

Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI) 5.15

Nutanix NCM-MCI-5.15

Version Demo

Total Demo Questions: 12

Total Premium Questions: 129

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Design a Nutanix Solution	25
Topic 2, Deploy a Nutanix Solution	11
Topic 3, Manage a Nutanix Solution	63
Topic 4, Optimize a Nutanix Solution	30
Total	129

QUESTION NO: 1

An administrator needs to replace an aging SAN and move to a hyper-converged infrastructure. The existing environment consists of the following hosts that are connected to the SAN:

- 5x AIX hosts
- 3x Hyper-V hosts
- 9x ESXi hosts
- 2x physical SQL Clusters (Windows Server 2012R2 hosts)

After deploying a Nutanix AHV cluster, which two actions should the administrator take to meet the requirements? (Choose two.)

- A.** Deploy Volumes to support the AIX and SQL workloads.
- B.** Migrate the ESXi workloads to AHV using Move.
- C.** Deploy Files to support the AIX hosts.
- D.** Migrate the ESXi and Hyper-V workloads using Move.

ANSWER: A D

Explanation:

Deploy Volumes to support the AIX and SQL workloads because Nutanix Volumes provides block storage from the Nutanix cluster to external guest operating systems and physical servers through iSCSI. This lets the AIX systems retain SAN-like block-storage access while the physical Windows Server SQL clusters can consume shared block storage from the new Nutanix environment. Nutanix Volumes is therefore the appropriate storage service for workloads that are not being converted into AHV virtual machines but still require external, centrally managed storage.

Migrate the ESXi and Hyper-V workloads using Move because Nutanix Move is the Nutanix migration service designed to move virtual machines from supported VMware vSphere and Microsoft Hyper-V source environments to AHV. Move performs the migration workflow, including replication and cutover, allowing the existing virtualized workloads to be brought onto the AHV cluster as part of the SAN replacement and HCI transition. Together, Volumes preserves block-storage service for the external AIX and physical SQL workloads, while Move transitions the VMware and Hyper-V virtual workloads to AHV. See the [Nutanix Volumes Guide](#) and the [Nutanix Move documentation](#).

QUESTION NO: 2 - (HOTSPOT)

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

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

Explanation:

For an all-flash Nutanix cluster, the appropriate data-reduction configuration is selected according to the I/O characteristics and data similarity of each workload. Full-clone virtual desktops are a strong fit for enabling both compression and deduplication in their container. Full clones contain substantial repeated operating-system, application, and desktop-image data. Deduplication removes duplicate blocks, while compression further reduces the physical capacity consumed by the remaining blocks. Using both capabilities provides the intended efficiency for a desktop pool made from full copies.

Hadoop is appropriately placed in a container with compression and Erasure Coding enabled. Hadoop commonly retains large amounts of distributed data, so reducing the stored footprint is valuable. Compression reduces the size of compressible data, and Erasure Coding provides a capacity-efficient resiliency method for eligible data extents. Nutanix describes Erasure Coding as a data-efficiency feature that reduces the capacity overhead associated with replicated data while maintaining fault tolerance. See the [Nutanix Prism Web Console Guide: Data Reduction](#) for the container-level data-reduction features.

Write-heavy databases should use a container with compression enabled. Compression provides space savings while being suited to this workload requirement; it is the selected data-efficiency feature for the database container. Keeping these workload classes in separate containers lets the administrator apply the storage-efficiency policy that matches each workload rather than applying one uniform policy across unrelated data types. Nutanix's [Storage Containers documentation](#) explains that storage policies, including data-reduction settings, are configured at the container level.

QUESTION NO: 3

An administrator needs to deploy VMs on a new AHV network. The VMs must receive addresses automatically from an IP address range managed by Nutanix.

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

- A. Create a managed network and assign the required VLAN.
- B. Define an IP address pool for the managed network.
- C. Configure the CVM external IP address as the VM default gateway.
- D. Enable Erasure Coding on the VM storage container.

ANSWER: A B**Explanation:**

The administrator must create a managed network in Prism and associate it with the required VLAN. A managed network allows Prism to provide IP address management services for VM NICs connected to that network while the VLAN identifies the Layer 2 network used by the AHV hosts and physical switching infrastructure.

