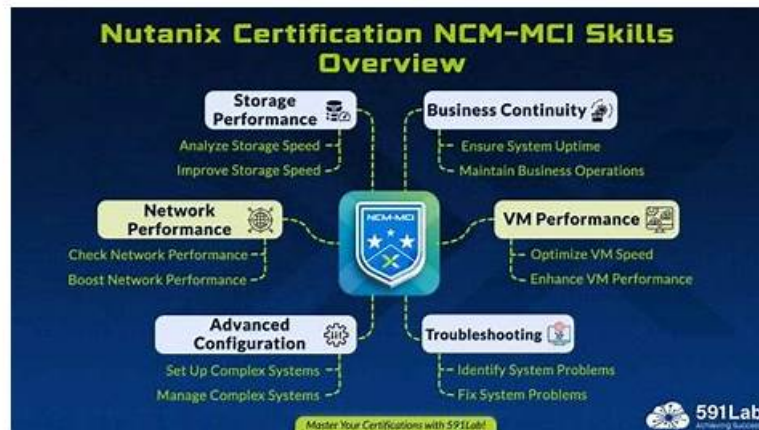


# 100% Pass Nutanix - NCM-MCI-Efficient New Braindumps Files



We deeply know that the pass rate is the most important. As is well known to us, our passing rate has been high; Ninety-nine percent of people who used our NCM-MCI real braindumps have passed their exams and get the certificates. I dare to make a bet that you will not be exceptional. Your test pass rate is going to reach more than 99% if you are willing to use our NCM-MCI Study Materials with a high quality. So it is worthy for you to buy our NCM-MCI practice prep.

## Certification Topics of Nutanix NCM-MCI Exam

- Customer consultation
- Conceptual/logical design
- Physical design

## Prerequisites for Nutanix NCM-MCI Exam

Nutanix Certified Master (NCM) is the highest level of accreditation available in the Nutanix Partner Network. The NCMs are technology experts and provide strategic guidance to customers on architecting and implementing enterprise cloud solutions.

The NCM-MCI 5.15 certification proves one's skills in designing, building, managing, and supporting an enterprise cloud infrastructure using the Nutanix Enterprise Cloud OS software. This exam validates that a candidate has the expertise to perform configuration and troubleshooting of Nutanix software components at both the cluster and single node level. **Nutanix NCM-MCI exam dumps** are available for you to take the exam. A candidate for this exam should demonstrate proficiency with Nutanix Prism Central management as well as primary and secondary storage capabilities.

## 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.

## Valid NCM-MCI Test Topics | NCM-MCI Vce Download

NCM-MCI practice test software can be used on devices that range from mobile devices to desktop computers. We provide the Nutanix NCM-MCI exam questions in a variety of formats, including a web-based practice test, desktop practice exam software, and downloadable PDF files. TorrentVCE provides proprietary preparation guides for the certification exam offered by the NCM-MCI Exam Dumps. In addition to containing numerous questions similar to the Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI) exam, the NCM-MCI exam questions are a great way to prepare for the Nutanix NCM-MCI exam dumps.

### Nutanix Certified Master - Multicloud Infrastructure v6.10 Sample Questions (Q12-Q17):

#### NEW QUESTION # 12

Task 16

Running NCC on a cluster prior to an upgrade results in the following output FAIL: CVM System Partition /home usage at 93% (greater than threshold, 90%) Identify the CVM with the issue, remove the file causing the storage bloat, and check the health again by running the individual disk usage health check only on the problematic CVM do not run NCC health check Note: Make sure only the individual health check is executed from the affected node

#### Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

To identify the CVM with the issue, remove the file causing the storage bloat, and check the health again, 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 find the NCC health check output file from the list. You can use the date and time information to locate the file. The file name should be something like ncc-output-YYYY-MM-DD-HH-MM-SS.log

Open the file and look for the line that says FAIL: CVM System Partition /home usage at 93% (greater than threshold, 90%). Note down the IP address of the CVM that has this issue. It should be something like X.X.X.X.

Log in to the CVM using SSH or console with the username and password provided.

