

Nutanix Certified Professional - Multi cloud Infrastructure (NCP-5.20)

Nutanix NCP-MCI-5.20

Version Demo

Total Demo Questions: 29

Total Premium Questions: 299

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Describe the Nutanix Cloud Platform	49
Topic 2, Plan a Nutanix Cloud Platform Deployment	21
Topic 3, Deploy a Nutanix Cloud Platform	14
Topic 4, Configure and Manage a Nutanix Cloud Platform	174
Topic 5, Troubleshoot a Nutanix Cloud Platform	41
Total	299

QUESTION NO: 1

A system 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. Enable Deduplication and Ensure Coding
- B. Utilize Optimize Resources for VM efficiency
- C. Determine utilization with Cluster Runway
- D. Perform an inventory with Life Cycle Management

ANSWER: B C

Explanation:

Utilize Optimize Resources for VM efficiency is appropriate because Prism's resource-optimization capabilities help an administrator identify VMs whose allocated CPU, memory, or other resources do not align with observed utilization. Reviewing these recommendations helps reclaim or right-size resources and provides a more realistic view of the capacity available for additional workloads. This evaluation should occur before making a placement decision, since configured resources alone may not represent the cluster's actual needs.

Determine utilization with Cluster Runway is also appropriate because Cluster Runway uses historical cluster consumption trends to estimate when CPU, memory, and storage capacity will be exhausted. This allows the administrator to assess whether the expected requirements of the new VMs fit within available capacity for the intended planning period, rather than relying only on a current utilization snapshot. Together, optimization recommendations and runway forecasting provide both an opportunity to recover inefficiently allocated resources and a forward-looking capacity assessment. Nutanix describes Prism as providing infrastructure monitoring, analytics, and planning capabilities at [Nutanix Prism](#). Additional Nutanix documentation is available through the [Nutanix Support Portal documentation library](#).

QUESTION NO: 2

An administrator installs Nutanix-specific PowerShell cmdlets. The cmdlets fail to run.

Which two requirements should the administrator verify first? (Choose two)

- A. Windows 2016 r2 or later is installed
- B. The cmdlets have the proper permissions
- C. PowerShell 2.0 or later is installed
- D. .Net Framework 4.0 is installed
- E. Python 2.7 or greater is installed

ANSWER: C D

Explanation:

The administrator should first verify that **PowerShell 2.0 or later is installed** and that **.Net Framework 4.0 is installed**. Nutanix PowerShell cmdlets depend on the Windows PowerShell runtime to load the module, interpret cmdlet syntax, establish sessions, and execute management operations against a Nutanix environment. If a supported PowerShell version is absent, the module cannot load or its cmdlets cannot be invoked successfully. The cmdlets also require the .NET Framework runtime because the module and its supporting libraries use .NET assemblies for connectivity, API communication, and command execution. These are foundational host prerequisites, so they should be confirmed before investigating cluster access, authorization, or command-specific behavior. After confirming both prerequisites, an administrator can verify that the module is installed and imported in the active PowerShell session and then test connectivity to the intended Prism endpoint. Nutanix continues to document PowerShell-based automation resources and examples

through its developer community: [Nutanix.dev PowerShell resources](#). For the Microsoft PowerShell platform and module-management guidance, see [Microsoft PowerShell documentation](#).

QUESTION NO: 3

How will an HDD failure affect VMs with data on the failed device?

- A. The VMs will crash, and will be restarted once the failed HDD has been replaced and the data has been restored.
- B. A live migration will be initiated, moving the affected VMs to a host that contains the replica data.
- C. The VMs will remain operational on that host and continue to function normally with no noticeable impact
- D. An HA event will occur, causing the affected VMs to restart on a node that contains the replica data.

ANSWER: C

Explanation:

The VMs will remain operational on that host and continue to function normally with no noticeable impact. Nutanix Distributed Storage Fabric (DSF) protects VM data through replication across failure domains. With the standard replication factor of two, each data extent has another replica on a different node, allowing the system to continue servicing reads and writes when an individual HDD becomes unavailable. The VM does not depend on a single local HDD for continued execution: its compute process remains running, while the storage layer transparently uses the available replica. Nutanix automatically detects the failed drive and begins restoring the configured redundancy by replicating data to healthy capacity in the cluster. This self-healing behavior is designed to maintain data availability and VM operation during a disk failure, without requiring an HA restart or migration event. The cluster should, however, have sufficient healthy capacity and nodes to satisfy its configured replication-factor requirements. Nutanix describes this distributed, replicated architecture in its [Acropolis infrastructure overview](#) and explains resiliency concepts in the [Prism Web Console Guide: Data Resiliency](#).

QUESTION NO: 4

Which prism dashboard should an administrator use to set the frequency of NCC?

