


## Get Real NCA-6.5 Questions Pass Nutanix Certification Exams Easily [Q42-Q63]



## Get Real NCA-6.5 Questions Pass Nutanix Certification Exams Easily [Q42-Q63]

Get Real NCA-6.5 Questions Pass Nutanix Certification Exams Easily  
NCA-6.5 Dumps are Available for Instant Access

**NO.42** What requirement needs to be considered before configuring Replication Factor of 1 (RF1)?

- \* Application resiliency
- \* Hypervisor
- \* RAID configuration
- \* Node count

To configure Replication Factor 1 (RF1), the cluster must have at least three nodes. In RF1, data is replicated to two other nodes within the cluster, so having at least three nodes is necessary to ensure that the data is replicated. Moreover, when you configure RF1, the cluster must have a minimum of three nodes to ensure that the data is accessible even if one of the nodes goes down.

Reference:

[https://portal.nutanix.com/page/documents/details?targetId=Prism-Admin-Guide-v5\\_6:pr-data-protection-replication-factor-c.html](https://portal.nutanix.com/page/documents/details?targetId=Prism-Admin-Guide-v5_6:pr-data-protection-replication-factor-c.html)

**NO.43** Refer to the exhibit.



The screenshot shows a window titled "Data Resiliency Status" with a close button (X) and a help button (?). Below the title bar, it says "FAULT DOMAIN TYPE: HOST". The main content is a table with three columns: "COMPONENT", "FAILURES TOLERABLE", and "MESSAGE".

COMPONENT	FAILURES TOLERABLE	MESSAGE
Static Configuration	1	
ZooKeeper	0	Desired fault tolerance for Zookeeper is 1 but we can tolerate only 0 node failure(s)
Stargate Health	0	Based on unhealthy node list (5,) the cluster can tolerate a maximum of 0 node failure(s)
Oplog	1	
Metadata	0	Metadata ring partitions with nodes: 192.168.1.33, 192.168.1.31, 192.168.1.32 are not fault tolerant.
Extent Groups	1	
Erasur Code	1	

An administrator notices that the Data Resiliency Statue of some components has been reduced to 0. making it impossible to support a node failure in the cluster.

What type of failure is most likely the cause of this issue?

- \* Block
- \* Disk
- \* CVM
- \* Nods

**NO.44** Which account is recommended for performing tasks and operations on the CVM via SSH?

- \* administrator
- \* admin
- \* nutanix
- \* root

**NO.45** An administrator receives an alert that a node has failed within a Nutanix AHV-based 10-node cluster. Before the failure, the cluster CPU and memory utilization was around 50%.

What actions will the cluster automatically take?

- \* VMs will be unavailable until the affected host is fully functional.
- \* VMs will migrate to other nodes in the cluster with no user impact.
- \* The cluster will enter read-only mode and the VMs will be powered down to preserve data.
- \* All HA-protected VMs will be automatically restarted on other nodes in the cluster.

Nutanix supports high availability (HA) for VMs by automatically restarting them on other nodes in the cluster in case of a node failure. The cluster CPU and memory utilization will increase temporarily until the failed node is recovered, but there will be no user impact. Reference: <https://www.nutanix.com/content/dam/nutanix/resources/misc/ebg-nca-6-5.pdf> (page 28)

**NO.46** An administrator is unable to discover a Nutanix-hosted iSCSI block device allocated to a bare-metal host.

What is causing this issue?

- \* Jumbo frames has not been configured on the switch
- \* External client access has not been enabled for the host
- \* Network segmentation has not been enabled
- \* External host IP address is not in the Filesystem Whitelist

Nutanix Volumes uses a Filesystem Whitelist to control which external hosts can access the iSCSI block devices allocated to them. If the external host IP address is not in the Filesystem Whitelist, it will not be able to discover or connect to the Nutanix-hosted iSCSI block device. Jumbo frames, external client access, and network segmentation are not required for Nutanix Volumes to work. Reference: <https://www.nutanix.com/content/dam/nutanix/resources/datasheets/ds-nutanix-acropolis-block-services.pdf> (page 4) <https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2049-Nutanix-Volumes:BP-2049-Nutanix-Volumes> (page 7)

