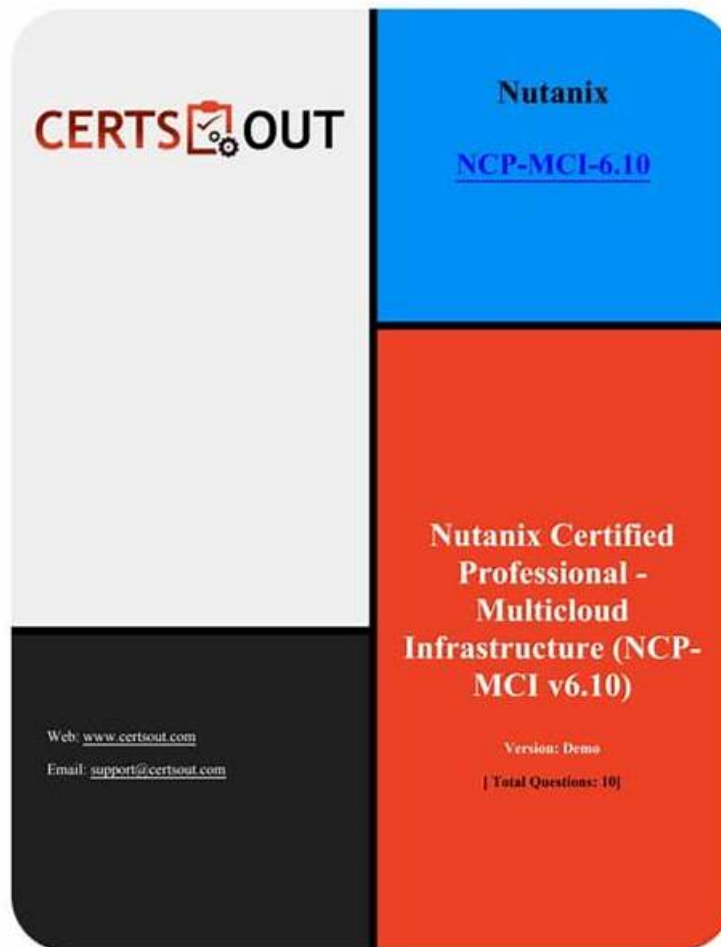


# Dumps NCM-MCI-6.10 Free Download - Exam NCM-MCI-6.10 Blueprint



BONUS!!! Download part of Pass4Leader NCM-MCI-6.10 dumps for free: <https://drive.google.com/open?id=1Lu48i4xl-1pat2CHN6SpsKEa5w7ZnspQ>

So we can say that the NCM-MCI-6.10 practice questions are the top-notch Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI) (NCM-MCI-6.10) dumps that will provide you with everything that you must need for instant Nutanix NCM-MCI-6.10 exam preparation. Take the right decision regarding your quick Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI) (NCM-MCI-6.10) exam questions preparation and download the real, valid, and updated NCM-MCI-6.10 exam dumps and start this journey.

On the pages of our NCM-MCI-6.10 study tool, you can see the version of the product, the updated time, the quantity of the questions and answers, the characteristics and merits of the product, the price of our product, the discounts to the client, the details and the guarantee of our NCM-MCI-6.10 study torrent, the methods to contact us, the evaluations of the client on our product, the related exams and other information about our Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI) test torrent. Thus you could decide whether it is worthy to buy our product or not after you understand the features of details of our product carefully on the pages of our NCM-MCI-6.10 Study Tool on the website.

>> **Dumps NCM-MCI-6.10 Free Download** <<

## **Exam Nutanix NCM-MCI-6.10 Blueprint, Printable NCM-MCI-6.10 PDF**

Pass4Leader's Nutanix NCM-MCI-6.10 web-based and desktop practice tests provide you with an Nutanix actual test scenario, allowing you to experience the NCM-MCI-6.10 final test conditions. Customizable Nutanix NCM-MCI-6.10 Practice Tests

(desktop and web-based) allow you to change the time and quantity of Nutanix NCM-MCI-6.10 practice questions.

## Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI) Sample Questions (Q16-Q21):

### NEW QUESTION # 16

An administrator wants to increase the performance of their Database virtual machine.

Database\_VM has a database that is spread across three vDisks in the volume group Database\_VM. The volume group is directly attached to the virtual machine. Previous performance analysis has indicated all storage requests are going to the same node. While this test environment has 1 node, the production environment has 3 nodes.

Configure the Volume Group Database\_VM so that it's optimized for the user's VM and the production environment. The virtual machine has been powered off and moved to this test cluster for the maintenance work.

Note: Do not power on the VM.

#### Answer:

Explanation:

See the Explanation below for detailed answer.

Explanation:

Here is the step-by-step solution to configure the Volume Group for optimized performance in the production environment.

This task is performed in Prism Central.

\* From the main dashboard, navigate to Compute & Storage > Volume Groups.

\* Find the Volume Group named Database\_VM in the list.

\* Select the checkbox next to Database\_VM.

\* Click the Actions dropdown menu and select Update.

\* In the "Update Volume Group" dialog, scroll to the bottom of the "Basic Configuration" section.

\* Find the checkbox labeled Enable Client Side Load Balancing and check it.

Note: This setting allows the iSCSI initiator within the guest VM to connect to all CVMs in the cluster, distributing the storage load from the three vDisks across all three nodes in the production environment instead of focusing all I/O on just one.

Click Save.

### NEW QUESTION # 17

Task 14

An administrator has been asked to configure a storage for a distributed application which uses large data sets across multiple worker VMs.

The worker VMs must run on every node. Data resilience is provided at the application level and low cost per GB is a Key Requirement.

Configure the storage on the cluster to meet these requirements. Any new object created should include the phrase Distributed\_App in the name.

#### Answer:

Explanation:

See the Explanation for step by step solution.

Explanation:

To configure the storage on the cluster for the distributed application, you can follow these steps:

Log in to Prism Element of cluster A using the credentials provided.

Go to Storage > Storage Pools and click on Create Storage Pool.

Enter a name for the new storage pool, such as Distributed\_App\_Storage\_Pool, and select the disks to include in the pool. You can choose any combination of SSDs and HDDs, but for low cost per GB, you may prefer to use more HDDs than SSDs.

Click Save to create the storage pool.

Go to Storage > Containers and click on Create Container.

Enter a name for the new container, such as Distributed\_App\_Container, and select the storage pool that you just created, Distributed\_App\_Storage\_Pool, as the source.

Under Advanced Settings, enable Erasure Coding and Compression to reduce the storage footprint of the data.

You can also disable Replication Factor since data resilience is provided at the application level. These settings will help you achieve low cost per GB for the container.

Click Save to create the container.

Go to Storage > Datastores and click on Create Datastore.

Enter a name for the new datastore, such as Distributed\_App\_Datastore, and select NFS as the datastore type.

Select the container that you just created, Distributed\_App\_Container, as the source.

Click Save to create the datastore.

The datastore will be automatically mounted on all nodes in the cluster. You can verify this by going to Storage > Datastores and clicking on Distributed\_App\_Datastore. You should see all nodes listed under Hosts.

You can now create or migrate your worker VMs to this datastore and run them on any node in the cluster.

The datastore will provide low cost per GB and high performance for your distributed application.

## NEW QUESTION # 18

The Database team is reporting performance degradation for a business-critical application on Saturdays.

The team is requesting monitoring of processor, memory and storage utilization for the cluster for the application: SQL01, SQL02.

The report should contain views for the following:

- \* At the cluster level, only for the Cluster 1: The maximum percentage of CPU used

- \* At the VM level, including any future VM with the prefix SQL: The maximum time taken, maximum percentage of time a VM waits to use the physical CPU, out of the total CPU The report should run on Sundays at 12:00 AM for the previous 7 days. The report should show when completed.

Create a report named SQL\_Batch\_Saturday that meets these requirements.

Generate an instance of the report named SQL\_Batch\_Saturday as a CSV and save the file.

Note: You must name the report SQL\_Batch\_Saturday to receive any credit. Any other name will not be accepted. SMTP is not configured.

