

# Professional NCM-MCI - Nutanix Certified Master - Multicloud Infrastructure v6.10 Test Answers



You can instantly access the practice material after purchasing it from Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI), so you don't have to wait to prepare for the Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI) examination. A free demo of the study material is also available at Exams4Collection. The 24/7 support system is available for the customers, so they can contact the team whenever they face any issue, and it will provide them with the solution.

## Nutanix NCM-MCI Exam Information

- Time Duration: 60 minutes
- The passing score: 73%
- Languages: English

## Nutanix NCM-MCI Exam Study Guide: What You Need To Know

### Which Are The Best Study Guides To Help Pass Nutanix NCM-MCI Exam?

**Nutanix NCM-MCI Exam: Pass with Ease!** a guide about Nutanix certification and tips to pass the exams

If you are eager to pass your Nutanix NCM-MCI Exam, then you've landed to the right place. We've got some of the best study guides with tips that have been proven and tested working by a number of individuals who have passed their exams using these study guides.

While the exams of different certification providers may differ, there are general things that can be done to ensure a passing grade. In this guide, we will look at some important tips, the best resources, and fantastic advice on how to pass an exam. Prepare yourself for the exam and learn how to overcome stress. **Nutanix NCM-MCI Exam Dumps** is an amazing guide that can help you pass the exam with ease.

Nutanix NCM-MCI Exam is a new and innovative platform for creating, deploying, and monitoring applications. In this article I will explain what Nutanix NCM-MCI Exam is and why you should be using it.

## Exam Nutanix NCM-MCI Registration | New NCM-MCI Test Simulator

The purchase process of our NCM-MCI question torrent is very convenient for all people. In order to meet the needs of all customers, our company is willing to provide all customers with the convenient purchase way. The PDF version of our NCM-MCI study tool is very practical, which is mainly reflected on the special function. As I mentioned above, our company are willing to provide all people with the demo for free. You must want to know how to get the trial demo of our NCM-MCI question torrent; the answer is the PDF version. You can download the free demo from the PDF version of our NCM-MCI exam torrent. Maybe you think it does not prove the practicality of the PDF version, do not worry, we are going to tell us another special function about the PDF version of our NCM-MCI study tool.

## Nutanix Certified Master - Multicloud Infrastructure v6.10 Sample Questions (Q15-Q20):

### NEW QUESTION # 15

Task 2

An administrator needs to configure storage for a Citrix-based Virtual Desktop infrastructure.

Two VDI pools will be created

Non-persistent pool names MCS\_Pool for tasks users using MCS Microsoft Windows 10 virtual Delivery Agents (VDAs)

Persistent pool named Persist\_Pool with full-clone Microsoft Windows 10 VDAs for power users

20 GiB capacity must be guaranteed at the storage container level for all power user VDAs. The power user container should not be able to use more than 100 GiB. Storage capacity should be optimized for each desktop pool.

Configure the storage to meet these requirements. Any new object created should include the name of the pool(s) (MCS and/or Persist) that will use the object.

Do not include the pool name if the object will not be used by that pool.

Any additional licenses required by the solution will be added later.

### Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

To configure the storage for the Citrix-based VDI, you can follow these steps:

Log in to Prism Central 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 VDI\_Storage\_Pool, and select the disks to include in the pool. You can choose any combination of SSDs and HDDs, but for optimal performance, you may prefer to use more SSDs than HDDs.

Click Save to create the storage pool.

Go to Storage > Containers and click on Create Container.

Enter a name for the new container for the non-persistent pool, such as MCS\_Pool\_Container, and select the storage pool that you just created, VDI\_Storage\_Pool, as the source.

Under Advanced Settings, enable Deduplication and Compression to reduce the storage footprint of the non-persistent desktops.

You can also enable Erasure Coding if you have enough nodes in your cluster and want to save more space. These settings will help you optimize the storage capacity for the non-persistent pool.

Click Save to create the container.

Go to Storage > Containers and click on Create Container again.

Enter a name for the new container for the persistent pool, such as Persist\_Pool\_Container, and select the same storage pool, VDI\_Storage\_Pool, as the source.

Under Advanced Settings, enable Capacity Reservation and enter 20 GiB as the reserved capacity. This will guarantee that 20 GiB of space is always available for the persistent desktops. You can also enter 100 GiB as the advertised capacity to limit the maximum space that this container can use. These settings will help you control the storage allocation for the persistent pool.