- A. data protection
- B. cluster health
- C. Alerts
- D. Tasks

ANSWER: B

Explanation:

The **cluster health** dashboard is the correct location because Nutanix Cluster Check (NCC) is the cluster-health diagnostic framework used to evaluate configuration, service, hardware, storage, and other operational health conditions. In Prism Element, the Cluster Health area provides the NCC controls, including access to health checks and the scheduling settings used to define how often NCC runs automatically. An administrator can use this area to configure the recurring NCC execution frequency so that cluster health is assessed at an appropriate operational interval, while also reviewing the results produced by those checks. This placement reflects NCC's purpose: it is a proactive health-validation capability rather than a workload, alert-management, task-history, or data-protection function. Regular NCC scheduling helps administrators identify conditions that require attention before they develop into service-impacting issues. Nutanix documents NCC as the health-check utility for Nutanix clusters and describes its integration with Prism's cluster-health workflows. See the official [Nutanix Cluster Check Guide](#) and the [Nutanix documentation portal](#) for the applicable Prism and NCC version documentation.

QUESTION NO: 5

A customer wants to isolate a group of VMs within their Nutanix environment for security reasons. The customer creates a VM with two NICs to act as a firewall and installs the appropriate software and certificates.

However, no one from the outside can access the application.

What is the likely cause of this problem?

- A. A shared volume group must be used by all isolated VMs
- B. More than one NIC cannot be added to a VM
- C. One of the NICs needs to be configured on the internal VLAN
- D. Wireshark is installed on the NAT VM

ANSWER: C

Explanation:

One of the NICs needs to be configured on the internal VLAN because the firewall VM must provide a routed, inspected path between two distinct network segments. In this design, one virtual NIC connects to the external or upstream application-access network, while the second virtual NIC connects to the VLAN used by the isolated application VMs. The firewall software can then receive traffic on the external-facing interface, apply its configured policies and certificates, and forward permitted traffic through its internal-facing interface to the protected workload network. The return path likewise passes back through the firewall, preserving a consistent and controllable security boundary. Simply adding two NICs does not establish this topology: each NIC must be attached to the appropriate Nutanix network/VLAN, and the isolated VMs must use the firewall's internal interface as the relevant gateway or have routes that direct external-bound traffic through it. Nutanix AHV supports attaching VM NICs to configured virtual networks and VLAN-backed segments, which enables this type of virtual appliance deployment. See the [Nutanix AHV Administration Guide networking overview](#) and the [Nutanix AHV product documentation resources](#).

QUESTION NO: 6

On nutanix cluster in which distribution is open ldap authentication management system.?

- A. Red Hat
- B. Ubuntu
- C. Suse
- D. CentOS

ANSWER: D

Explanation:

CentOS is correct. In the Nutanix AOS and Prism architecture covered by the NCP-MCI 5.20 track, the Controller VM (CVM) is built on a customized CentOS-based Linux environment. Nutanix uses services within that CVM environment, including OpenLDAP components, to support cluster-level authentication and authorization functions. OpenLDAP provides the directory-service foundation used for managing local user and group information that Prism and cluster services can use. This is an implementation detail of the Nutanix software stack rather than a choice of guest operating system for user workloads. Administrators normally configure and manage authentication through Prism rather than administering the underlying directory service directly. Knowing that the CVM platform is CentOS-based is useful when identifying Nutanix service locations, understanding platform architecture, and following Nutanix support procedures. Nutanix documentation describes CVMs as the software appliances that provide the cluster's core data and control-plane services, while Prism documentation covers cluster authentication and user-management capabilities. See the [Nutanix Prism Web Console Guide](#) and the [Nutanix Security Guide](#) for the relevant user-management and authentication context.

QUESTION NO: 7

through what interfaces is/are (anti-)affinity rules configured. (Choose three)

- A. affinity :cli
- B. anti-affinity :cli
- C. affinity :prism
- D. anti-affinity :prism

ANSWER: A B C

Explanation:

For the Nutanix release scope covered by this question, affinity and anti-affinity placement controls are available through the command-line interface, while Prism provides the graphical workflow for configuring affinity. The command-line interface is appropriate for both affinity and anti-affinity administration because it exposes the relevant Acropolis command set for defining placement behavior and incorporating it into repeatable operational procedures or automation. This enables administrators to apply rules consistently without requiring manual interaction with the Prism user interface.

Prism also supports configuring affinity through its VM-management workflow. This provides an administrative interface for associating a virtual machine's placement preference with the intended host, making the setting accessible during normal VM operations. These capabilities implement workload-placement requirements while working with AHV scheduling behavior. Nutanix documentation for the applicable product generation is available through the [AHV Administration Guide](#); Nutanix also maintains its official documentation catalog at the [Nutanix Documentation portal](#). Therefore, the supported selections are command-line affinity configuration, command-line anti-affinity configuration, and Prism affinity configuration.

