

# NCM-MCI-6.5<sup>Q&As</sup>

Nutanix Certified Master - Multicloud Infrastructure (NCM-MCI)v6.5

## Pass NCM-MCI-6.5 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/ncm-mci-6-5.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Official  
Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



**QUESTION 1**

**CORRECT TEXT** 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

A. Answer: See the for step by step solution.

Correct Answer: A

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>

---

**QUESTION 2**

**CORRECT TEXT**

Task 7

An administrator has environment that will soon be upgraded to 6.5. In the meantime, they need to implement log and apply a security policy named Staging\_Production, such that not VM in the Staging Environment can communicate with any

VM in the production Environment,



Configure the environment to satisfy this requirement.

Note: All other configurations not indicated must be left at their default values.

A. Answer: See the for step by step solution.

Correct Answer: A

To configure the environment to satisfy the requirement of implementing a security policy named Staging\_Production, such that no VM in the Staging Environment can communicate with any VM in the production Environment, you need to do the following steps: Log in to Prism Central and go to Network > Security Policies > Create Security Policy. Enter Staging\_Production as the name of the security policy and select Cluster A as the cluster. In the Scope section, select VMs as the entity type and add the VMs that belong to the Staging Environment and the Production Environment as the entities. You can use tags or categories to filter the VMs based on their environment. In the Rules section, create a new rule with the following settings: Direction: Bidirectional Protocol: Any Source: Staging Environment Destination: Production Environment Action: Deny Save the security policy and apply it to the cluster. This will create a security policy that will block any traffic between the VMs in the Staging Environment and the VMs in the Production Environment. You can verify that the security policy is working by trying to ping or access any VM in the Production Environment from any VM in the Staging Environment, or vice versa. You should not be able to do so.



Name

Purpose

Isolate This Category

From This Category

☐ Apply the isolation only within a subset of the data center

Advanced Configuration  
 Policy Hit Logs ☐ Disabled



## QUESTION 3

### CORRECT TEXT

#### 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 `toappdev@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. Answer: See the for step by step solution.

Correct Answer: A

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 ORA. 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 as the Email Recipient. Select CSV as the Report Output Format.

Click Next.

Review the report details and click Finish.

**ADD VIEWS**

Custom Predefined All

Q Type to filter...

**CUSTOM VIEWS**

- Bar Chart
- Line Chart
- Histogram
- Data Table
- Configuration Summary
- Metric Summary
- Entity Count
- Title and Description
- Group

**Report Preview**

**Add Data Table**

Select the entities that need to be reported in the view.

ENTITY TYPE

Nutanix Entities : VM

☐ All vms

☒ Specific vms

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

Name : Starts with : ORA **OR**

**Columns**

FOCUS Custom Columns

Custom

Column Name	Aggregation
CPU Usage	Max
Controller Read IO Latency	Max
CPU Ready Time	Average
Name	-

**Sorting**

Cancel Add

## QUESTION 4

### CORRECT TEXT

#### Task 10

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

\*

VM specifications:

\*

vCPUs: 2

\*

Memory: BGb

\*

Disk Size: 50Gb

\*

Cluster: Cluster A

\*

Network: default- net

```

{
  "metadata": {
    "spec": {
      "message": "Request could not be processed.",
      "reason": "INVALID_REQUEST"
    }
  }
}

```

The API call is falling, 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.

A. Answer: See the for step by step solution.

Correct Answer: A

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

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

accli net.list(uuid network default\_net)

ncli cluster info(uuid cluster)

Put Call: <https://Prism Central IP address : 9440/api/nutanix/v3vms> 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": ,

"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

---

## QUESTION 5

### CORRECT TEXT

#### 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.

A. Answer: See the for step by step solution.

Correct Answer: A

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.

[Latest NCM-MCI-6.5 Dumps](#)   [NCM-MCI-6.5 Study Guide](#)   [NCM-MCI-6.5 Braindumps](#)