**NO.47** What happens to images that are imported into Prism Central?

- \* Copied to the Prism Central image repository
- \* Remain in the source cluster after import
- \* Copied to all clusters managed by Prism Central
- \* Deleted from the source cluster after import

The Prism Central image repository is a centralized location for storing virtual machine images that can be used to deploy VMs to any cluster managed by Prism Central. When an image is imported into Prism Central, a copy of the image is made and stored in the repository, so that it can be easily accessed and deployed to any cluster as needed.

The original image file typically remains in the source cluster after it is imported into Prism Central, unless it is explicitly deleted. Similarly, the image copy in the Prism Central repository is not automatically copied to all clusters managed by Prism Central, but can be used to deploy new VMs on any cluster as needed.

**NO.48** Which Nutanix storage efficiency feature is suitable for nearly every workload?

- \* Erasure Coding
- \* Deduplication
- \* Compression
- \* Thick Provisioning

Compression is a Nutanix storage efficiency feature that is suitable for nearly every workload, as it reduces the amount of data stored on disk without impacting performance. Erasure Coding, Deduplication, and Thick Provisioning are other storage efficiency features, but they have different trade-offs and use cases. Reference:

<https://www.nutanix.com/content/dam/nutanix/resources/misc/ebg-nca-6-5.pdf> (page 30)

**NO.49** What is the purpose of Discoveries within the Nutanix Support Portal?

- \* To identify clusters that are affected by known issues.
- \* To provide an audit trail of cluster logins.
- \* To identify VMs that require OS updates or patching.
- \* To provide forecasting of cluster resource utilization.

Discoveries within the Nutanix Support Portal are used to identify clusters that are affected by known issues. This can help Nutanix Support quickly identify and address any potential problems that may be caused by a specific issue. Reference:

[https://portal.nutanix.com/page/documents/details?targetId=Support-Portal-User-Guide-v2\\_0:sp-discoveries-c.html](https://portal.nutanix.com/page/documents/details?targetId=Support-Portal-User-Guide-v2_0:sp-discoveries-c.html)

**NO.50** Which tool provides diagnostic system data for Nutanix to deliver proactive support? (Choose Correct answer and give explanation/References from Nutanix Certified Associate (NCA) v6.5 official study resources)

- \* NCC
- \* Cluster Heart

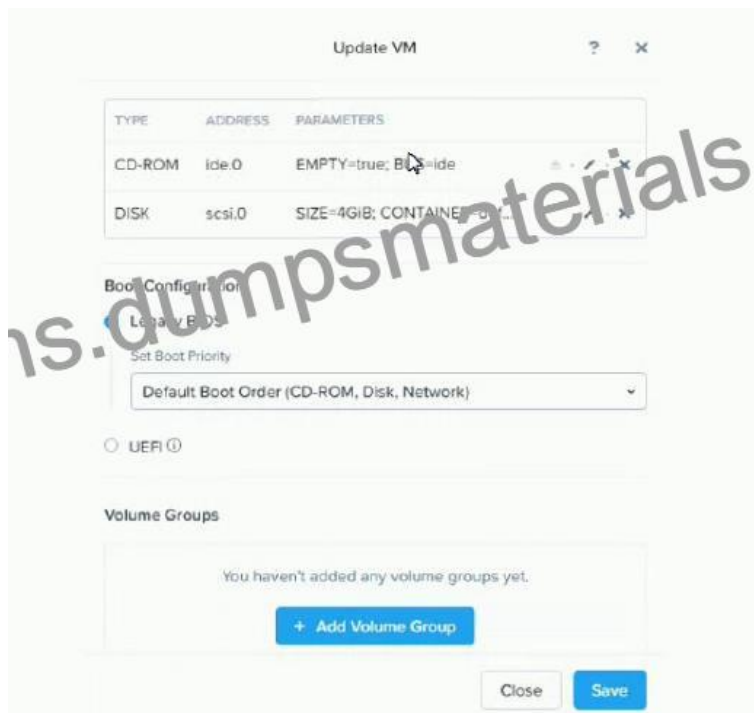
- \* Pulse
- \* Cluster health

Pulse is a built-in tool that periodically collects and sends diagnostic data to Nutanix support for proactive monitoring and troubleshooting. This data includes cluster configuration, performance metrics, and error logs, among other things. Nutanix support can use this data to identify and resolve issues before they become critical, and to provide recommendations for optimizing cluster performance.