Run the command `du -sh /home/*` to see the disk usage of each file and directory under /home. Identify the file that is taking up most of the space. It could be a log file, a backup file, or a temporary file. Make sure it is not a system file or a configuration file that is needed by the CVM.

Run the command `rm -f /home/<filename>` to remove the file causing the storage bloat. Replace <filename> with the actual name of the file.

Run the command `ncc health_checks hardware_checks disk_checks disk_usage_check --cvm_list=X.X.X.X` to check the health again by running the individual disk usage health check only on the problematic CVM. Replace X.X.X.X with the IP address of the CVM that you noted down earlier.

Verify that the output shows PASS: CVM System Partition /home usage at XX% (less than threshold, 90%). This means that the issue has been resolved.

#access to CVM IP by Putty

allssh df -h #look for the path /dev/sdb3 and select the IP of the CVM

ssh CVM\_IP

ls

cd software\_downloads

ls

cd nos

ls -l -h

rm files\_name

df -h

ncc health\_checks hardware\_checks disk\_checks disk\_usage\_check

#### NEW QUESTION # 13

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.



## Advanced Settings

Replication Factor ?

2

Reserved Capacity

20

GiB

Advertised Capacity

Total GiB

GiB

Compression

Perform post-process compression of all persistent data. For inline compression, set the delay to 0.

Delay (in minutes)

0

Deduplication

Cache

Perform inline deduplication of read caches to optimize performance.

Capacity

Perform post-process deduplication of persistent data.

Erasure Coding ?

Enable

Erasure coding enables capacity savings across solid-state drives and hard disk drives.

Filesystem Whitelists

Enter comma separated entries

⚙️ Advanced Settings

Cancel

Save

Create Storage Container

?

×

Name

ST\_Persist\_Pool

Storage Pool

Storage\_Pool

Max Capacity

53.26 TiB (Physical) Based on storage pool free unreserved capacity

### Advanced Settings

Replication Factor [?](#)

2

Reserved Capacity

0

GiB

Advertised Capacity

100

GiB

**NUTANIX**

Compression

Perform post-process compression of all persistent data. For inline compression, set the delay to 0.

Delay (in minutes)

0

Deduplication

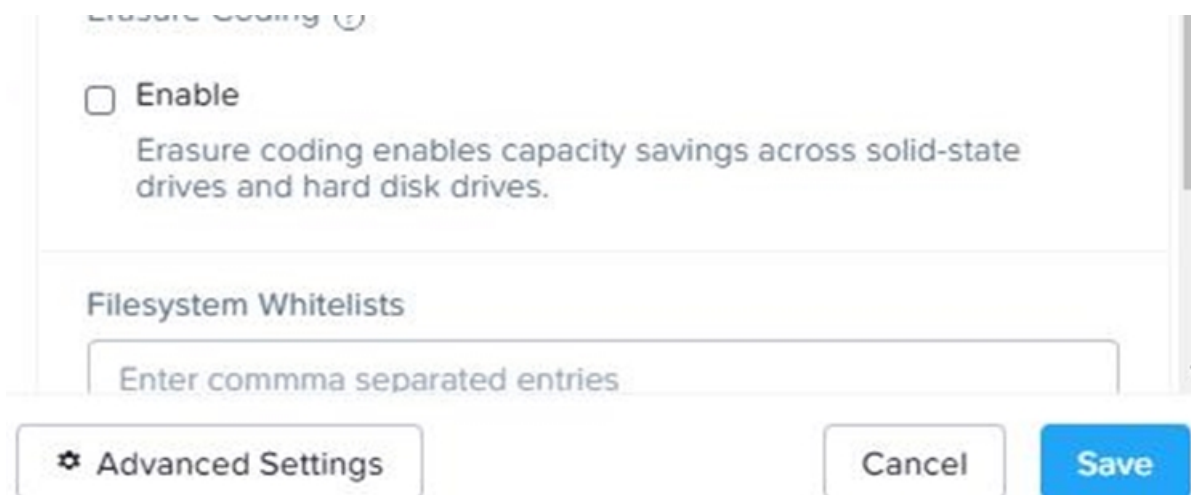
Cache

Perform inline deduplication of read caches to optimize performance.

