

Nutanix

NCP-MCI-6.10 Exam


Nutanix Certified Professional - Multicloud Infrastructure v6.10

**Questions & Answers
Demo**

Version: 4.0

Question: 1

The customer expects to maintain a cluster runway of 9 months. The customer doesn't have a budget for 6 months but they want to add new workloads to the existing cluster.



Based on the exhibit, what is required to meet the customer's budgetary timeframe?

- A. Add resources to the cluster.
- B. Postpone the start of new workloads.
- C. Delete workloads running on the cluster.
- D. Change the target to 9 months.

Answer: A

Explanation:

The exhibit shows that the overall runway is only 66 days, meaning that the current cluster does not have enough capacity to sustain workloads for 6 months, let alone 9 months.

The best solution is to add resources to the cluster (Option A), such as CPU, memory, or storage, to extend the runway.

Postponing new workloads (Option B) may help in the short term but does not align with the business need to continue adding workloads.

Deleting workloads (Option C) is not a viable option because the customer wants to add more, not remove them.

Changing the target to 9 months (Option D) does not change the actual resource constraints; it only alters the target timeframe.

Reference:

Nutanix Prism Central → Capacity Planning and Runway Analysis
Nutanix Bible → Cluster Resource Management and Scaling
Nutanix Support KB → How to Extend Cluster Runway with Resource Scaling

Question: 2

An administrator is trying to configure Metro Availability between Nutanix ESXi-based clusters. However, the Compatible Remote Sites screen does not list all required storage containers. Which two reasons could be a cause for this issue? (Choose two.)

- A. Source and destination hardware are from different vendors.
- B. The remote site storage container has compression enabled.
- C. The destination storage container is not empty.
- D. Both storage containers must have the same name.

Answer: C, D

Explanation:

For Metro Availability to work properly, the storage containers at both the primary and secondary sites must meet certain requirements:

The storage containers must have the same name (Option D). This ensures that replication and failover work seamlessly. If the names do not match, the storage containers will not be listed as compatible.

The destination storage container must be empty (Option C). Metro Availability requires a clean storage container at the secondary site to receive data. If the container already contains data, it cannot be used.

Option A is incorrect: Different vendor hardware does not affect Metro Availability compatibility.

Option B is incorrect: Compression does not affect Metro Availability compatibility. However, it is recommended to keep compression settings aligned between sites.

Reference:


Nutanix Documentation: Metro Availability Deployment Guide

Nutanix Best Practices for Metro Availability

Nutanix KB 2093: Troubleshooting Metro Availability Storage Container Issues

Question: 3

An administrator receives complaints about VM performance.



After reviewing the VM's CPU Ready Time data shown in the exhibit, which step should the administrator take to diagnose the issue further?

- A. Check the number of vCPUs assigned to each CVM.
- B. Review host CPU utilization.
- C. Assess cluster SSD capacity.
- D. Enable VM memory oversubscription.

Answer: B

Explanation:

Understanding the Issue

The administrator is investigating VM performance complaints and is analyzing CPU Ready Time data.

CPU Ready Time is a crucial metric in Nutanix and virtualization environments (AHV, ESXi, or Hyper-V).

It measures the amount of time a VM is waiting for CPU scheduling due to resource contention. High CPU Ready Time indicates that VMs are ready to run but are waiting because the host lacks available CPU resources.

Analysis of the Exhibit

The graph shows CPU Ready Time spikes for multiple VMs.

Some VMs have CPU Ready Time exceeding 18% to 21.5%, which is very high.

A healthy CPU Ready Time should be below 5%.

Values above 10% indicate CPU contention, and anything above 20% is critical and requires

immediate troubleshooting.

Evaluating the Answer Choices

✗ (A) Check the number of vCPUs assigned to each CVM. (Incorrect)

CVMs (Controller VMs) have fixed CPU allocation, and modifying their vCPU count is not recommended unless advised by Nutanix Support.

The issue is related to VM CPU contention, not CVM configuration.

✓ (B) Review host CPU utilization. (Correct Answer)

High CPU Ready Time suggests CPU overcommitment or host saturation.

The administrator should check host CPU usage in Prism Central to determine if the cluster is overloaded.