QUESTION NO: 8

Which two methods should an administrator employ to set up the AHV host VLAN in a cluster to follow Nutanix recommendations? (Choose two)

- A. Avoid adding guest VMs to the CVM and host VLAN.
- B. Add all VMs to the default VLAN0 on bridge br1.
- C. Add the CVM and the AHV host to the same VLAN.
- D. Add all VMs to a new VLAN other than VLAN 0.

ANSWER: C D

Explanation:

Nutanix recommends placing the Controller VM (CVM) and its AHV host on the same management VLAN. This provides the required Layer 2 connectivity for host and CVM management communication and gives the cluster a consistent, supportable network foundation. The management VLAN is normally configured as the native, untagged network used by the AHV host and CVM interfaces, with the physical switch ports configured consistently across all nodes.

Guest workloads should be placed on a separate VLAN rather than VLAN 0. Creating and assigning a dedicated, non-default VLAN for guest VMs separates tenant or application traffic from the cluster-management network. This segmentation is a Nutanix networking best practice because it makes traffic policy, addressing, routing, and security controls clearer while preserving the management network for cluster components. AHV uses Open vSwitch bridges to connect host, CVM, and guest traffic, and VLAN-backed networks can be attached to VMs according to the intended workload network design.

For AHV networking concepts and VLAN-backed guest networks, see the [Nutanix AHV Administration Guide](#) and the [Nutanix Documentation portal](#).

QUESTION NO: 9

What are the resource requirements for a single small Prism Central instance?

- A. 4 vCPUs, 16 GB of memory, and 400 GiB of storage
- B. 2 vCPUs, 8 GB of memory, and 500 GiB of storage
- C. 4 vCPUs, 16 GB of memory, and 500 GiB of storage
- D. 8 vCPUs, 32 GB of memory, and 500 GiB of storage

ANSWER: A

Explanation:

A single small Prism Central instance requires 4 vCPUs, 16 GB of memory, and 400 GiB of storage. These are the deployment resources allocated to the Prism Central virtual machine when the Small sizing option is selected during deployment. This size is intended for comparatively small environments and provides the baseline compute, memory, and disk capacity needed for Prism Central services, including its management functions, internal databases, monitoring data, and associated platform processes. Administrators should ensure that the selected Nutanix cluster has sufficient unreserved capacity before deployment, because Prism Central runs as a controller-managed virtual machine and its allocated resources must remain available for reliable operation. Prism Central sizing should also be evaluated against the expected scale of managed clusters, hosts, virtual machines, and enabled services; environments that exceed the supported Small-instance scope should use the appropriate larger deployment size. Nutanix documents the Small Prism Central deployment profile and its resource values in the Prism Central deployment sizing guidance. See the [Nutanix Prism Central Guide: Deployment Sizing](#) and the [Nutanix Prism product overview](#) for official platform context.

QUESTION NO: 10

what information is shared through pulse? multi answer

- A. system level statistics
- B. metadata
- C. identification data
- D. system alerts

ANSWER: A D

Explanation:

Nutanix Pulse shares system-level statistics and system alerts with Nutanix to provide proactive support and operational insight. The telemetry includes health, performance, capacity, and configuration-related statistics collected from the cluster. This information helps Nutanix Support identify trends, diagnose conditions more efficiently, and provide recommendations based on the environment's observed state. Pulse also transmits alert information, allowing support workflows and health monitoring to account for active or recent system events that could affect cluster availability or performance. The service is designed around supportability and product-improvement telemetry rather than sharing customer workload contents. Administrators can review and manage Pulse participation and related data-sharing settings from Prism, in accordance with their organization's support and privacy requirements. Nutanix describes Pulse as a mechanism for securely sending diagnostic and operational information to Nutanix for analysis and proactive support. See the [Nutanix Product Support overview](#) and the [Prism Central documentation for Pulse](#) for details on its support telemetry and configuration.

QUESTION NO: 11

When license needs to be reclaimed from Host: (Choose 2)

- A. While Destroying the Cluster
- B. If you move nodes from one cluster to another , First reclaim the licenses, move the nodes, Then Re-Apply licences
- C. Return licences to your Inventory when you remove one or more nodes from cluster

ANSWER: A B

Explanation:

Licenses should be reclaimed while destroying the cluster because the cluster will no longer consume the assigned entitlement. Reclaiming first returns the license capacity to the organization's available inventory, where it can be assigned to another eligible cluster or retained for later use. This also prevents inactive or decommissioned infrastructure from remaining associated with license consumption in the licensing workflow.

Licenses should also be reclaimed before moving nodes from one cluster to another, then reapplied after the nodes are part of the destination cluster. Nutanix licensing is associated with the cluster's licensed resource requirements rather than simply following hardware automatically between clusters. Reclaiming before the move ensures that the source cluster's assignment is updated appropriately; applying the required license after the move ensures that the destination cluster is correctly licensed for its resulting configuration. Administrators should plan this as part of the node migration or cluster lifecycle procedure and verify available license inventory before making the change. Nutanix documents licensing administration and cluster lifecycle operations through the [Nutanix Support Portal documentation library](#) and provides licensing information at [Nutanix Software Options](#).

