

Certification NCM-MCI Exam Cost, NCM-MCI Dump File



NCM-MCI
Nutanix Certified
Master-
Multicloud
Infrastructure

Certification Questions
& Exams Dumps

www.edurely.com

What's more, part of that iPassleader NCM-MCI dumps now are free: <https://drive.google.com/open?id=1fz9bY-dJLEZhaTXQoqVQvF60oLa4ZMlq>

Do you want to get the NCM-MCI certification to boost your career? Do you desire to feel competent and confident going into your real Nutanix Certified Master - Multicloud Infrastructure v6.10 certification exam? Real NCM-MCI Exam Questions are available right here at iPassleader, so don't waste your time going elsewhere. By practicing with our Real NCM-MCI Exam Questions, which are offered in NCM-MCI PDF, web-based practice test, and desktop practice exam software formats, you can crack your Nutanix Certified Master - Multicloud Infrastructure v6.10 (NCM-MCI) certification test on first attempt and advance in the Nutanix industry.

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

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

What is the salary of the Nutanix NCM-MCI Exam

The Average salary of different countries of Nutanix Certified Expert

- United States - USD 116,178 per year
- India - INR 776,970 per year
- UK - Pounds 45,280 per year

>> Certification NCM-MCI Exam Cost <<

NCM-MCI Dump File, Certification NCM-MCI Exam Dumps

Our NCM-MCI practice exam simulator mirrors the NCM-MCI exam experience, so you know what to anticipate on NCM-MCI

exam day. Our Nutanix NCM-MCI features various question styles and levels, so you can customize your NCM-MCI exam questions preparation to meet your needs.

Nutanix Certified Master - Multicloud Infrastructure v6.10 Sample Questions (Q11-Q16):

NEW QUESTION # 11

Task 11

An administrator has noticed that after a host failure, the SQL03 VM was not powered back on from another host within the cluster. The Other SQL VMs (SQL01, SQL02) have recovered properly in the past.

Resolve the issue and configure the environment to ensure any single host failure affects a minimal number of SQL VMs.

Note: Do not power on any VMs

Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

One possible reason why the SQL03 VM was not powered back on after a host failure is that the cluster was configured with the default (best effort) VM high availability mode, which does not guarantee the availability of VMs in case of insufficient resources on the remaining hosts. To resolve this issue, I suggest changing the VM high availability mode to guarantee (reserved segments), which reserves some memory on each host for failover of VMs from a failed host. This way, the SQL03 VM will have a higher chance of being restarted on another host in case of a host failure.

To change the VM high availability mode to guarantee (reserved segments), you can follow these steps:

Log in to Prism Central and select the cluster where the SQL VMs are running.

Click on the gear icon on the top right corner and select Cluster Settings.

Under Cluster Services, click on Virtual Machine High Availability.

Select Guarantee (Reserved Segments) from the drop-down menu and click Save.

To configure the environment to ensure any single host failure affects a minimal number of SQL VMs, I suggest using anti-affinity rules, which prevent VMs that belong to the same group from running on the same host. This way, if one host fails, only one SQL VM will be affected and the other SQL VMs will continue running on different hosts.

To create an anti-affinity rule for the SQL VMs, 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 Group.

Enter a name for the group, such as SQL Group, and click Next.

Select the SQL VMs (SQL01, SQL02, SQL03) from the list and click Next.

Select Anti-Affinity from the drop-down menu and click Next.

Review the group details and click Finish.

I hope this helps. How else can I help?

https://portal.nutanix.com/page/documents/details?targetId=AHV-Admin-Guide-v6_5:ahv-affinity-policies-c.html

NEW QUESTION # 12

Task 3

An administrator needs to assess performance gains provided by AHV Turbo at the guest level. To perform the test the administrator created a Windows 10 VM named Turbo with the following configuration.

1 vCPU

8 GB RAM

SATA Controller

40 GB vDisk

The stress test application is multi-threaded capable, but the performance is not as expected with AHV Turbo enabled. Configure the VM to better leverage AHV Turbo.

Note: Do not power on the VM. Configure or prepare the VM for configuration as best you can without powering it on.

Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

To configure the VM to better leverage AHV Turbo, you can follow these steps:

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

Go to VM > Table and select the VM named Turbo.

Click on Update and go to Hardware tab.

Increase the number of vCPUs to match the number of multiqueues that you want to enable. For example, if you want to enable 8 multiqueues, set the vCPUs to 8. This will improve the performance of multi-threaded workloads by allowing them to use multiple processors.

Change the SCSI Controller type from SATA to VirtIO. This will enable the use of VirtIO drivers, which are required for AHV Turbo.

Click Save to apply the changes.

Power off the VM if it is running and mount the Nutanix VirtIO ISO image as a CD-ROM device. You can download the ISO image from Nutanix Portal.