NCC (Nutanix Cluster Check) is another tool provided by Nutanix, but it is used for performing manual health checks and diagnostics on the cluster, rather than providing proactive support.

Cluster Heart and Cluster Health are not tools provided by Nutanix, and are not relevant to this question.

**NO.51** Refer to the exhibit.



A team member has built a CentOS VM image for use as a template.

An administrator deploys a new VM using the CentOS image as the 05 drive. No other drives are presented to the VM.

Following the deployment, the administrator attempts to boot the VM, but is unsuccessful.

What could be causing this issue?

- \* The VMS Boot Priority must reference disk first before CD-ROW
- \* The VM's Disk (scsi.0) setting must be set to make this disk bootable
- \* No Volume Groups have been created for the VM
- \* The CentOS VM image was built with UEFI:

**NO.52** Which product or feature is most suitable for deploying a Microsoft SQL cluster in a Nutanix environment?

- \* Volumes

- \* Objects
- \* NFS shares
- \* Files

In a Nutanix environment, the product or feature most suitable for deploying a Microsoft SQL cluster is Nutanix Objects. Nutanix Objects is a S3-compatible object storage solution that provides a highly scalable and durable storage for unstructured data. It is built on top of the Nutanix Distributed File System (NDFS) and provides a high-performance, scalable, and highly available storage solution.

One of the key features of Nutanix Objects is its ability to support Microsoft SQL Server Always On Availability Groups (AAG) on S3 object storage. This allows you to deploy a highly available SQL cluster on Nutanix, with the ability to failover SQL services to a different node in the event of a failure, and provide high availability and disaster recovery for SQL databases.

**NO.53** When should Pulse be disabled?

- \* When the cluster is connected to the Internet with less than 1GbE.
- \* When the cluster is deployed in 3 dark-site.
- \* Before any cluster maintenance activities requiring node reboots.
- \* Only under the guidance of Nutanix Support.

Pulse should only be disabled under the guidance of Nutanix Support and not for any other reason. Pulse is a feature of the Nutanix cluster that enables the cluster to report back health and performance data to Nutanix for troubleshooting and performance optimization purposes. Disabling Pulse without the guidance of Nutanix Support is not recommended, as it could lead to potential issues with the cluster. Reference:

[https://portal.nutanix.com/page/documents/details?targetId=Platform-Administrator-Guide-v5\\_7:pag-pulse-c.html](https://portal.nutanix.com/page/documents/details?targetId=Platform-Administrator-Guide-v5_7:pag-pulse-c.html)

**NO.54** Refer to the exhibit.

Overview • Diagram • Table

+ Expand Cluster

Repair Host Bo



After an abrupt power outage, an administrator receives a number of alerts indicating disks are failed. Prism Element shows a large number of disks offline as seen in the diagram. This single block, four node cluster is configured FT1 with an RF2 container holding all user data.

What number of drive failures on different domains (node, block, or rack) could this configuration have tolerated prior to seeing data unavailability?

- \* 1
- \* 4
- \* 3
- \* 2

**NO.55** Which Windows technology does AHV support to prevent authentication theft?

- \* Endpoint Manager
- \* Azure AD Multi-Factor
- \* Windows Defender Credential Guard
- \* Active Directory federation Services

**NO.56** Which storage protocol is used to attach a container to an ESXi host?

- \* CIFS
- \* FCoE
- \* SMB
- \* NFS

**NO.57** A single host in a four-node AHV-based cluster experiences a complete network failure.

If more than enough resources exist in the cluster to tolerate a node failure, what happens to the user VMs running on that host?