QUESTION NO: 12 - (HOTSPOT)

An administrator needs to shut down an AHV cluster to relocate hardware. The administrator upgrades NCC and runs health checks.

Which steps should the administrator perform next?

Item instructions: For each procedure, indicate the order in which that procedure must take place to meet the item requirements.

Procedure

Step

Shut down CVMs

	▼
Step 1	
Step 2	
Step 3	
Step 4	

Shut down Nodes

	▼
Step 1	
Step 2	
Step 3	
Step 4	

Shut down Guest VMs

	▼
Step 1	
Step 2	
Step 3	
Step 4	

Stop the Cluster

	▼
Step 1	
Step 2	
Step 3	
Step 4	

ANSWER:

Procedure

Step

Shut down CVMs

	▼
Step 1	
Step 2	
Step 3	
Step 4	

Shut down Nodes

	▼
Step 1	
Step 2	
Step 3	
Step 4	

Shut down Guest VMs

	▼
Step 1	
Step 2	
Step 3	
Step 4	

Stop the Cluster

	▼
Step 1	
Step 2	
Step 3	
Step 4	

Explanation:

The correct shutdown sequence is to shut down guest VMs first, stop the Nutanix cluster second, shut down the Controller VMs third, and shut down the AHV nodes last. This order preserves the management and storage-control functions needed to perform a controlled shutdown. Guest VMs are powered off first so that application writes are completed and workloads are no longer dependent on the cluster's storage and virtualization services.

After guest workloads are stopped, the cluster is stopped while the Controller VMs remain online. Stopping the cluster cleanly terminates Nutanix distributed services and places the cluster into a controlled state before its controller infrastructure is powered down. The Controller VMs can then be shut down because they are no longer required to coordinate active cluster services. Finally, the physical AHV nodes are shut down only after their Controller VMs have stopped, preventing an abrupt loss of the CVM operating systems and allowing the hardware to be safely relocated.

This sequencing is consistent with Nutanix guidance for planned cluster shutdown operations: protect and stop workload activity, stop cluster services, power off CVMs, then power off hosts. Review the applicable cluster shutdown and startup procedures in the [Nutanix AHV Administration Guide](#) before carrying out the maintenance window.

QUESTION NO: 13

An Engineer logs into a CVM to troubleshoot connectivity loss between two CVMs and find the following interfaces subnets and VLANs configured.

-eth0 – VLAN10-10.1.10.0/24

eth01–No VLAN-192.168.5.0/24

eth2 – LAN20-10.1.20.0/24

Which two Nutanix networking features are being utilized? (Choose two)

- A. Internal vSwitch
- B. IP address management (IPAM)
- C. Network Load balancing
- D. Backplane LAN

ANSWER: A D

Explanation:

Internal vSwitch and **Backplane LAN** are being utilized. Nutanix AHV networking uses Open vSwitch-based virtual switching on each host to provide local Layer 2 connectivity among the host, Controller VM, and guest virtual machines. This internal virtual-switching capability supports the locally attached CVM networking required for Nutanix services while also permitting external networks to be presented through VLAN-backed uplinks. The displayed VLAN-associated networks therefore indicate that virtual switching is participating in the CVM's network connectivity.

The 192.168.5.0/24 network without a VLAN corresponds to the Nutanix backplane network. This is a private, non-routable communication network used for intra-node and CVM-related traffic. The backplane LAN provides an internal path that is distinct from the externally reachable management and VLAN-based guest or service networks. As a result, a configuration that shows both VLAN-connected subnets and an untagged 192.168.5.0/24 internal subnet reflects the combined use of AHV's internal virtual switching and the Nutanix backplane LAN.

For background on AHV's integrated virtualization and networking architecture, see [Nutanix AHV](#) and the [Nutanix Cloud Platform](#).

QUESTION NO: 14

An administrator needs to deploy an application with a large amount of data connected via Nutanix volumes.

Which two actions should the administrator take when designing the Volume Group? (Choose two.)

- A. Distribute workload across multiple virtual disks
- B. Enable RSS (Receive Side Scaling)
- C. Use multiple subnets for iSCSI traffic
- D. Enable thick provisioning on the Volume Group(s)

ANSWER: A B

Explanation:

For an application that requires a large amount of data through Nutanix Volumes, the appropriate design actions are **Distribute workload across multiple virtual disks** and **Enable RSS (Receive Side Scaling)**. A Volume Group can contain multiple virtual disks, allowing application data and I/O to be spread across separate disks rather than concentrating activity on a single virtual disk. This design increases available parallelism for I/O processing and helps the application use storage

resources more effectively as its data set and I/O demand grow. The application operating system can discover and use these disks through the Nutanix iSCSI attachment, subject to its own disk-management and application-layout requirements.