Power on the VM and install the latest Nutanix VirtIO drivers for Windows 10. You can follow the instructions from Nutanix Support Portal.

After installing the drivers, power off the VM and unmount the Nutanix VirtIO ISO image.

Power on the VM and log in to Windows 10.

Open a command prompt as administrator and run the following command to enable multiqueue for the VirtIO NIC:

```
ethtool -L eth0 combined 8
```

Replace eth0 with the name of your network interface and 8 with the number of multiqueues that you want to enable. You can use ipconfig /all to find out your network interface name.

Restart the VM for the changes to take effect.

You have now configured the VM to better leverage AHV Turbo. You can run your stress test application again and observe the performance gains.

<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e000000LKPdCAO> change vCPU to 2/4 ?

Change SATA Controller to SCSI:

```
acli vm.get Turbo
```

Output Example:

```
Turbo {
  config {
    agent_vm: False
    allow_live_migrate: True
    boot {
      boot_device_order: "kCdrom"
      boot_device_order: "kDisk"
      boot_device_order: "kNetwork"
      uefi_boot: False
    }
    cpu_passthrough: False
    disable_branding: False
    disk_list {
      addr {
        bus: "ide"
        index: 0
      }
      cdrom: True
      device_uuid: "994b7840-dc7b-463e-a9bb-1950d7138671"
      empty: True
    }
    disk_list {
      addr {
        bus: "sata"
        index: 0
      }
    }
    container_id: 4
    container_uuid: "49b3e1a4-4201-4a3a-8abc-447c663a2a3e"
    device_uuid: "622550e4-fb91-49dd-8fc7-9e90e89a7b0e"
    naa_id: "naa.6506b8dcda1de6e9ce911de7d3a22111"
    storage_vdisk_uuid: "7e98a626-4cb3-47df-a1e2-8627cf90cae6"
    vmdisk_size: 10737418240
    vmdisk_uuid: "17e0413b-9326-4572-942f-68101f2bc716"
  }
  flash_mode: False
}
```

```

hwclock_timezone: "UTC"
machine_type: "pc"
memory_mb: 2048
name: "Turbo"
nic_list {
connected: True
mac_addr: "50:6b:8db2:a5:e4"
network_name: "network"
network_type: "kNativeNetwork"
network_uuid: "86a0d7ca-acfd-48db-b15c-5d654ff39096"
type: "kNormalNic"
uuid: "b9e3e127-966c-43f3-b33c-13608154c8bf"
vlan_mode: "kAccess"
}
num_cores_per_vcpu: 2
num_threads_per_core: 1
num_vcpus: 2
num_vnuma_nodes: 0
vga_console: True
vm_type: "kGuestVM"
}
is_rfl_vm: False
logical_timestamp: 2
state: "Off"
uuid: "9670901f-8c5b-4586-a699-41f0c9ab26c3"
}
acli vm.disk_create Turbo clone_from_vmdisk=17e0413b-9326-4572-942f-68101f2bc716 bus=scsi remove the old disk acli
vm.disk_delete 17e0413b-9326-4572-942f-68101f2bc716 disk_addr=sata.0

```

NEW QUESTION # 13

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 xxxxxxxxxxxx --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 # 14

TASK2

The security team has provided some new security requirements for cluster level security on Cluster 2.

Security requirements:

Update the password for the root user on the Cluster 2 node to match the admin user password.

Note: The 192.168.x.x network is not available. To access a node use the host IP (172.30.0.x) from the CVM.

Output the cluster-wide configuration of the SCMA policy to desktop\output.txt before changes are made.

Enable the Advanced Intrusion Detection Environment (AIDE) to run on a weekly basis for the hypervisor and cvms for Cluster 2.

Enable high-strength password policies for the hypervisor and cluster.

Ensure CVMs require SSH keys for login instead of passwords. (SSH keys are located in the desktop\Files\SSH folder.) Ensure the cluster meets these requirements. Do not reboot any cluster components.

Note: Please ensure you are modifying the correct components.

Answer:

Explanation:

See the Explanation

Explanation:

This task focuses on Security Technical Implementation Guides (STIGs) and general hardening of the Nutanix cluster. Most of these tasks are best performed via the Nutanix Command Line Interface (ncli) on the CVM, though the SSH key requirement is often easier to handle via the Prism GUI.

Here is the step-by-step procedure to complete Task 2.

Prerequisites: Connection

Open PuTTY (or the available terminal) from the provided Windows Desktop.

SSH into the Cluster 2 CVM. (If the Virtual IP is unknown, check Prism Element for the CVM IP).

Log in using the provided credentials (usually nutanix / nutanix/4u or the admin password provided in your instructions).

Step 1: Output SCMA Policy (Do this FIRST)

Requirement: Output the cluster-wide configuration of the SCMA policy to desktop\output.txt before changes are made.