- \* VMs remain down until network connectivity is restored.
- \* VMs stay up, but are unreachable until connectivity is restored.
- \* VMs automatically restart on the three remaining hosts.
- \* VMs are live-migrated to the three remaining hosts.

If a single host in an AHV-based cluster experiences a complete network failure, and if there are enough resources available in the cluster to tolerate a node failure, then the user VMs running on the failed host will be automatically live-migrated to the remaining three hosts in the cluster. This ensures that the VMs stay up and running even if one of the hosts fails. Reference:

[https://portal.nutanix.com/page/documents/details?targetId=Platform-Administrator-Guide-v5\\_7:pag-vm-migrate-c.html](https://portal.nutanix.com/page/documents/details?targetId=Platform-Administrator-Guide-v5_7:pag-vm-migrate-c.html)

**NO.58** Which feature enables Image Placement Policies to be mapped to target clusters?

- \* YAML
- \* Labels
- \* JSON
- \* Categories

Image Placement Policies can be mapped to target clusters using labels. Labels are key-value pairs that can be used to associate arbitrary metadata with clusters and VMs. Labels can be used to define Image Placement Policies, which specify which clusters can be used for a given image. Reference:

[https://portal.nutanix.com/page/documents/details?targetId=Prism-Admin-Guide-v5\\_6:pr-images-placement-policies-c.html](https://portal.nutanix.com/page/documents/details?targetId=Prism-Admin-Guide-v5_6:pr-images-placement-policies-c.html)

**NO.59** Which AHV feature proactively monitors a Nutanix cluster for compute and storage I/O contention or hot-spots over a period of time?

- \* Genesis
- \* ADS
- \* Prism
- \* Cluster Maintenance Utility

ADS is a feature of Nutanix AHV (Acropolis Hypervisor) that proactively monitors a Nutanix cluster for compute and storage I/O contention or hotspots over a period of time [1]. It provides dynamic and intelligent scheduling of workloads on the cluster for optimal performance and utilization.

**NO.60** User are complaining that web application hosted on a Nutanix cluster are running slow. After reviewing the performance metric, it is determined that the CPU Ready time is high in the cluster.

What entity is impacted by CPU ready time?

- \* Opllog
- \* NIC
- \* vDisk
- \* VM

The opllog is a critical component of the Nutanix distributed storage fabric and can be impacted by high CPU Ready time. When the CPU Ready time is high, there can be delays in writing data to the opllog, resulting in slower overall cluster performance. Additionally, high CPU Ready time can also lead to increased write amplification, further reducing I/O performance.

**NO.61** What is Nutanix's scale-out storage technology that makes its hyperconverged platform possible?

- \* Distributed File Services
- \* Distributed Storage Fabric
- \* HDD/SDD Storage Tiers
- \* Hot swappable NVMe/SSD