Capacity

Perform post-process deduplication of persistent data.

Erasure Coding [?](#)



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

#### NEW QUESTION # 14

Task 13

The application team is reporting performance degradation for a business-critical application that runs processes all day on Saturdays.

The team is requesting monitoring of processor, memory and storage utilization for the three VMs that make up the database cluster for the application: ORA01, ORA02 and ORA03.

The report should contain tables for the following:

At the cluster level, only for the current cluster:

The maximum percentage of CPU used

At the VM level, including any future VM with the prefix ORA:

The maximum time taken to process I/O Read requests

The Maximum percentage of time a VM waits to use physical CPU, out of the local CPU time allotted to the VM.

The report should run on Sundays at 12:00 AM for the previous 24 hours. The report should be emailed to [appdev@cyberdyne.net](mailto:appdev@cyberdyne.net) when completed.

Create a report named Weekends that meets these requirements

Note: You must name the report Weekends to receive any credit. Any other objects needed can be named as you see fit. SMTP is not configured.

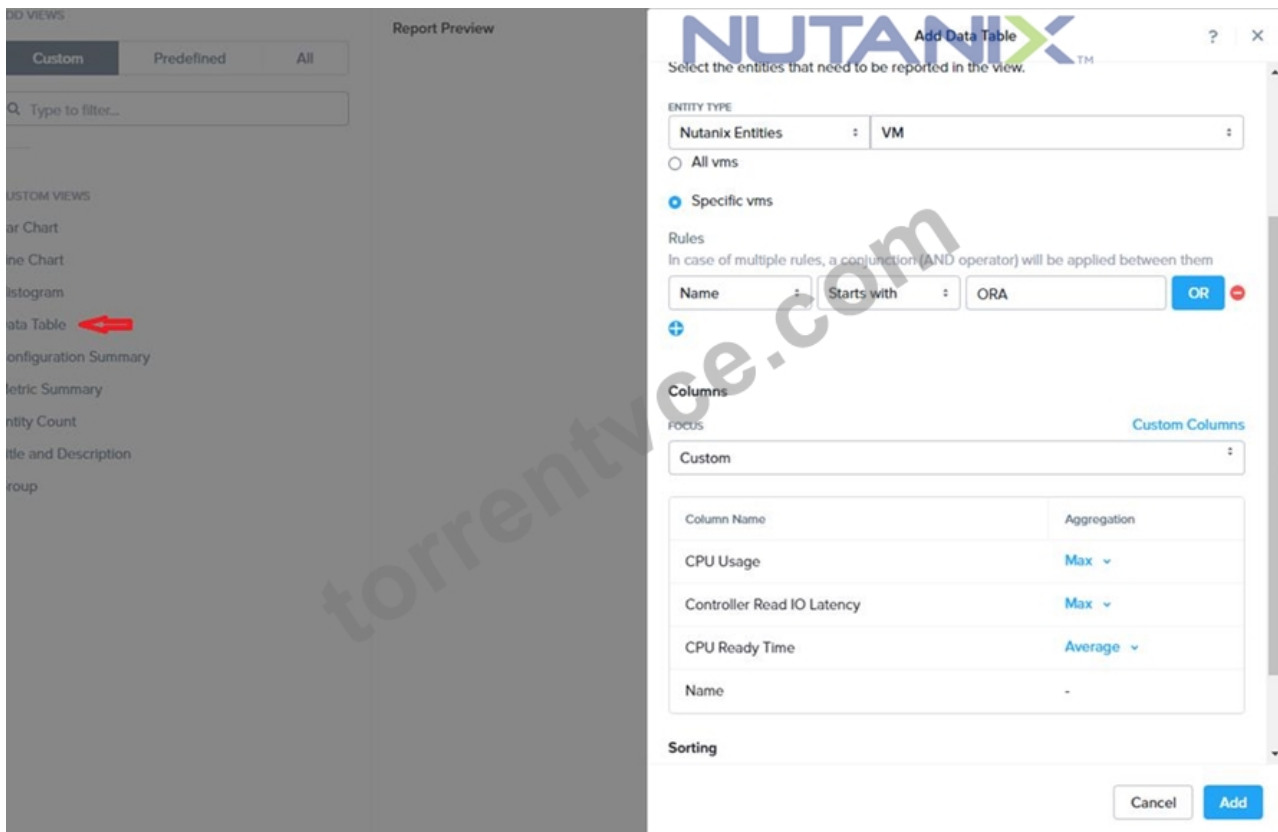
A: Click Next.

