Incorrect

ANSWER:

Procedure

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

One container with compression and Erasure Coding enabled for Hadoop

One container with compression enabled for heavy databases

One container with deduplication enabled for full clone virtual desktops

One container with Erasure Coding enabled for Hadoop

One container with deduplication enabled for Hadoop

Correct/Incorrect

Select	Select	Correct	Incorrect
Select	Select	Correct	Incorrect
Select	Select	Correct	Incorrect
Select	Select	Correct	Incorrect
Select	Select	Correct	Incorrect
Select	Select	Correct	Incorrect

Explanation:

For an All Flash Nutanix cluster, storage efficiency settings should be selected per container according to the workload's data characteristics and access pattern. Full-clone virtual desktops are a strong fit for deduplication because many desktops contain the same operating-system blocks, application binaries, and common image data. Enabling deduplication in the desktop container eliminates repeated copies of those shared blocks and provides meaningful capacity savings while preserving the workload-specific container design. Nutanix documents deduplication as a data-reduction technology intended to reduce duplicated data in storage; see the [Prism Web Console Guide: Data Reduction](#).

Hadoop stores large distributed data sets and is an appropriate use case for Erasure Coding when capacity efficiency is required. Erasure Coding protects data using parity information instead of retaining full mirrored copies, reducing storage overhead for suitable Hadoop data. Creating a dedicated Hadoop container with Erasure Coding enabled aligns the data-protection scheme with this distributed-data workload. Nutanix explains the capacity-saving role and requirements of this feature in the [Prism Web Console Guide: Erasure Coding](#).

Write-heavy databases are appropriately assigned to a container with compression enabled. Compression reduces the physical capacity consumed by database data and is configured at the container level, allowing the administrator to apply the efficiency policy specifically where it is needed. Maintaining separate containers also prevents a policy selected for one workload's data profile from being imposed on the others. Nutanix's [compression documentation](#) describes compression as a container-level storage-efficiency capability. Thus, deduplication for the full-clone desktop container, Erasure Coding for the Hadoop container, and compression for the write-heavy database container are the three correct procedures.

QUESTION NO: 8

Which two advanced features Requires you to install Nutnaix Guest Tools NGT ? (Choose Two)

- A. Disaster Recovery of a VM between AHV and ESXi Hypervisors
- B. Self-service file-level recovery from VM snapshots

- C. File recovery from a third-party backup
- D. Live Migration of a VM between AHV & other Hyervisors

ANSWER: A B

Explanation:

NGT is required for **Disaster Recovery of a VM between AHV and ESXi Hypervisors** because cross-hypervisor disaster recovery requires guest-aware functionality to prepare and recover supported workloads when the VM is replicated and restored across the AHV and ESXi environments. Installing NGT places the necessary Nutanix guest components in the virtual machine, allowing Nutanix data-protection workflows to coordinate with the guest operating system where required.

NGT is also required for **Self-service file-level recovery from VM snapshots**. This feature relies on the NGT guest components to expose and mount the appropriate snapshot recovery point within the guest context, so authorized users can browse and restore individual files rather than performing a full virtual-machine restore. This is distinct from basic VM snapshot operations: the guest tooling provides the integration needed for granular, self-service recovery. Nutanix documents NGT as the guest-side component that enables advanced VM data-protection and recovery capabilities, including self-service restore workflows. See the [Nutanix Prism Web Console Guide: Nutanix Guest Tools](#) and the [Nutanix self-service restore documentation](#).

QUESTION NO: 9

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 can 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 appropriate because Prism's Capacity Runway feature evaluates historical resource-consumption trends and forecasts when cluster capacity is likely to be exhausted. The administrator can use this analysis to review available compute and storage headroom, projected growth, and the likely impact of continued VM expansion. This provides a capacity-planning basis for deciding whether the existing cluster has sufficient resources for additional workloads.

“Utilize Optimize Resources for VM efficiency” is also appropriate because Prism can identify VMs whose allocated CPU, memory, or storage resources are not efficiently matched to observed utilization. Reviewing and applying relevant rightsizing recommendations can reclaim capacity from overprovisioned workloads before new VMs are deployed. Combining these actions gives the administrator both a forward-looking cluster-capacity forecast and an opportunity to improve current resource efficiency, producing a more reliable assessment of whether the cluster can accommodate the planned VMs. Nutanix documents Capacity Runway and related capacity planning in the [Prism Central Guide](#), and describes rightsizing and efficiency recommendations within [Nutanix Cloud Platform Operations](#).