In the SSH session on the CVM, run:

Bash

```
ncli cluster get-software-config-management-policy
```

Copy the output from the terminal window.

Open Notepad on the Windows Desktop.

Paste the output.

Save the file as output.txt on the Desktop.

Step 2: Enable AIDE (Weekly)

Requirement: Enable the Advanced Intrusion Detection Environment (AIDE) to run on a weekly basis for the hypervisor and CVMs.

In the same CVM SSH session, run the following command to modify the SCMA policy:

Bash

```
ncli cluster edit-software-config-management-policy enable-aide=true schedule-interval=WEEKLY (Note: This single command applies the policy to both Hypervisor and CVMs by default in most versions).
```

Step 3: Enable High-Strength Password Policies

Requirement: Enable high-strength password policies for the hypervisor and cluster.

Run the following command:

Bash

```
ncli cluster set-high-strength-password-policy enable=true
```

Step 4: Update Root Password for Cluster Nodes

Requirement: Update the password for the root user on the Cluster 2 node to match the admin user password.

Method A: The Automated Way (Recommended)

Use ncli to set the password for all hypervisor nodes at once without needing to SSH into them individually.

Run:

Bash

```
ncli cluster set-hypervisor-password
```

When prompted, enter the current admin password (this becomes the new root password).

Method B: The Manual Way (If NCLI fails or manual access is required)

Note: Use this if the exam specifically wants you to touch the node via the 172.x network.

From the CVM, SSH to the host using the internal IP:

Bash

```
ssh root@172.30.0.x (Replace x with the host ID, e.g., 4 or 5)
```

Run the password change command:

Bash

passwd

Enter the admin password twice.

Repeat for other nodes in Cluster 2.

Step 5: Cluster Lockdown (SSH Keys)

Requirement: Ensure CVMs require SSH keys for login instead of passwords.

It is safest to do this via the Prism Element GUI to prevent locking yourself out.

Open Prism Element for Cluster 2 in the browser.

Click the Gear Icon (Settings) -> Cluster Lockdown.

Uncheck the box "Enable Remote Login with Password".

Click New Public Key (or Add Key).

Open the folder Desktop\Files\SSH on the Windows desktop.

Open the public key file (usually ends in .pub) in Notepad and copy the contents.

Paste the key into the Prism "Key" box.

Click Save.

Note: Do not reboot the cluster. The SCMA and Password policies take effect immediately without a reboot.

NEW QUESTION # 15

Task 5

An administrator has been informed that a new workload requires a logically segmented network to meet security requirements.

Network configuration:

VLAN: 667

Network: 192.168.0.0

Subnet Mask: 255.255.255.0

DNS server: 34.82.231.220

Default Gateway: 192.168.0.1

Domain: cyberdyne.net

IP Pool: 192.168.9.100-200

DHCP Server IP: 192.168.0.2

Configure the cluster to meet the requirements for the new workload if new objects are required, start the name with 667.

Answer:

Explanation:

See the Explanation for step by step solution

Explanation:

To configure the cluster to meet the requirements for the new workload, you need to do the following steps:

Create a new VLAN with ID 667 on the cluster. You can do this by logging in to Prism Element and going to Network Configuration > VLANs > Create VLAN. Enter 667 as the VLAN ID and a name for the VLAN, such as 667_VLAN.

Create a new network segment with the network details provided. You can do this by logging in to Prism Central and going to Network > Network Segments > Create Network Segment. Enter a name for the network segment, such as

667_Network_Segment, and select 667_VLAN as the VLAN. Enter 192.168.0.0 as the Network Address and 255.255.255.0 as the Subnet Mask. Enter 192.168.0.1 as the Default Gateway and 34.82.231.220 as the DNS Server. Enter cyberdyne.net as the Domain Name.

Create a new IP pool with the IP range provided. You can do this by logging in to Prism Central and going to Network > IP Pools > Create IP Pool. Enter a name for the IP pool, such as 667_IP_Pool, and select 667_Network_Segment as the Network Segment. Enter 192.168.9.100 as the Starting IP Address and 192.168.9.200 as the Ending IP Address.

Configure the DHCP server with the IP address provided. You can do this by logging in to Prism Central and going to Network > DHCP Servers > Create DHCP Server. Enter a name for the DHCP server, such as 667_DHCP_Server, and select 667_Network_Segment as the Network Segment. Enter 192.168.0.2 as the IP Address and select 667_IP_Pool as the IP Pool.

□

NEW QUESTION # 16

.....

All purchases at iPassleader are protected by paypal system which is the most reliable payment system all over the world. So when you buy Nutanix NCM-MCI exam dumps, you won't worry about any leakage or mistakes during the deal. iPassleader puts