The administrator must also define an IP address pool for the managed network. The pool identifies the addresses available for allocation to VM NICs. Prism can then assign an unused address from the configured range when a VM is created or updated to use IP address management.

See the [Nutanix Documentation Portal](#) for AHV network and IP address management guidance.

QUESTION NO: 4

An administrator needs to protect a Windows Server VM using application-consistent snapshots. The VM hosts a supported application that uses Microsoft Volume Shadow Copy Service (VSS).

Which two actions are required to create application-consistent recovery points? (Choose two.)

- A. Install and run Nutanix Guest Tools in the Windows VM.
- B. Select application-consistent snapshots in the Protection Domain schedule.
- C. Enable Erasure Coding on the storage container.

D. Attach the VM disks through a Volume Group.

ANSWER: A B

Explanation:

Nutanix Guest Tools must be installed and running in the Windows VM. NGT includes the Nutanix VSS components that coordinate with the Windows guest operating system and registered VSS writers during snapshot creation.

The administrator must select application-consistent snapshots in the Protection Domain schedule. This instructs Nutanix to quiesce the application through VSS before taking the snapshot. Without application-consistent snapshot configuration, protection points are crash-consistent rather than coordinated with the application. Nutanix describes NGT and application-consistent snapshot workflows in the [Nutanix Data Protection documentation](#).

QUESTION NO: 5

A customer has Metro Availability configured in Automatic Resume Mode between Site A and Site B. What happens to VMs if Primary Site A has full network failure?

- A. Automatic protection domain promotion in Site B. VMs automatically restart in Site B
- B. VMs continue to run on Site A following an automatic disable of protection domains.
- C. An administrator must unwitness in Site B and promote protection domains in Site B for VMs to restart
- D. An administrator must promote protection domains in Site B for VMs to restart.


ANSWER: D

Explanation:

When Metro Availability is configured in Automatic Resume mode and the primary site experiences a full network failure, an administrator must promote the protection domains at Site B before the protected VMs can restart there. Promotion makes Site B the active owner of the Metro-protected entities and enables VM recovery at the surviving site. This is an intentional administrative action because a full network failure can create uncertainty about the original site's state; manual promotion helps prevent an unintended dual-active condition and preserves the expected protection-domain ownership process.

Automatic Resume mode applies to resuming Metro Availability protection after the failed site is restored and communication is re-established. It does not eliminate the requirement to initiate protection-domain promotion during this failure scenario. After Site B has been promoted, the administrator can recover or restart the VMs at Site B, then follow the documented Metro Availability recovery procedure when Site A returns to service. Nutanix documents the witness and Metro Availability failover/recovery workflow in the [Prism Web Console Guide](#). Additional Metro Availability architecture and operational guidance is available in the [Nutanix Disaster Recovery documentation](#).

QUESTION NO: 6



An administrator logged into Prism Element, notices the following in the home dashboard, and showing the exhibit.

Which action should the administrator take next?

- A. Configure the fqdn for the hosts indicated
- B. Troubleshoot the zookeeper server issue
- C. Re-run NCC specifying data_resiliency_check
- D. Check the hardware page for disk failures

ANSWER: C

Explanation:

Re-run NCC specifying data_resiliency_check is the appropriate next action because a Prism Element dashboard indication involving data resiliency should be investigated through Nutanix Cluster Check (NCC). NCC is the Nutanix-supported diagnostic framework for validating cluster health and identifying conditions that can affect the availability, protection, or recoverability of data. Running the relevant data-resiliency health check narrows the investigation to the affected resiliency condition and returns actionable findings, including the relevant cluster objects and recommended remediation steps. This is preferable to making an immediate configuration change based only on a dashboard status, because the dashboard is an alerting and summary interface rather than the detailed diagnostic source. The administrator should review the resulting NCC output, determine the precise resiliency condition reported, and then perform the remediation recommended for that finding. Nutanix documents NCC as the utility used to run cluster health checks and diagnose operational issues; see the [Nutanix documentation portal](#) and the [Nutanix Community overview of NCC](#).

QUESTION NO: 7

An administrator migrates a VM onto a new Nutanix cluster. After the migration, the administrator observes the following conditions:

- Cluster memory utilization: 64%
- Cluster CPU utilization: 19%
- Cluster storage utilization: 32%
- Average VM CPU utilization: 25%
- Average VM CPU ready%: 24%
- Average VM memory utilization: 60%

Which two changes should the administrator make to improve VM performance? (Choose two.)