QUESTION NO: 10

An administrator is creating an AHV network that will carry VM traffic on VLAN 330. The physical switches use 802.1Q trunking between the AHV hosts and the network infrastructure.

Which two configuration requirements must be met? (Choose two.)

- A. Create the AHV network with VLAN ID 330.
- B. Allow VLAN 330 on the physical-switch ports connected to the AHV hosts.
- C. Configure VLAN 330 as the CVM management subnet.
- D. Create a storage container using VLAN 330.

ANSWER: A B

Explanation:

Create the AHV network with VLAN ID 330 and **allow VLAN 330 on the physical-switch ports connected to the AHV hosts** are correct. The AHV network definition associates the VM network with its VLAN tag. VM traffic leaving an AHV host is tagged for that VLAN when it uses the configured network.

The connected switch ports must also permit VLAN 330 on the trunk. If the VLAN is absent from the allowed VLAN list, tagged VM traffic cannot reach the upstream network even though the network is correctly defined in Prism. Configuring an NTP server or creating an additional storage container does not establish Layer 2 connectivity for the VLAN. See the [Nutanix Documentation portal](#) for AHV network configuration guidance.

QUESTION NO: 11

An administrator reviews the error in the exhibit, the examines the state of the cluster, the administrator finds that the AHV host is up and Functional , but the cluster capacity is low.

What will happen to the VMs on the host with this controller VM?

- A. They will remain on the current host
- B. they will reboot onto other nodes in the cluster
- C. the will shutdown and be orphaned on the host
- D. they will be shutdown unless protected by a Protection Policy

ANSWER: A

Explanation:

They will remain on the current host. A Controller VM (CVM) issue by itself is not equivalent to an AHV host failure. In this scenario, the AHV host is explicitly reported as up and functional, so the user VMs running on that host continue to be powered on and scheduled by the local AHV hypervisor. Nutanix distributes storage services across the cluster's CVMs; therefore, a workload VM does not need its local CVM to be healthy in order to remain running. The VM can continue to access the Nutanix Distributed Storage Fabric through available CVMs in the cluster, assuming the cluster retains quorum and storage availability.

Nutanix HA actions that restart workload VMs elsewhere are associated with an actual host failure or loss of host communication, rather than an isolated CVM condition while the host remains functional. The stated low cluster capacity is operationally important because it can constrain remediation and placement choices, but it does not by itself cause powered-on workload VMs to be evacuated, shut down, or restarted. The appropriate immediate result is continued operation of the VMs on their current AHV host while the administrator investigates and restores the affected CVM. See the [Nutanix documentation portal](#) for AHV and cluster administration guidance and [Nutanix AHV documentation and product resources](#) for AHV availability architecture.

QUESTION NO: 12

Which two access protocols are supported by Files? (Choose two.)

- A. NFS

- B. FCOE
- C. SMB
- D. iSCSI

ANSWER: A C

Explanation:

NFS and SMB are the client access protocols provided by Nutanix Files. NFS enables UNIX and Linux clients to mount file shares using standard Network File System semantics. Nutanix Files supports NFS exports for clients that need shared, POSIX-oriented file access. SMB enables Windows clients and applications to access file shares through the Server Message Block protocol, including use with Microsoft Active Directory environments for identity-based authentication and permissions. Together, these protocols allow the same Nutanix Files platform to serve common enterprise file-sharing requirements across heterogeneous operating-system environments. Nutanix Files is a scale-out file-storage service hosted on Nutanix infrastructure; clients connect to file-server shares using SMB or NFS rather than block-storage protocols. For product documentation describing Files features and supported client protocols, see the [Nutanix Files Overview](#) and the [Nutanix Files product page](#).

QUESTION NO: 13

The networking team makes changes to the Top of Rack switches to which the Nutanix cluster are attached. A few VMs are able to communicate with each other on the same node but are unable to connect to other parts of the network.

What is the likely cause of this issue?

- A. There is a VLAN misconfiguration on the VMs.
- B. The vNIC port is disconnected on the affected VMs.
- C. There is a VLAN misconfiguration on the switch.
- D. Jumbo Frame are misconfigured across ports.