Receive Side Scaling should be enabled on supported guest operating systems and network adapters used for the iSCSI connection. RSS distributes receive-side network processing across multiple CPU cores rather than processing incoming storage traffic on one core. Because Nutanix Volumes uses iSCSI connectivity, this can prevent a single CPU from becoming the limiting factor during sustained or high-throughput storage I/O, helping the guest efficiently process the network traffic associated with the attached Volume Group. Nutanix documents Volume Group configuration and host-attachment considerations in its [Nutanix Volumes Guide](#); RSS behavior and configuration are documented in [Microsoft Receive Side Scaling documentation](#).

QUESTION NO: 15

An administrator finds that guest VMs are offline and storage is unavailable. The administrator needs to know if any process within the CVMs on the nutanix cluster are experiencing crashes.

Which log file should the administrator locate to determine if any processes have crashed?

- A. ERROR logs
- B. INFO logs
- C. FATAL logs
- D. WARN logs

ANSWER: C

Explanation:

FATAL logs are the appropriate source to locate when determining whether a process in a Controller VM has crashed. In Nutanix CVM logging, fatal-level entries represent severe, unrecoverable events that cause a service or process to terminate, abort, or become unable to continue normal operation. Reviewing these records helps an administrator correlate a process failure with the time at which guest VM availability and storage access were lost. This is particularly useful during an outage because CVM services provide the distributed storage and control-plane functions required by the cluster; a fatal event affecting one of these services can be a key indicator for escalation and root-cause analysis. Administrators should collect the relevant fatal log entries, including timestamps, process names, crash signatures, and any associated stack traces, then use that evidence alongside cluster health checks and support-log collection. Nutanix's troubleshooting workflow emphasizes reviewing CVM logs and collecting diagnostic information before engaging support. See the [Nutanix Files and Logs Guide](#) and the [Prism Web Console Guide](#) for Nutanix logging and support-data guidance.

QUESTION NO: 16

What is the function of the virbr0 bridge on AHV?

- A. To carry all traffic between the user VMs and the upstream network.
- B. To carry management and storage communication between user VMs and the CVM.
- C. To carry management and storage communication between user VMs and AHV host.
- D. To carry storage communication between the guest VMs and the CVM
- E. To provide an internal private bridge for management communication between the AHV host and the CVM.

ANSWER: E

Explanation:

The correct function is: "To provide an internal private bridge for management communication between the AHV host and the CVM." On an AHV node, `virbr0` is the private virtual bridge used for host-to-Controller-VM communication. It is not the

uplink-facing bridge used to connect guest virtual machines to the physical network. The Controller VM has a private interface connected to this bridge, and the AHV host also has an interface on the same internal network. This path supports the node's local control and management interactions with its CVM, independently of guest VM network connectivity. Nutanix documents the AHV networking model and distinguishes the private CVM/host networking from the bridges and uplinks used for external VM networking. This separation is important because the CVM is the local Nutanix services controller for the node and requires reliable host-level communication even when guest network configurations vary. See the [Nutanix AHV Administration Guide networking overview](#) and the [Nutanix product documentation portal](#) for AHV network architecture details.

QUESTION NO: 17

Which process should a cluster administrator follow to determine the uplink bond configuration node?

- A. Log into the Nutanix CVM and execute the command `manage_ovs_update_uplinks`
- B. Log into the Nutanix AHV host and execute the command `ovs-appctl bond/show`
- C. Log into the Nutanix AHV host and execute the command `ifconfig br0`
- D. Log into the Nutanix CVM and execute the command `ovs-vsctl show`

ANSWER: B

Explanation:

Logging in to the Nutanix AHV host and executing `ovs-appctl bond/show` is the appropriate process for determining the uplink bond configuration. AHV uses Open vSwitch (OVS) for virtual switching, and an uplink bond is an OVS construct that groups physical NICs to provide network redundancy and, depending on the configured mode, traffic distribution. The `ovs-appctl bond/show` command queries the local OVS daemon directly and displays bond-specific operational details, including the bond name, member interfaces, active slave, bond mode, link status, and other forwarding information. An administrator can optionally specify the bond name, such as `ovs-appctl bond/show br0-up`, when the uplink bond name is known. Because the physical uplink interfaces and the OVS service are local to each AHV node, this check must be performed on the AHV host whose uplink configuration is being examined. Nutanix documents AHV networking as using Open vSwitch, while the Open vSwitch command reference documents `bond/show` as the application-control command for viewing bond status and configuration details. See the [Nutanix AHV Administration Guide](#) and the [Open vSwitch ovs-appctl reference](#).

QUESTION NO: 18

Prism Central will be installed manually on an AHV cluster.

Which three disk images must be downloaded from the portal for the Prism Central VM? (Choose three.)