Click Save to create the container.

Go to Storage > Datastores and click on Create Datastore.

Enter a name for the new datastore for the non-persistent pool, such as MCS\_Pool\_Datastore, and select NFS as the datastore type. Select the container that you just created, MCS\_Pool\_Container, as the source.

Click Save to create the datastore.

Go to Storage > Datastores and click on Create Datastore again.

Enter a name for the new datastore for the persistent pool, such as Persist\_Pool\_Datastore, and select NFS as the datastore type.

Select the container that you just created, `Persist_Pool_Container`, as the source.

Click **Save** to create the datastore.

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

You can now use Citrix Studio to create your VDI pools using MCS or full clones on these datastores. For more information on how to use Citrix Studio with Nutanix Acropolis, see **Citrix Virtual Apps and Desktops on Nutanix** or **Nutanix virtualization environments**.

 <https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2079-Citrix-Virtual-Apps-and-Desktops:bp-nutanix-storage-configuration.html>

## NEW QUESTION # 16

Refer to the exhibit.



### Task1

A newly created Windows VM "SQL02" is experiencing poor storage performance when compared to "SQL01" running within the same cluster, on the same storage container.

The cluster is in a healthy state.

Create a new session named `Monitor SQL02` with meaningful metrics. Right click on the session page and click **Select All** then paste this into Notepad and save it as `Task 1.txt` on the desktop.

Also, save the analysis as a report named "MonitorSQL02" and send the report as a PDF on a daily basis to `perf_group@ACME.org`. Reports should not be retained. If any new objects need to be created, use `monitorvm2` in the name.

Finally, correct the issue within "SQL02"

Notes:

Do not power on the VMs.

While you will be creating a session, you will need to examine the VM configurations to determine the issue.

Do not delete the VM to resolve the issue, any other destructive change is acceptable

## Answer:

Explanation:

See the Explanation

Explanation:

This is a classic Nutanix performance troubleshooting scenario. The issue is almost certainly that the VM was created using the wrong Disk Bus Type (IDE or SATA instead of SCSI).

Here is the step-by-step solution to complete Task 1.

Part 1: Analysis and Reporting

Create the Session

Log in to Prism Central (or Prism Element, depending on the exam environment, but Analysis is usually a PC feature).

Navigate to Operations -> Analysis.

Click New Session.

Name: `Monitor SQL02`

Entity: Search for and select the VM named `SQL02`.

Metrics: Since the issue is storage performance, search for and add these specific metrics:

Hypervisor IOPS (or Controller IOPS)

Hypervisor IO Latency (or Controller IO Latency)

Hypervisor IO Bandwidth

Click Save.

Save Session Data (Task 1.txt)

Open the "Monitor SQL02" session you just created.

(Per instructions): Right-click anywhere on the chart/data area -> Click **Select All**.

Copy the selected text (Ctrl+C).

Open Notepad on the provided desktop.

Paste the data.

Save the file as `Task 1.txt` on the Desktop.

Create and Schedule the Report

While still in the Analysis session, click the **Create Report** (or "Add to Report") button.

Report Name: `MonitorSQL02`

Report Settings:

Format: PDF

Frequency: Daily

Email Recipient: perf\_group@ACME.org

Retention: 0 (or "Do not retain", as requested).

Note: If the system forces you to create a new Report object and MonitorSQL02 is rejected, use monitorvm2 as the name per the instructions.

Save/Schedule the report.

Part 2: Diagnose and Fix the Issue

The Issue:

VM SQL02 was likely created with its data disks set to IDE or SATA.

Why this causes poor performance: IDE/SATA are emulated hardware with high CPU overhead and low queue depths (single-threaded).

The Standard: SQL01 (the healthy VM) is using SCSI, which is multithreaded and optimized for virtualization.

The Fix (Steps):

Navigate to the VM list in Prism

Select SQL02 and click Update (or Edit).

Scroll down to the Disks section.

Identify the data disk(s). You will see the Bus Type listed as IDE or SATA.

Do not delete the VM. instead, perform a disk conversion (destructive change to the disk is allowed, but we want to keep the data).

Method to Convert (Clone to SCSI):

Hover over the IDE/SATA disk to see the path/filename of the vDisk (or write it down).

Click Add New Disk.

Operation: select Clone from ADSF file.