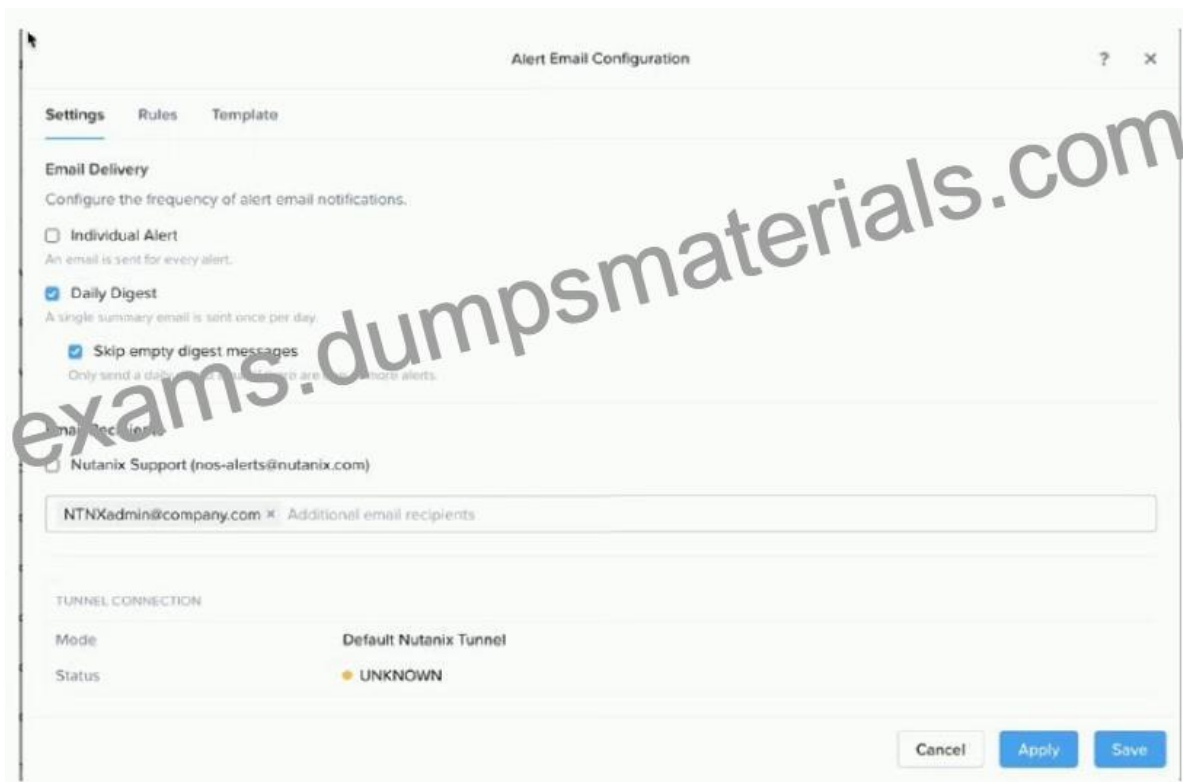
Nutanix's scale-out storage technology, the Distributed Storage Fabric, makes its hyperconverged platform possible. It provides a single, unified layer of storage to multiple nodes in a cluster, allowing data to be stored, accessed, and processed in a distributed, fault-tolerant manner. Reference: <https://www.nutanix.com/what-is-nutanix/distributed-storage-fabric/>

**NO.62** What data is stored as a file on storage devices owned by a CVM?

- \* Storage Pool
- \* vDisk
- \* Extent Group
- \* Container

A vDisk is a file that is stored on storage devices owned by a CVM. A vDisk represents a virtual disk that is attached to a VM or a CVM. A Storage Pool is a logical grouping of physical disks in a cluster. An Extent Group is a collection of data blocks that belong to a vDisk. A Container is a logical grouping of vDisks that share the same storage policies and settings. Reference: <https://www.nutanix.com/content/dam/nutanix/resources/misc/ebg-nca-6-5.pdf> (pages 31-32)

**NO.63** Refer to the exhibit.



An administrator has configured Prism Central to email daily digests of alerts on a cluster. After a week, the administrator notices that digests are not being received.

What is the most likely cause of the issue?

- \* The recipient address is not registered with Nutanix.
- \* Nutanix support is not enabled as a recipient.
- \* The tunnel connection has not been enabled.
- \* The SMTP server is not configured properly.

The NCA-6.5 exam covers a wide range of topics, including the basics of Nutanix solutions, installation and configuration, cluster management, and troubleshooting. NCA-6.5 exam is designed to test the candidate's knowledge and skills in all areas of Nutanix solutions, from setting up a new cluster to troubleshooting complex issues. NCA-6.5 exam also covers advanced topics such as data protection, disaster recovery, and automation. Passing the NCA-6.5 exam requires a deep understanding of Nutanix solutions and the ability to apply this knowledge to real-world scenarios.

**Get Instant Access REAL NCA-6.5 DUMP Pass Your Exam Easily:**



<https://www.dumpsmaterials.com/NCA-6.5-real-torrent.html>