- A. Add more memory to the VMs.
- B. Reduce the number of vCPUs assigned to VMs.
- C. Replace high core count CPUs with high clock rate CPUs. (also can be but is physical invas)
- D. Reduce the number of VMs on the hosts.

ANSWER: B D

Explanation:

High CPU ready time indicates that virtual CPUs are spending a substantial amount of time waiting to be scheduled on physical CPU resources. Here, the average VM CPU ready value of 24% is the key performance indicator: it is far above a healthy baseline even though aggregate cluster CPU utilization appears low. This pattern commonly occurs when virtual machines have been allocated more vCPUs than their workloads need. A multi-vCPU VM must be scheduled appropriately, and excessive vCPU allocation can create scheduling delay and contention that is not apparent from the cluster-wide CPU utilization percentage alone.

Reduce the number of vCPUs assigned to VMs. right-sizes the VMs to their actual CPU demand, reducing scheduling complexity and CPU ready time. **Reduce the number of VMs on the hosts.** decreases the density of competing workloads

on the affected hosts, giving the remaining VMs more scheduling opportunities and improving responsiveness. These changes should be validated by monitoring CPU ready time and VM latency after rebalancing or resizing. Nutanix Prism provides VM and host performance metrics to identify CPU contention and assess the effect of configuration changes. See the [Nutanix Prism metrics documentation](#) and [Nutanix virtualization best practices](#).

QUESTION NO: 8

An administrator is deploying an application that requires maximum I/O throughput for scratch data. The administrator is concerned that the throughput requirement is greater than what can be provided by a single cluster node.

What should the administrator do to meet this requirement?

- A. Create a volume group with several virtual disks and attach to the VM
- B. Create a volume group with one virtual disk attached to the VM
- C. Create a virtual disk attached to the VM's PCI interface
- D. Create several virtual disks attached to the VM's virtual SCSI controller

ANSWER: A

Explanation:

Create a volume group with several virtual disks and attach to the VM. For a workload whose scratch-data I/O demand exceeds the capacity available through one cluster node, an AHV Volume Group can be configured with multiple virtual disks and used with the Volume Group Load Balancer capability. This distributes the virtual disks across multiple Controller VMs, allowing the guest's I/O to use storage resources from more than one node rather than concentrating all traffic through a single node. The application must issue I/O concurrently across the presented disks to realize the aggregate throughput benefit; therefore, the guest should use an appropriate striping or volume-management configuration when the workload supports it. Volume Groups are especially suitable for high-performance block storage presented to a VM, including temporary or scratch-data use cases where the application requires scalable bandwidth. Nutanix describes this approach and the associated scaling behavior in its Volume Group Load Balancer technical discussion: [Scaling Up VM Storage Performance on AHV Volume Group Load Balancer](#). Additional architecture context for Nutanix distributed storage and Controller VM data services is available in the [Nutanix Bible](#).

QUESTION NO: 9

An administrator needs to expand a cluster based on AHV and running on Nutanix G5 hardware with a new Nutanix G7 system. The cluster is running AOS 5.10 but the operation fails when expanding the cluster. AOS is running on the new system. There are no issues discovering the node.

Why is the operation failing?

- A. AOS version running on G7 is older than the version on the cluster.
- B. Foundation software has not been updated on the cluster.
- C. G7 hardware cannot be added to a G5 cluster.
- D. EVC needs to be configured before cluster expansion

ANSWER: A

Explanation:

AOS version running on G7 is older than the version on the cluster. Nutanix cluster expansion validates software compatibility after the candidate node has been discovered. Successful discovery confirms network reachability and that the cluster can communicate with the new node, but it does not mean that the node is eligible to join. The Controller VM and AOS software on an expansion node must meet the version requirements of the existing cluster. In this case, the G5-based cluster is already running AOS 5.10, while the G7 node is on an earlier AOS release. The older node software cannot safely

join a cluster that is running the later release because it may not contain the required services, APIs, metadata handling, and hardware-platform support expected by the cluster. The administrator should use Foundation or the applicable Nutanix upgrade procedure to install a supported, compatible AOS version on the G7 node before repeating the expansion workflow. Nutanix also documents platform-mixing and version requirements as prerequisites for expansion, including the need to verify supported hardware and software combinations. See the Nutanix community discussion on [product mixing restrictions](#) and the [Prism Web Console documentation for adding nodes](#).