If host CPU usage is consistently above 85–90%, VMs are competing for CPU resources, leading to high CPU Ready Time.

✗ (C) Assess cluster SSD capacity. (Incorrect)

SSD capacity impacts storage performance (latency, read/write speeds) but does not affect CPU Ready Time.

High CPU Ready Time is a CPU scheduling issue, not a storage bottleneck.

✗ (D) Enable VM memory oversubscription. (Incorrect)

Memory oversubscription does not impact CPU scheduling.

Enabling memory oversubscription affects RAM allocation, but CPU Ready Time is strictly related to CPU contention.

Next Steps to Diagnose & Resolve the Issue

Review Host CPU Utilization:

Navigate to Prism Central → Analysis → CPU Usage per Host.

Identify hosts experiencing high CPU load.

Check VM vCPU Allocation:

Ensure that VMs do not have excessive vCPUs assigned, which can lead to scheduling inefficiencies. Overprovisioning vCPUs can cause unnecessary contention.

Balance Workload Across Hosts:

Use Nutanix AHV DRS (Dynamic Scheduling) or VMware DRS to redistribute VMs across hosts.

Check if certain hosts are overloaded while others have spare CPU capacity.

Consider Scaling Out the Cluster:

If CPU usage is consistently high, adding more nodes may be required to reduce CPU contention.

Multicloud Infrastructure Reference & Best Practices

CPU Ready Time Best Practices:

Keep CPU Ready Time below 5%.

Avoid overcommitting vCPUs on heavily loaded hosts.

Monitor Prism Central Runway Metrics to predict future CPU resource needs.

Nutanix AHV CPU Scheduling Optimization:

Ensure proper VM sizing (avoid excessive vCPU allocation).

Balance workloads using Nutanix AHV DRS.

Reference:

Nutanix Prism Central: Performance Analysis and CPU Metrics

Nutanix Bible: VM Performance and Resource Management

Nutanix KB: Troubleshooting High CPU Ready Time in AHV

Question: 4

Refer to Exhibit:

Cluster Details

Virtual IP / FQDN is used to access the PC VM Cluster.

Cluster Name

Unnamed

FQDN

Virtual IP

In a scale-out Prism Central deployment, what additional functionality does configuring an FQDN instead of a Virtual IP provide?

- A. Load balancing
- B. Resiliency
- C. Segmentation
- D. SSL Certificate

Answer: A

Explanation:

When using FQDN instead of a Virtual IP in a scale-out Prism Central deployment, Nutanix enables load balancing across multiple Prism Central instances.

Option A (Load balancing) is correct because it ensures that requests are distributed among multiple Prism Central nodes, improving performance and redundancy.

Option B (Resiliency) is incorrect because resiliency is achieved through HA and replication, not through FQDN configuration.

Option C (Segmentation) is incorrect because network segmentation is handled at the VLAN or security policy level.


Option D (SSL Certificate) is incorrect because SSL certificates can be applied regardless of whether FQDN or Virtual IP is used.

Reference:

Nutanix Prism Central Deployment Guide
Nutanix Best Practices for Scale-Out Prism Central
Nutanix Support KB: Configuring FQDN for Prism Central

Question: 5

Refer to Exhibit:



After adding new workloads, why is Overall Runway below 365 days and the scenario still shows the cluster is in good shape?

- A. Because Storage Runway is still good.
- B. Because new workloads are sustainable.
- C. Because there are recommended resources.
- D. Because the Target is 1 month.

Answer: B

Explanation:

In Nutanix Capacity Planning, Overall Runway represents how long the cluster can support current and new workloads before resources are exhausted.

Even if the runway is below 365 days, the system considers the cluster to be in good shape if new workloads are sustainable (Option B).

Option A is incorrect: Storage runway alone is not the only factor; CPU and memory are equally important.

Option C is incorrect: The presence of recommended resources does not mean the cluster is in good

shape.

Option D is incorrect: The target of 1 month affects projections but does not explain why the cluster is in good shape.

Reference:

Nutanix Prism Central → Capacity Runway and Planning

Nutanix Bible → Workload Placement and Cluster Sizing

Nutanix Support KB → Capacity Planning Best Practices