### Answer:

Explanation:

See the Explanation below for detailed answer.

Explanation:

Here is the step-by-step solution to create and run the report, performed entirely within Prism Central.

#### 1. Create the Analysis Session

First, we will build the charts and entities for the report using the Analysis tool.

- \* From the Prism Central main menu, navigate to Operations > Analysis.

- \* Click the + New Session button.

- \* Add the required entities:

- \* In the "Entities" search box, type Cluster 1 and select Cluster: Cluster 1.

- \* In the "Entities" search box, type SQL01 and select VM: SQL01.

- \* In the "Entities" search box, type SQL02 and select VM: SQL02.

- \* Click Add Charts > New Chart to add the Cluster CPU chart:

- \* Title: Cluster 1 Max CPU Usage

- \* Entity Type: Cluster

- \* Metric: Cluster CPU Usage %

- \* Aggregation: Select Maximum

- \* Click Add.

- \* Click Add Charts > New Chart to add the VM CPU Ready Time chart:

- \* Title: VM Max CPU Ready Time

- \* Entity Type: VM

- \* Metric: CPU Ready Time %

- \* Aggregation: Select Maximum

- \* Click Add.

#### 2. Save and Configure the Report

Now, save the session as a report and configure it to dynamically include all VMs with the SQL prefix.

- \* Click the Save as Report icon (the bookmark icon in the upper right).

- \* Name the report SQL\_Batch\_Saturday and click Save.

- \* Navigate to Operations > Reports.

- \* Find the SQL\_Batch\_Saturday report you just created and click its name to open the report editor.

- \* In the Entities tile, click the Edit (pencil) icon.

- \* By default, it will list "Cluster 1", "SQL01", and "SQL02".

- \* Change the VM selection:

- \* Select the radio button for All VMs prefixed with.

- \* In the text box, enter SQL.

- \* Ensure "Cluster 1" is still listed under "Clusters".

- \* Click Save.

### 3. Schedule the Report

- \* While still in the report editor for SQL\_Batch\_Saturday, click the Schedule button.
- \* Configure the schedule:
  - \* Recurrence: Weekly
  - \* Repeat on: Sunday
  - \* Start Time: 12:00 AM
  - \* Time Range: Previous 7 Days
- \* Configure the notification (as SMTP is not available):
  - \* Expand the Notification Settings section.
  - \* Check the box for Notify when ready (this enables the bell icon notification).
  - \* Ensure "Email Report" is not checked.
  - \* Click Save.

### 4. Generate and Save the CSV Instance

Finally, run the report now and download the CSV as requested.

- \* Navigate back to the main Operations > Reports list.
- \* Select the checkbox next to SQL\_Batch\_Saturday.
- \* Click the Actions dropdown and select Run Now.
- \* In the dialog, confirm the time range (e.g., "Last 7 Days") and click Run.
- \* Click the Report Instances tab.
- \* Wait for the report instance "SQL\_Batch\_Saturday" to finish running (the status will change from "Running" to Succeeded).
- \* Once it has succeeded, click the Download (arrow) icon for that instance.
- \* Select the CSV format.
- \* Save the file to the desktop.

## NEW QUESTION # 19

Following new security guidelines, it must be ensured that the storage of critical virtual machines will be encrypted in future.

The assignment is to be made by a new category called VM-Storage with a value of softwareencrypted in Prism Central. Make sure a second value of SEDencrypted is also created for future use.

Create the above-mentioned category and perform further configurations in Prism Central for VM-based storage encryption.

Assign the name Encrypted-Storage to the newly created policy.

### Answer:

Explanation:

See the Explanation below for detailed answer.

Explanation:

Here is the step-by-step solution to create the category and the corresponding storage encryption policy within Prism Central.

#### 1. Create the Category

First, you must create the category and the two values requested.

- \* In Prism Central, navigate to Administration > Categories.
- \* Click New Category.
- \* In the Name field, enter VM-Storage.
- \* In the Add a Value field, type softwareencrypted and click the Add (plus) button.
- \* In the Add a Value field again, type SEDencrypted and click the Add (plus) button.
- \* Click Save.