Path: Browse to the storage container and select the file associated with the current IDE disk.

Bus Type: Select SCSI (This is the critical fix).

Index: Ensure it doesn't conflict with existing disks (usually index 1 or higher for data).

Click Add.

Once the new SCSI disk is added, find the original IDE/SATA disk and click the X to remove it.

Click Save.

Note: You do not need to power on the VM to verify. The change from IDE to SCSI allows the VM to use the Nutanix VirtIO drivers for maximum storage performance.

## NEW QUESTION # 17

Task 12

An administrator needs to create a report named VMs\_Power\_State that lists the VMs in the cluster and their basic details including the power state for the last month.

No other entities should be included in the report.

The report should run monthly and should send an email to admin@syberdyne.net when it runs.

Generate an instance of the report named VMs\_Power\_State as a CSV and save the zip file as

Desktop\Files\VMs\_Power\_state.zip Note: Make sure the report and zip file are named correctly. The SMTP server will not be configured.

### Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

To create a report named VMs\_Power\_State that lists the VMs in the cluster and their basic details including the power state for the last month, you can follow these steps:

Log in to Prism Central and click on Entities on the left menu.

Select Virtual Machines from the drop-down menu and click on Create Report.

Enter VMs\_Power\_State as the report name and a description if required. Click Next.

Under the Custom Views section, select Data Table. Click Next.

Under the Entity Type option, select VM. Click Next.

Under the Custom Columns option, add the following variables: Name, Cluster Name, vCPUs, Memory, Power State. Click Next.

Under the Time Period option, select Last Month. Click Next.

Under the Report Settings option, select Monthly from the Schedule drop-down menu. Enter admin@syberdyne.net as the Email Recipient. Select CSV as the Report Output Format. Click Next.

Review the report details and click Finish.

To generate an instance of the report named VMs\_Power\_State as a CSV and save the zip file as

Desktop\Files\VMs\_Power\_state.zip, you can follow these steps:

Log in to Prism Central and click on Operations on the left menu.

Select Reports from the drop-down menu and find the VMs\_Power\_State report from the list. Click on Run Now.

Wait for the report to be generated and click on Download Report. Save the file as Desktop\Files\VMs\_Power\_state.zip.

1. Open the Report section on Prism Central (Operations > Reports)
2. Click on the New Report button to start the creation of your custom report
3. Under the Custom Views section, select Data Table
4. Provide a title to your custom report, as well as a description if required.
5. Under the Entity Type option, select VM
6. This report can include all as well as a selection of the VMs
7. Click on the Custom Columns option and add the below variables:
  - a. Name - Name of the listed Virtual Machine
  - b. vCPUs - A combination of the vCores and vCPU's assigned to the Virtual Machine
  - c. Memory - Amount of memory assigned to the Virtual Machine
  - d. Disk Capacity - The total amount of assigned virtual disk capacity
  - e. Disk Usage - The total used virtual disk capacity
  - f. Snapshot Usage - The total amount of capacity used by snapshots (Excluding Protection Domain snapshots)
8. Under the Aggregation option for Memory and Disk Usage accept the default Average option
9. Click on the Add button to add this custom selection to your report
10. Next click on the Save and Run Now button on the bottom right of the screen
11. Provide the relevant details on this screen for your custom report:
12. You can leave the Time Period For Report variable at the default of Last 24 Hours
13. Specify a report output of preference (PDF or CSV) and if required Additional Recipients for this report to be mailed to. The report can also simply be downloaded after this creation and initial run if required
14. Below is an example of this report in a CSV format:

## NEW QUESTION # 18

Task 10

An administrator is working to create a VM using Nutanix V3 API calls with the following specifications.

\* VM specifications:

- \* vCPUs: 2
- \* Memory: 8Gb
- \* Disk Size: 50Gb
- \* Cluster: Cluster A
- \* Network: default- net

The API call is failing, indicating an issue with the payload:

The body is saved in Desktop/ Files/API\_Create\_VM.text

Correct any issues in the text file that would prevent from creating the VM. Also ensure the VM will be created as speeded and make sure it is saved for re-use using that filename.

Deploy the vm through the API

Note: Do not power on the VM.

### Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e000000LLEzCAO>

<https://jsonformatter.curiousconcept.com/>

acli net.list (uuid network default\_net)

ncli cluster info (uuid cluster)

Put Call: <https://Prism Central IP address : 9440/api/nutanix/v3/vms>

Edit these lines to fix the API call, do not add new lines or copy lines.