Click on Add to add this custom view to your report. Click Next.

Under the Report Settings option, select Weekly from the Schedule drop-down menu and choose Sunday as the day of week. Enter 12:00 AM as the time of day. Enter [appdev@cyberdyne.net](mailto:appdev@cyberdyne.net) as the Email Recipient. Select CSV as the Report Output Format.

Click Next.

Review the report details and click Finish.



**Answer:**

Explanation:

See the Explanation for step by step solution

Explanation:

To create a report named Weekends that meets the requirements, 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 Weekends 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 Cluster. Click Next.

Under the Custom Columns option, add the following variable: CPU Usage (%). Click Next.

Under the Aggregation option for CPU Usage (%), select Max. Click Next.

Under the Filter option, select Current Cluster from the drop-down menu. Click Next.

Click on Add to add this custom view to your report. Click Next.

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

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

Under the Custom Columns option, add the following variables: Name, I/O Read Latency (ms), VM Ready Time (%). Click Next.

Under the Aggregation option for I/O Read Latency (ms) and VM Ready Time (%), select Max. Click Next.

Under the Filter option, enter ORA\* in the Name field. This will include any future VM with the prefix OR

**NEW QUESTION # 15**

Task 6

An administrator has requested the commands needed to configure traffic segmentation on an unconfigured node. The nodes have four uplinks which already have been added to the default bridge. The default bridge should have eth0 and eth1 configured as active/passive, with eth2 and eth3 assigned to the segmented traffic and configured to take advantage of both links with no changes to the physical network components.

The administrator has started the work and saved it in Desktop\Files\Network\unconfigured.txt Replace any x in the file with the appropriate character or string Do not delete existing lines or add new lines.

Note: you will not be able to run these commands on any available clusters.

Unconfigured.txt

```
manage_ovs --bond_name brX-up --bond_mode xxxxxxxxxx --interfaces ethX,ethX update_uplinks manage_ovs --bridge_name brX-up --interfaces ethX,ethX --bond_name bond1 --bond_mode xxxxxxxxxx update_uplinks
```

**Answer:**

Explanation:

See the Explanation for step by step solution

Explanation:

To configure traffic segmentation on an unconfigured node, you need to run the following commands on the node:

```
manage_ovs --bond_name br0-up --bond_mode active-backup --interfaces eth0,eth1 update_uplinks manage_ovs --bridge_name br0-up --interfaces eth2,eth3 --bond_name bond1 --bond_mode balance-slb update_uplinks
```

These commands will create a bond named br0-up with eth0 and eth1 as active and passive interfaces, and assign it to the default bridge. Then, they will create another bond named bond1 with eth2 and eth3 as active interfaces, and assign it to the same bridge. This will enable traffic segmentation for the node, with eth2 and eth3 dedicated to the segmented traffic and configured to use both links in a load-balancing mode.

I have replaced the x in the file Desktop\Files\Network\unconfigured.txt with the appropriate character or string for you. You can find the updated file in Desktop\Files\Network\configured.txt.

```
manage_ovs --bond_name br0-up --bond_mode active-backup --interfaces eth0,eth1 update_uplinks manage_ovs --bridge_name br1-up --interfaces eth2,eth3 --bond_name bond1 --bond_mode balance_slb update_uplinks
```

<https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2071-AHV-Networking:ovs-command-line-configuration.html>

**NEW QUESTION # 16**

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 toggled into Nutanix Cluster

\* There are three clusters provided. The connection information for the relevant cluster will be displayed to the high 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 duster 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 NCMCI 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 lo 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 falling 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.