#### 2. Create the Encryption Policy

Next, you will create the security policy that uses the new category.

- \* In Prism Central, navigate to Security > Data-at-Rest Encryption.
- \* Click the + Create Security Policy button.
- \* In the Policy Name field, enter Encrypted-Storage.
- \* Ensure the Encryption Type is set to Software-based.
- \* For Target VMs, select the radio button for VMs matching a category.
- \* In the Select Category dropdown, choose the VM-Storage category you just created.
- \* In the Select Value dropdown, choose softwareencrypted.
- \* Click Save.

This policy will now automatically apply software-based encryption to any new or existing VMs that are assigned the VM-Storage: softwareencrypted category.

## NEW QUESTION # 20

The DB team is requesting an SQL database instance and has requested it be configured for best performance.

This VM has been migrated from a 3 tier solution into Nutanix.

The database VM hosts 4 databases, each set to a 20 GB limit. Logs are expected to not grow beyond 20 GB and should be limited to within 25% to avoid runaway processes. Do not configure more storage than is needed.

The VM that has been migrated is identified as sql3532. Once the VM has been properly reconfigured, the DBA team will reconfigure the OS and database.

The VM should be configured as per KB-3532.

While this VM is being tested, make sure it is the first VM to power up in the event the node it is on goes down.

To maximize performance, ensure as much of the VM as possible will be kept on SSD drives.

Note: The VM does not need to be powered on. The VM should remain on the default container and should not be configured with a volume group. No network is required at this time.

### Answer:

Explanation:

See the Explanation below for detailed answer.

Explanation:

Here is the step-by-step solution to reconfigure the sql3532 virtual machine.

This task is performed from the Prism Element interface for the cluster the VM is on (e.g., Cluster 1).

#### 1. Locate and Update the VM

- \* From the Prism Element main dashboard, navigate to the VM view.

- \* Find the VM named sql3532 in the VM table.

- \* Select the checkbox next to sql3532 and click the Update button.

#### 2. Configure HA Priority and Flash Mode

In the "Update VM" dialog, configure the HA and SSD performance settings:

- \* HA Priority:

- \* Find the VM High Availability section.

- \* Select the High Priority radio button. This ensures it is one of the first VMs to power on during an HA event.

- \* Flash Mode (SSD Performance):

- \* Scroll down to the Flash Mode section.

- \* Check the box to Enable Flash Mode. This pins the VM's vDisks to the SSD tier, satisfying the requirement to keep as much of the VM as possible on SSDs, especially since it's on the default (hybrid) container.

#### 3. Reconfigure Disks (per KB-3532)

While still in the "Update VM" dialog, scroll to the Disks section to add the new data and log disks. The key to "best performance" (KB-3532) is to place Data and Logs on separate vSCSI controllers.

- \* (The VM already has an OS disk, which we will assume is on scsi.0.)

- \* Add Data Disk:

- \* Click the + Add New Disk button.

- \* Storage Container: default (as required).

- \* Size: 80 GB (for the 4 x 20 GB databases).

- \* Bus Type: SCSI.

- \* Device Index: 1. (This creates a new vSCSI controller, scsi.1, for the data disk).

- \* Click Add.

- \* Add Log Disk:

- \* Click the + Add New Disk button.

- \* Storage Container: default (as required).

- \* Size: 20 GB.

- \* Bus Type: SCSI.

- \* Device Index: 2. (This creates a third vSCSI controller, scsi.2, for the log disk).

- \* Click Add.

#### 4. Save Configuration

- \* After adding the disks and setting HA/Flash Mode, click the main Save button at the bottom of the "Update VM" dialog.

The VM is now configured with high availability, its storage is pinned to SSD, and its disk layout follows performance best practices by separating the OS, Data, and Log I/O paths onto three different controllers.

## NEW QUESTION # 21

.....



2026 Latest Pass4Leader NCM-MCI-6.10 PDF Dumps and NCM-MCI-6.10 Exam Engine Free Share:  
<https://drive.google.com/open?id=1Lu48i4xl-1pat2CHN6SpsKEa5w7ZnspQ>