- A. var
- B. tmp
- C. boot
- D. home
- E. data

ANSWER: C D E

Explanation:

A manually deployed Prism Central VM on AHV requires the **boot**, **home**, and **data** disk images from the Nutanix Portal. These images represent the separate virtual disks that make up the Prism Central appliance. The boot image supplies the VM's bootable operating-system disk and allows the appliance to start. The home image provides the filesystem area used

for Prism Central application components and related persistent configuration. The data image provides the dedicated persistent storage used by Prism Central services and operational data. During manual deployment, the administrator uploads these required QCOW2 images to an AHV image service and attaches them to the newly created Prism Central VM in the prescribed configuration. Obtaining the matching set of images for the selected Prism Central release is important so that the appliance disks and software version are consistent. Nutanix provides Prism Central software downloads and product documentation through the [Nutanix Support Portal](#). For support entitlement, downloads, and documentation access, see [Nutanix Product Support](#).

QUESTION NO: 19

Which three upgrades should an administrator be able to perform using Lifecycle Management? (Choose Three)

- A. AOS
- B. BMC
- C. BIOS
- D. Hypervisor
- E. HBA Firmware

ANSWER: B C E

Explanation:

Lifecycle Management (LCM) is Nutanix's integrated workflow for assessing a cluster's current component versions, identifying applicable update bundles, performing compatibility checks, and orchestrating supported platform firmware updates. It is designed to simplify the operational lifecycle of Nutanix nodes by centrally managing validated updates rather than requiring administrators to update each hardware component manually. BMC, BIOS, and HBA Firmware are supported firmware categories that LCM can inventory and update when the node platform and installed bundle provide the applicable payloads. The BMC update maintains the node's out-of-band management-controller firmware, while the BIOS update maintains the server's system firmware. HBA Firmware updates maintain the firmware for host-bus adapters used for storage connectivity. Before applying updates, LCM runs prechecks and presents the available updates so an administrator can review readiness and sequence the maintenance operation appropriately. Hardware and firmware availability can vary by OEM model and Nutanix compatibility support, so the LCM inventory is the authoritative view for a particular cluster. See the official [Nutanix Lifecycle Manager documentation](#) and the [Nutanix documentation portal](#) for supported LCM workflows and release-specific compatibility information.

QUESTION NO: 20

Write I/O that is deemed sequential will enter which component of the Distributed Storage Fabric first?

- A. Oplog
- B. Extent store
- C. Cache
- D. Metadata Ring

ANSWER: B

Explanation:

Extent store is correct. Nutanix Distributed Storage Fabric (DSF) classifies incoming write I/O according to its access pattern. When a write is identified as sequential, DSF directs it to the extent store, where data is organized in larger extents and written in a manner suited to sustained, streaming-style workloads. This path avoids treating sequential data as small, random writes and enables the platform to use its storage resources efficiently for workloads such as large file transfers, log streams, and sequential database activity.

The extent store is a core DSF component responsible for storing the data extents that make up virtual disk content. DSF uses its data-placement and I/O-handling mechanisms to provide a distributed storage layer across the cluster while maintaining the expected resiliency and performance characteristics. Nutanix documentation describes DSF as the distributed software layer that presents storage services across cluster nodes and manages data placement and access. See the [Nutanix Cloud Infrastructure overview](#) and the [Nutanix documentation portal](#) for platform and storage architecture documentation.

QUESTION NO: 21

An administrator needs to view and compare IOPs for VMs across multiple clusters.

Which two dashboards can the administrator use to find this information? (Choose two)

- A. Planning
- B. Apps
- C. Explore
- D. Analysis
- E. Home

ANSWER: C D

Explanation:

Explore and **Analysis** are the appropriate Prism Central dashboards for examining and comparing VM IOPS across registered clusters. Explore provides an interactive performance-investigation view in which an administrator can select VM entities, choose performance metrics such as IOPS, and evaluate the resulting time-series data. Because Prism Central aggregates information from its managed clusters, this view supports a centralized investigation rather than requiring the administrator to open each individual cluster separately.

Analysis is intended for building and saving customized performance dashboards. An administrator can create charts that use VM performance metrics, select the relevant VMs from different clusters, and place their IOPS data together for comparison over a chosen time range. This is particularly useful for repeatable operational monitoring and for sharing a consistent visualization with other administrators. Together, these dashboards provide both ad hoc metric exploration and customizable comparative reporting for VM I/O activity in a multicluster environment. Nutanix describes Prism Central as the centralized management plane for monitoring and managing multiple clusters; see the [Nutanix Prism overview](#) and the [Nutanix Documentation portal](#) for Prism Central administration documentation.

