

Exam Dumps NCM-MCI Zip & Relevant NCM-MCI Answers



2026 Latest Actual4Labs NCM-MCI PDF Dumps and NCM-MCI Exam Engine Free Share: <https://drive.google.com/open?id=1fDnjpzuQMbXn5MiDFwuuC-vQpeRD7P51>

With our outstanding NCM-MCI exam questions, we can assure you a 99% percent pass rate. Due to continuous efforts of our experts, we have exactly targeted the content of the NCM-MCI exam. You will pass the exam after 20 to 30 hours' learning with our NCM-MCI Study Material. Many users have witnessed the effectiveness of our NCM-MCI guide exam you surely will become one of them. Try it right now!

Nutanix NCM-MCI Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">• Business Continuity: The topic of business continuity measures knowledge about analyzing BCDR plans for compliance and evaluating BCDR plans for specific workloads.
Topic 2	<ul style="list-style-type: none">• Analyze and Optimize Storage Performance: It covers storage settings, workload requirements, and storage internals.
Topic 3	<ul style="list-style-type: none">• Analyze and Optimize VM Performance: Manipulation of VM configuration for resource utilization is discussed in this topic. It also explains interpreting VM, node, and cluster metrics.
Topic 4	<ul style="list-style-type: none">• Advanced Configuration and Troubleshooting: This topic covers sub-topics of executing API calls, configuring third-party integrations, analyzing AOS security posture, and translate business needs into technical solutions. Lastly, it discusses troubleshooting Nutanix services as well.
Topic 5	<ul style="list-style-type: none">• Analyze and Optimize Network Performance: Focal points of this topic are overlay networking, physical networks, virtual networks, network configurations, and flow policies. Moreover, questions about configurations also appear.

What is the exam cost of the Nutanix NCM-MCI Exam

The NCM-MCI certification exam cost is USD 199 for each attempt.

>> Exam Dumps NCM-MCI Zip <<

Relevant NCM-MCI Answers & NCM-MCI Relevant Answers

One of the most effective strategies to prepare for the Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI) exam successfully is to prepare with actual Nutanix NCM-MCI exam questions. It would be difficult for the candidates to pass the NCM-MCI exam on the first try if the NCM-MCI study materials they use are not updated. Studying with invalid NCM-MCI practice material results in a waste of time and money. Therefore, updated Nutanix NCM-MCI practice questions are essential for the preparation of the NCM-MCI exam.

Nutanix Certified Master - Multicloud Infrastructure v6.10 Sample Questions (Q16-Q21):

NEW QUESTION # 16

Task 15

An administrator found a CentOS VM, Cent_Down, on the cluster with a corrupted network stack. To correct the issue, the VM will need to be restored from a previous snapshot to become reachable on the network again.

VM credentials:

Username: root

Password: nutanix/4u

Restore the VM and ensure it is reachable on the network by pinging 172.31.0.1 from the VM.

Power off the VM before proceeding.

Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

To restore the VM and ensure it is reachable on the network, you can follow these steps:

Log in to the Web Console of the cluster where the VM is running.

Click on Virtual Machines on the left menu and find Cent_Down from the list. Click on the power icon to power off the VM.

Click on the snapshot icon next to the power icon to open the Snapshot Management window.

Select a snapshot from the list that was taken before the network stack was corrupted. You can use the date and time information to choose a suitable snapshot.

Click on Restore VM and confirm the action in the dialog box. Wait for the restore process to complete.

Click on the power icon again to power on the VM.

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

Run the command ping 172.31.0.1 to verify that the VM is reachable on the network. You should see a reply from the destination IP address.

Go to VMS from the prism central gui

Select the VM and go to More -> Guest Shutdown

Go to Snapshots tab and revert to latest snapshot available

power on vm and verify if ping is working

NEW QUESTION # 17

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 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 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 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 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 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.

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

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 xxxxxxxxxxxx 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 # 19

Task 2

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

Two VDI pools will be created

Non-persistent pool named 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.

Create Storage Container ? x

Name
ST_MCS_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
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

nutanix

Create Storage Container ? x

Name

Storage Pool

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

Advanced Settings

Replication Factor ?

Reserved Capacity
 GiB

Advertised Capacity
 GiB

Compression
 Perform post-process compression of all persistent data. For inline compression, set the delay to 0.
 Delay (in minutes)

Deduplication

Cache
 Perform inline deduplication of read caches to optimize performance.

Capacity
 Perform post-process deduplication of persistent data.

Erasur Coding ?

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

Filesystem Whitelists

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

NEW QUESTION # 20

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

Column Name	Aggregation
Name	-
vCPUs	-
Memory	Average ▾
Disk Capacity	-
Disk Usage	Average ▾
Snapshot Usage	-

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:

Run Report

Report

REPORT INSTANCE NAME

DESCRIPTION

TIME PERIOD FOR REPORT

Last 24 Hours

TIMEZONE

Report Format

PDF

CSV

NUTANIX
Email Report

Report will be emailed to the following recipients

-

ADDITIONAL RECIPIENTS

Cancel Run

12. You can leave the Time Period For Report variable at the default of Last 24 Hours