QUESTION NO: 10

An administrator needs to maximize storage potential in a six-node cluster. The container is configured with the following:

- ☞ Post-process compression
- ☞ Capacity deduplication

What should the administrator configure to increase space savings?

- A. Erasure coding
- B. Inline compression
- C. Reserve capacity
- D. Cache dedupe

ANSWER: A

Explanation:

Erasure coding increases space savings by reducing the capacity overhead required to provide data resiliency. In a six-node Nutanix cluster, erasure coding can be used for eligible data in the capacity tier, distributing data and parity information across nodes rather than maintaining full replica copies. This allows the cluster to preserve fault tolerance while consuming less raw storage capacity than a replication-based layout. The container already uses post-process compression and capacity deduplication, which reduce the physical footprint of stored data; erasure coding complements those efficiency features by optimizing how the remaining data is protected. It is particularly valuable for data that is not latency-sensitive and has become eligible for placement in the capacity tier, helping maximize usable storage without removing resiliency. A six-node cluster provides sufficient nodes for Nutanix erasure coding, making it an appropriate configuration to increase overall space efficiency. Nutanix documentation describes data efficiency technologies and their interaction with storage capacity in the [Nutanix Product Documentation](#) and the [Nutanix Unified Storage overview](#).

QUESTION NO: 11

An administrator is currently using both 10g uplink with LACP and balance-tcp. A single VM should not be allowed to use more than a single 10G uplink, and both uplinks should be utilized.

Which two command should be used to configure the bond? (choose two)

- A. `Ovs-vsctl set port br0-up bond_mode=balance-slb`
- B. `Ovs-vsctl set port br0-up bond_mode=balance-tcp`
- C. `Ovs-vsctl set port br0-up bond_mode=active-backup`
- D. `Ovs-vsctl set port br0-up other_config: bond-rebalance-interval=30000`

ANSWER: A D

Explanation:

`Ovs-vsctl set port br0-up bond_mode=balance-slb` is appropriate because source-load balancing selects an uplink using the source MAC address and VLAN. For VM traffic, this keeps a given VM's traffic assigned to one physical member of the bond at a time, so that VM cannot aggregate bandwidth across both 10 GbE uplinks. At the same time, different VMs can be assigned to different uplinks, allowing the available links to be used collectively.

`Ovs-vsctl set port br0-up other_config: bond-rebalance-interval=30000` enables periodic load rebalancing at a 30-second interval. This is important with source-load balancing because VM workloads can change over time; periodic rebalancing can move eligible MAC/VLAN assignments between bond members to improve distribution while preserving the one-uplink-per-VM behavior at any given time. The configuration therefore meets both requirements: per-VM traffic remains limited to a single 10 GbE uplink, while the bond can distribute multiple VM flows across both uplinks. Open vSwitch documents the behavior of the `balance-slb` mode and the `bond-rebalance-interval` setting in its [ovs-vswitchd configuration database reference](#). Nutanix AHV uses Open vSwitch for host virtual switching, as described in the [Nutanix AHV networking documentation](#).

QUESTION NO: 12

An administrator manages a Prism Central environment with several development teams. Each team must be able to create and manage its own VMs, but must not be able to consume more than its assigned CPU, memory, and storage resources.

Which Prism Central feature should the administrator use?

- A. Create a Project with resource quotas.
- B. Create an image placement policy with Hard enforcement.
- C. Configure a Protection Domain for each development team.
- D. Enable Capacity Deduplication on the development container.

ANSWER: A

Explanation:

Create a Project with resource quotas. Prism Central Projects provide a logical tenant boundary for a group of users and can define limits for resources such as vCPUs, memory, and storage. Users assigned to the Project can deploy and manage workloads within the Project scope, while the configured quotas prevent the team from consuming resources beyond the approved allocation.

Projects can also define the clusters, networks, images, and categories available to project users. This provides controlled self-service access without granting unrestricted access to all infrastructure resources in Prism Central. Nutanix documents Projects and quota-based governance in the [Prism Central documentation](#).