You can test using the Prism Element API explorer or PostMan

Body:

```
{  
  "spec": {  
    "name": "Test_Deploy",  
    "resources": {  
      "power_state": "OFF",  
      "num_vcpus_per_socket": ,  
      "memory_gb": 8,  
      "disk_gb": 50,  
      "cluster": "Cluster A",  
      "network": "default- net"  
    }  
  }  
}
```

```

"num_sockets": 1,
"memory_size_mib": 8192,
"disk_list": [
{
"disk_size_mib": 51200,
"device_properties": {
"device_type": "DISK"
}
},
{
"device_properties": {
"device_type": "CDROM"
}
}
],
"nic_list": [
{
"nic_type": "NORMAL_NIC",
"is_connected": true,
"ip_endpoint_list": [
{
"ip_type": "DHCP"
}
],
"subnet_reference": {
"kind": "subnet",
"name": "default_net",
"uuid": "00000000-0000-0000-0000-000000000000"
}
},
{
"cluster_reference": {
"kind": "cluster",
"name": "NTNXDemo",
"uuid": "00000000-0000-0000-0000-000000000000"
}
},
{
"api_version": "3.1.0",
"metadata": {
"kind": "vm"
}
}
]
},

```

<https://www.nutanix.dev/2019/08/26/post-a-package-building-your-first-nutanix-rest-api-post-request/> Reference

## NEW QUESTION # 19

Topic 1, Performance Based Questions

Environment

You have been provisioned a dedicated environment for your assessment which includes the following:

Workstation

\* windows Server 2019

\* All software/tools/etc to perform the required tasks

\* Nutanix Documentation and whitepapers can be found in desktop\files\Documentation

\* Note that the workstation is the system you are currently logged into Nutanix Cluster

\* There are three clusters provided. The connection information for the relevant cluster will be displayed to the right of the question

Please make sure you are working on the correct cluster for each item. Please ignore any licensing violations

\* Cluster A is a 3-node cluster with Prism Central 2022.6 where most questions will be performed

\* Cluster B is a one-node cluster and has one syslog item and one security item to perform

\* Cluster D is a one-node cluster with Prism Central 5.17 and has a security policy item to perform Important Notes

\* If the text is too small and hard to read, or you cannot see an of the GUI. you can increase/decrease the zoom of the browser with CTRL + ,and CTRL + (the plus and minus keys) You will be given 3 hours to complete the scenarios for Nutanix NCMMCI Once you click the start button below, you will be provided with:

- A Windows desktop A browser page with the scenarios and credentials (Desktop\instructions) Notes for this exam delivery:

The browser can be scaled to Improve visibility and fit all the content on the screen.

- Copy and paste hot-keys will not work Use your mouse for copy and paste.

- The Notes and Feedback tabs for each scenario are to leave notes for yourself or feedback for

- Make sure you are performing tasks on the correct components.

- Changing security or network settings on the wrong component may result in a failing grade.

- Do not change credentials on an component unless you are instructed to.

- All necessary documentation is contained in the Desktop\Files\Documentation directory Task 1 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 # 20

.....

Students are given a fixed amount of time to complete each test, thus Nutanix Exam Questions candidate's ability to control their time and finish the Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI) exam in the allocated time is a crucial qualification. Obviously, this calls for lots of practice. Taking Exams4Collection NCM-MCI Practice Exam helps you get familiar with the Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI) exam questions and work on your time management skills in preparation for the real Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI) exam.

**Exam NCM-MCI Registration:** <https://www.exams4collection.com/NCM-MCI-latest-braindumps.html>

- Get Excellent Marks in One Go with Nutanix NCM-MCI Real Dumps □ Enter ➔ [www.dumpsmaterials.com](http://www.dumpsmaterials.com) □ and search for ▷ NCM-MCI ▷ to download for free □ NCM-MCI New Study Guide
- 100% Pass Quiz Nutanix - NCM-MCI –High-quality Test Answers □ Search for 【 NCM-MCI 】 and download exam materials for free through ➔ [www.pdfvce.com](http://www.pdfvce.com) □ □ NCM-MCI Test Questions Pdf
- NCM-MCI Test Answers | Latest Nutanix NCM-MCI: Nutanix Certified Master - Multicloud Infrastructure v6.10 □