QUESTION NO: 22

Which two features require the Nutanix Guest Tools package is installed (Choose Two)

- A. Self-Service Restore
- B. Volume Snapshot Services
- C. Network Visualization
- D. Image service
- E. Prism Central

ANSWER: A B

Explanation:

Self-Service Restore requires Nutanix Guest Tools (NGT) because the guest-side components enable the virtual machine to participate in the restore workflow. NGT provides the communication and integration needed for an authorized user to

browse and restore eligible files from a recovery point without requiring an administrator to perform the entire operation. Installing and maintaining NGT in supported guest operating systems is therefore a prerequisite when this guest-aware self-service capability is needed.

Volume Snapshot Services also requires NGT. In Windows guests, NGT includes the Nutanix VSS provider, which integrates with Microsoft Volume Shadow Copy Service to create application-consistent snapshots. This coordination lets applications and Windows flush or quiesce relevant I/O before the snapshot is taken, helping ensure that data captured by protection workflows is in a consistent state. The NGT installation and its VSS component must be present and operational in the guest for this integration to be used. Nutanix documents NGT deployment and supported guest features in its [product documentation portal](#); installation media and associated resources are available through the [Nutanix Support Portal downloads page](#).

QUESTION NO: 23

An administrator needs to convert an existing node to be a Compute Only node. Which two steps should the administrator take? (Choose two.)

- A. Update the license file
- B. Enable CO mode using aCLI
- C. Image the node as CO using Foundation
- D. Add the node to the cluster using Prism Element

ANSWER: B C

Explanation:

Converting a node for compute-only use requires both configuring the node's compute-only designation and reinstalling it with the appropriate compute-only imaging workflow. **Enable CO mode using aCLI** is required because the Acropolis command-line interface is used to set the node's compute-only mode before it is brought back into service. This establishes that the node is intended to provide CPU and memory resources rather than participate as a storage-bearing node.

Image the node as CO using Foundation is also required. Nutanix Foundation performs the supported reimaging process and deploys the software configuration appropriate to a compute-only node. After imaging, the node can be added back to the cluster through the applicable node-expansion workflow. Performing both actions ensures that the node is consistently configured at both the software-installation and cluster-role levels, allowing it to contribute compute capacity while avoiding unintended storage participation. For Nutanix documentation and supported procedures, consult the [Nutanix Support Portal](#) and the [Nutanix Documentation portal](#).

QUESTION NO: 24

An administrator manages a 4-block 16 node Nutanix cluster. One block experiences a Power Delivery Unit (PDU) failure. All nodes in that block power off. VMs on the affected nodes restart as expected. The cluster continues normal operation until power is restored.

Which Nutanix feature is responsible for maintaining operations in this scenario?

- A. IO Redirection
- B. Block Awareness
- C. Async Disaster Recovery
- D. Nutanix App Mobility Fabric

ANSWER: B

Explanation:

Block Awareness is the Nutanix resiliency capability that protects a cluster from the loss of an entire block, such as the simultaneous outage caused by a block-level PDU failure. A Nutanix block contains multiple nodes that can share common physical infrastructure, including power components. Block Awareness ensures that data replicas are placed across separate blocks rather than being concentrated within the same block. Therefore, when every node in one block becomes unavailable, the cluster retains accessible replicas of the affected data on nodes in the remaining blocks.

In this four-block cluster, the surviving nodes and Controller VMs continue to provide storage services and cluster functions. Nutanix AHV high availability detects that the hosts in the failed block are down and restarts their virtual machines on available hosts elsewhere in the cluster. Because the required VM data remains available from replicas outside the failed block, the restarted VMs can resume operation and the cluster can continue serving workloads while the failed block is repaired or power is restored.

This behavior is a core aspect of Nutanix distributed storage and infrastructure fault-domain resilience. See Nutanix's [Nutanix Cloud Platform overview](#) and [Nutanix AHV overview](#) for platform resiliency and availability capabilities.

QUESTION NO: 25

An administrator notices that most of the VMs in the cluster are on one host. Users report that an application seems to respond slowly. The application server VM has significantly more memory assigned to it than other VMs.

How should the administrator fix this issue?

- A. Reduce the amount of memory assigned to the VM.
- B. Migrate the VM to a different host.
- C. Add more memory to the VM.
- D. Increase the memory on the CVM.

ANSWER: B

Explanation:

Migrate the VM to a different host. This situation indicates an uneven distribution of virtual-machine workloads: most VMs are concentrated on a single host, and the application VM is comparatively memory intensive. That host can consequently experience memory contention and greater CPU, storage, and network scheduling pressure, which can increase application response time even when the VM itself has a large configured memory allocation. Moving the application server VM to a host with more available resources rebalances the cluster and gives the workload access to less-contended compute and memory capacity. In AHV, VM migration is a standard administrative operation and can be performed while preserving workload availability through live migration when the cluster and VM conditions support it. After the migration, the administrator should validate host utilization and application performance in Prism to confirm that the workload is no longer affected by host-level resource contention. Nutanix describes AHV as providing integrated VM management and mobility capabilities at [Nutanix AHV](#). Nutanix's documentation portal, including AHV administration guidance for VM operations, is available at [Nutanix Support Portal](#).