ANSWER: C

Explanation:

There is a VLAN misconfiguration on the switch. Communication between VMs located on the same AHV host can remain local to that host: the Open vSwitch (OVS) bridges and forwards eligible traffic internally without requiring the packet to traverse the physical uplink to the Top-of-Rack switch. Consequently, those VMs may still communicate even when the switch configuration prevents their VLAN traffic from reaching the rest of the network. After a Top-of-Rack switch change, a likely issue is that the VM VLAN is missing from the uplink trunk's allowed VLAN list, assigned incorrectly as a native/access VLAN, or is inconsistently configured between switch ports used by cluster hosts. Traffic from the affected VMs then cannot leave the host correctly or return from external network devices, while local same-host traffic continues to work. Nutanix recommends validating the physical-switch VLAN trunk configuration for every host uplink, including the VLANs used by guest VMs, and ensuring the configuration is consistent across all relevant ports. This behavior is aligned with the AHV networking design in which OVS provides virtual switching on each host and physical uplinks carry external traffic. See the [Nutanix AHV Administration Guide: Networking Overview](#) and [Nutanix AHV Administration Guide: VLAN Configuration](#).

QUESTION NO: 14

A Nutanix cluster is deployed with the following configuration:

- Three four-node blocks (A, B, and C)
- All Flash Redundancy Factor 2

What is the effect of simultaneous disk failures on two nodes located in block A?

- A. VMs are migrated off of the nodes with disk failures.
- B. Each node with a failed disk is marked as degraded.
- C. VM read and write operations continue normally.
- D. VMs on the nodes with failed disks are unable to write data

ANSWER: C

Explanation:

VM read and write operations continue normally. Redundancy Factor 2 provides a higher level of data resiliency by maintaining sufficient protected copies of data and metadata to tolerate the loss of two failure domains while preserving data availability. In this configuration, the cluster has twelve nodes distributed across three separate four-node blocks, giving the distributed storage fabric enough nodes and block-level placement choices to maintain access to surviving replicas when disks fail simultaneously on two nodes in the same block.

Nutanix distributes data across the cluster rather than depending on the local disks of a single node. When a disk becomes unavailable, the surviving replicas continue to service I/O. The system identifies the lost storage and initiates recovery or re-protection activity as capacity and resources permit, while the redundant copies satisfy guest read and write requests. RF2 is therefore designed for continued service during this two-node fault scenario, assuming the cluster was healthy and had adequate available capacity before the failures. Nutanix distinguishes Redundancy Factor, which controls tolerated storage failures, from replication settings used for other data-protection purposes; see [Nutanix Community: Redundancy Factor vs. Replication Factor](#) and [Nutanix Data Resiliency](#).

QUESTION NO: 15

An Administrator Wants a VM to Communicate with two Separate VLANs.

What steps should the administrator complete?

- A. Add a single NIC to the VM and Configure it for both VLANs
- B. Add two NICs to the VM and Configure one for each VLAN
- C. Turn on IPAM for auto configuration of VLANs
- D. Configure a security Policy to assign VLANs to the VM

ANSWER: B

Explanation:

Add two NICs to the VM and Configure one for each VLAN is correct. In Nutanix AHV, a virtual NIC connects a guest VM to a single configured virtual network/VLAN. To give one VM Layer 2 connectivity to two distinct VLANs, the administrator adds two virtual NICs and selects the appropriate AHV network for each NIC. The guest operating system then detects two network interfaces, each of which can receive an IP address and subnet configuration appropriate to its connected VLAN.

This design gives the VM a separate attachment point for each broadcast domain and allows it to communicate with hosts or services on both networks. The administrator should also configure the guest's IP addressing, routes, and—where applicable—DNS settings so that traffic uses the intended interface. For example, a VM might use one NIC for an application VLAN and another NIC for a management, backup, or replication VLAN. Nutanix networking documentation describes attaching VM NICs to configured AHV networks and managing VM network connectivity through Prism. See the [Nutanix Documentation portal](#) and the [Nutanix Support Portal documentation catalog](#) for the applicable AHV Administration Guide for the deployed AOS release.

QUESTION NO: 16

An administrator needs to install prism central for their 2550 multi-cluster environment.

How should they deploy the application?

- A. Deploy a large 3-VM Prism Central instance
- B. Deploy a large 1-VM Prism Central instance.
- C. Deploy a small 3-VM Prism Central instance.
- D. Deploy a small 1-VM Prism Central instance.