QUESTION NO: 26

An administrator wants to make sure that the health of a cluster is in a good state.

Which command the utility will provide information related to the health of the cluster?

- A. ncc
- B. generic status
- C. top
- D. iostat

ANSWER: A

Explanation:

ncc is the Nutanix Cluster Check utility and is the appropriate utility for assessing overall Nutanix cluster health. NCC runs a collection of built-in checks across cluster services, storage, networking, hardware, configuration, and other operational dependencies. Its output identifies checks that pass, produce warnings, or fail, giving an administrator a consolidated view of conditions that could affect cluster stability or availability. A common health-validation workflow is to run NCC checks from a Controller VM, such as `ncc health_checks run_all`, and then review the reported findings before or after maintenance and during troubleshooting. NCC is designed specifically for Nutanix environments and is regularly updated to include checks appropriate to supported AOS and platform configurations. Nutanix documents NCC as the cluster-health validation tool and provides guidance on running checks and interpreting results in the Nutanix Cluster Check Guide. See the [Nutanix Cluster Check \(NCC\) documentation](#) and the [Nutanix Support Portal documentation library](#) for version-specific command guidance.

QUESTION NO: 27

Which two conditions can cause an AHV VM live migration to fail? (Choose two.)

- A. The destination host lacks sufficient CPU or memory resources
- B. The VM network is unavailable on the destination host
- C. The VM has an IP address assigned by IPAM
- D. The VM has a snapshot

ANSWER: A B

Explanation:

A live migration can fail when the destination host does not have sufficient resources to run the VM. AHV must be able to satisfy the VM's CPU and memory requirements on the destination host. A migration can also fail when the destination host does not provide the guest network required by the VM, because the VM must retain connectivity to its configured network after it moves.

Administrators should validate destination capacity, host connectivity, applicable affinity policies, and network availability before initiating a migration. See the [Nutanix Documentation portal](#) for AHV live migration requirements and troubleshooting information.

QUESTION NO: 28

Which Nutanix Product combines the Distributed Storage Fabric & App Mobility Fabric?

- A. Xtract
- B. Prism
- C. Acropolis
- D. Calm

ANSWER: C

Explanation:

Acropolis is the correct answer. In Nutanix architecture and product terminology, Acropolis brings together the platform's distributed storage capabilities and its application/VM mobility and virtualization capabilities. The Distributed Storage Fabric pools local storage resources across cluster nodes into a resilient, scale-out storage fabric that presents storage services to workloads. The App Mobility Fabric provides the compute and virtualization layer used to run and manage workloads with mobility across the cluster. Together, these capabilities form the foundational infrastructure layer that allows Nutanix clusters to deliver integrated compute and storage rather than treating them as separately managed silos.

In current Nutanix naming, these functions are commonly discussed through Nutanix Cloud Infrastructure, including Nutanix Cloud Platform software such as AOS for distributed storage services and AHV for virtualization. However, the exam wording uses the established Acropolis terminology for the combination of these core fabrics. This integrated design is central to Nutanix's scale-out architecture: nodes contribute resources to the cluster, and workloads can use shared distributed services while being operated through a unified platform. Nutanix describes its cloud infrastructure and core platform capabilities at [Nutanix Cloud Infrastructure](#) and provides platform information at [Nutanix Acropolis](#).

QUESTION NO: 29

An administrator is managing a highly active SQL workload that is writing I/Os in 2 MB chunks to a vDisk. The administrator wants to optimize capacity setting for the computer.

Which storage optimization technology should the administration select to meet these requirement?

- A. Erasure Coding
- B. Compression
- C. Capacity Deduplication
- D. Cache Deduplication

ANSWER: B

Explanation:

Compression is the appropriate capacity optimization technology for this workload because Nutanix compression is designed to reduce the physical capacity consumed by data written in larger blocks. A SQL workload issuing writes in 2 MB chunks provides suitably large, contiguous data regions for the compression engine to process efficiently. This can substantially reduce the vDisk's physical storage footprint while maintaining a storage policy appropriate for an actively used database workload.

In Nutanix AOS, compression is a data-reduction capability that operates at the extent-store level and can be enabled through storage efficiency settings or the relevant data-reduction policy. It is particularly useful where data has enough redundancy within larger write regions to yield meaningful savings. The administrator should validate the actual compression ratio and performance behavior for the specific SQL data set, since savings vary by data type and already-compressed or encrypted data generally has limited compressibility. Nutanix also provides per-container and per-VM controls so the policy can be applied to the intended workload rather than indiscriminately across all data.

For additional implementation and monitoring guidance, consult the [Nutanix Bible: AOS storage efficiency overview](#) and the [Nutanix Support Portal documentation library](#).