ANSWER: A

Explanation:

Deploy a large 3-VM Prism Central instance. Prism Central sizing must be selected according to the scale of the managed environment, including the number of managed entities and the operational load generated by multi-cluster monitoring, inventory, alerts, analytics, and lifecycle workflows. A 2,550-scale multi-cluster deployment requires the Large Prism Central configuration rather than a small deployment. The Large configuration uses a three-VM Prism Central scale-out deployment, distributing Prism Central services across the controller VMs and providing the capacity and resiliency needed for a production environment of this size.

Deploying Prism Central as three VMs also aligns with Nutanix high-availability design: Prism Central services can continue operating through the loss of an individual Prism Central VM, subject to the remaining healthy quorum and infrastructure conditions. Administrators should confirm that the hosting cluster has sufficient resources for the Prism Central VMs and follow the version-specific sizing guidance before deployment. Nutanix documents Prism Central deployment and scale-out configuration in the [Prism Central Guide](#); the associated requirements and supported deployment considerations are available in the [Prism Central deployment documentation](#).

QUESTION NO: 17

An administrator has two clusters: Site A and Site B.

Async DR is configured between the clusters as follows:

- Site A is the Primary Site
- Site B is the Secondary inactive Site, no VMs running
- Site A is fully in Production
- Site B used as Backup/Disaster Recovery

The customer wants to perform a planned failover between the sites.

How should the administrator accomplish this?

- A. Click Activate on the Secondary site within the Data Protection Dashboard
- B. Click Activate on the Secondary site and run deactivate NCLI command on Primary site
- C. Choose Migrate on the Secondary site within the Data Protection Dashboard
- D. Choose Migrate on the Primary site within the Data Protection Dashboard

ANSWER: D

Explanation:

Choose Migrate on the Primary site within the Data Protection Dashboard. In Nutanix asynchronous disaster recovery, migration is the supported workflow for a planned transition of workloads from the active primary site to the recovery site. Because the primary cluster is available and production is operating normally, it can coordinate the final replication activity before the role transition. This helps ensure that the most current protected data is transferred to Site B before the protected virtual machines are started there.

The migration operation is initiated from the current primary cluster because that cluster owns the active protection domain and the replication relationship. Nutanix then performs the controlled role reversal: Site B becomes the active location for the protected workloads, while the prior primary can become the recovery target after the relationship is reversed. This process is intended for maintenance events, datacenter moves, or other planned DR exercises where both sites remain reachable and operational.

Nutanix documents the migration workflow as part of protection-domain failover operations in the [Prism Web Console Guide](#). See also Nutanix's [Disaster Recovery overview](#) for the role of planned migration in DR operations.

QUESTION NO: 18

An administrator is managing multiple Windows and Linux VMs connected to Nutanix Volumes. The Linux VMs are experiencing intermittent connectivity issues. The Windows VMs do not experience the same issue.

Which option should the administrator use to resolve this issue?

- A. Utilize Jumbo Frames
- B. Add additional disks to the Volume Group
- C. Set the SCSI timeout value to 60
- D. Utilize a separate subnet for the Linux VMs

ANSWER: C

Explanation:

Set the SCSI timeout value to 60. Nutanix Volumes present Volume Group disks to guest VMs as SCSI-attached block devices. During a brief storage-path interruption, controller event, or an operation that temporarily delays I/O completion, the Linux guest OS decides whether to continue waiting for a response based on its configured SCSI command timeout. A timeout that is too short can cause the Linux guest to declare the device unavailable even though the Nutanix volume and its connection recover shortly afterward. This behavior can appear intermittent because it depends on the duration of the transient delay and the I/O activity at that moment.

Configuring the Linux guest SCSI timeout to 60 seconds gives the storage path sufficient time to recover and complete outstanding I/O without the operating system prematurely failing the attached Volume Group disk. The setting should be applied persistently for the relevant SCSI devices so it survives reboots and device rediscovery. This is especially important for application disks, where an unnecessary device timeout can cascade into filesystem or application-level errors. VMware also documents the relationship between guest SCSI timeout behavior and temporary storage disruption in its guidance: [VMware KB 1009465](#). Nutanix Volumes documentation is available at [Nutanix Volumes Guide](#).