

NCM-MCI-6.5^{Q&As}

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 10

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

*

VM specifications:

*

vCPUs: 2

*

Memory: 8Gb

*

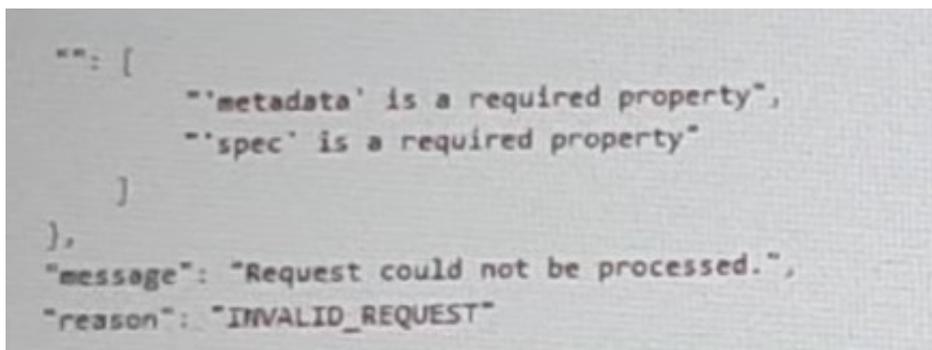
Disk Size: 50Gb

*

Cluster: Cluster A

*

Network: default- net



```
{}: {
  "'metadata' is a required property",
  "'spec' is a required property"
},
"message": "Request could not be processed.",
"reason": "INVALID_REQUEST"
```

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.

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

Correct Answer: A

<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e00000LLEzCAO>

<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/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 2

CORRECT TEXT Task 14 The application team has requested several mission-critical VMs to be configured for disaster recovery. The remote site (when added) will not be managed by Prism Central. As such, this solution should be built using the Web Console.

Disaster Recovery requirements per VM: Mkt01 RPO: 2 hours Retention: 5 snapshots Fin01 RPO: 15 minutes Retention: 7 days Dev01 RPO: 1 day Retention: 2 snapshots Configure a DR solution that meets the stated requirements. Any objects created in this item must start with the name of the VM being protected. Note: the remote site will be added later

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

Correct Answer: A

To configure a DR solution that meets the stated requirements, you can follow these steps:

Log in to the Web Console of the source cluster where the VMs are running. Click on Protection Domains on the left menu and click on Create Protection Domain. Enter a name for the protection domain, such as PD_Mkt01, and a description

if required.

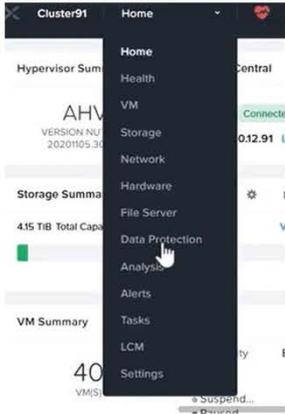
Click Next.

Select Mkt01 from the list of VMs and click Next. Select Schedule Based from the drop-down menu and enter 2 hours as the interval. Click Next.

Select Remote Site from the drop-down menu and choose the remote site where you want to replicate the VM. Click Next.

Enter 5 as the number of snapshots to retain on both local and remote sites. Click Next. Review the protection domain details and click Finish. Repeat the same steps for Fin01 and Dev01, using PD_Fin01 and PD_Dev01 as the protection

domain names, and adjusting the interval and retention values according to the requirements.



+ Protection Domain



A protection domain is a grouping of Virtual Machines for disaster recovery purposes. Enter a name (using alpha numeric characters only) for the protection domain you would like to create. You will then be guided into assigning Virtual Machines to it, and scheduling it.

Name

Protection Domain

Name **Entities** Schedule

Unprotected Entities (49) ?

Protected

Auto protect related entities. ?

Previous

Next

Auto protect related entities. ?

Protected Entities (1)

Search by Entity Name

Search by CG Name

<input type="checkbox"/>	Entity Name	CG
<input type="checkbox"/>	Mkt01	Mkt01
<input type="checkbox"/>		

Unprotect Selected Entities

Next

New Schedule

Protection Domain

? X

Name Entities Schedule

Configure your local schedule

Repeat every minute(s) ?

Repeat every hour(s) ?

Repeat every day(s) ?

Repeat weekly

S M T W T F S

Repeat monthly

Day of month: ?

Start on at

End on at

Retention policy

Local keep the last snapshots

Remote sites have not been defined for this cluster.

Create application consistent snapshots

Cancel Create Schedule

QUESTION 3

CORRECT TEXT 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.

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

Correct Answer: A

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-a1e2-8627cf90eae6" 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:8d:b2: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_rf1_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
```

QUESTION 4

CORRECT TEXT

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.

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

Correct Answer: A

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

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

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

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

QUESTION 5

CORRECT TEXT

Task 8

Depending on the order you perform the exam items, the access information and credentials could change. Please refer to the other item performed on Cluster B if you have problems accessing the cluster.

The infosec team has requested that audit logs for API Requests and replication capabilities be enabled for all clusters for the top 4 severity levels and pushed to their syslog system using highest reliability possible. They have requested no other logs to be included.

Syslog configuration:

Syslog Name: Corp_syslog

Syslog IP: 34.69.43.123

Port: 514

Ensure the cluster is configured to meet these requirements.

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

Correct Answer: A

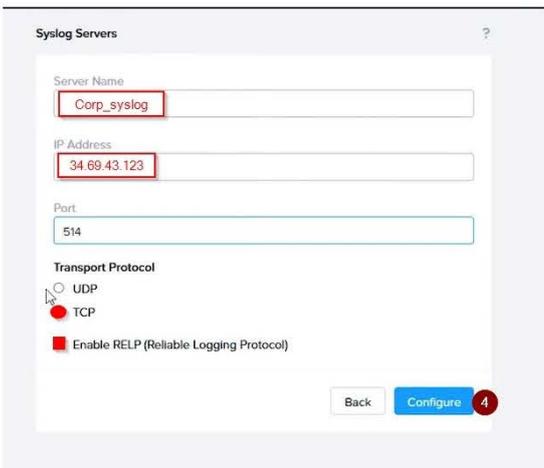
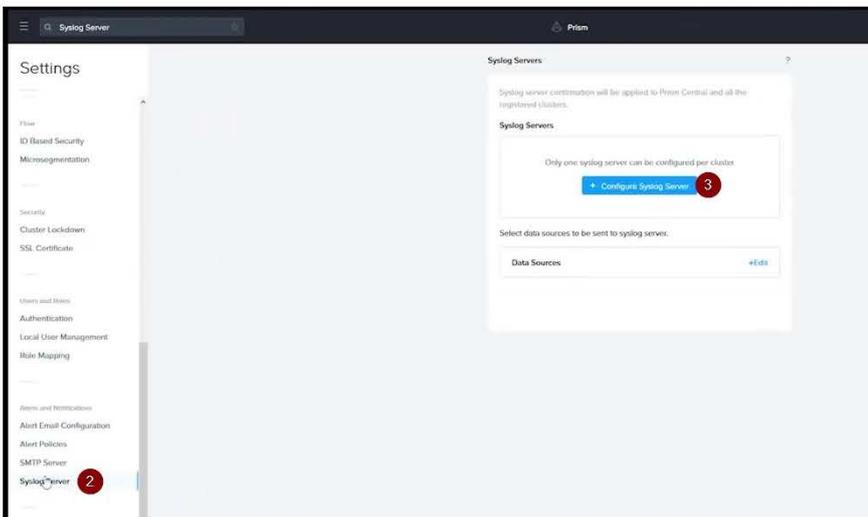
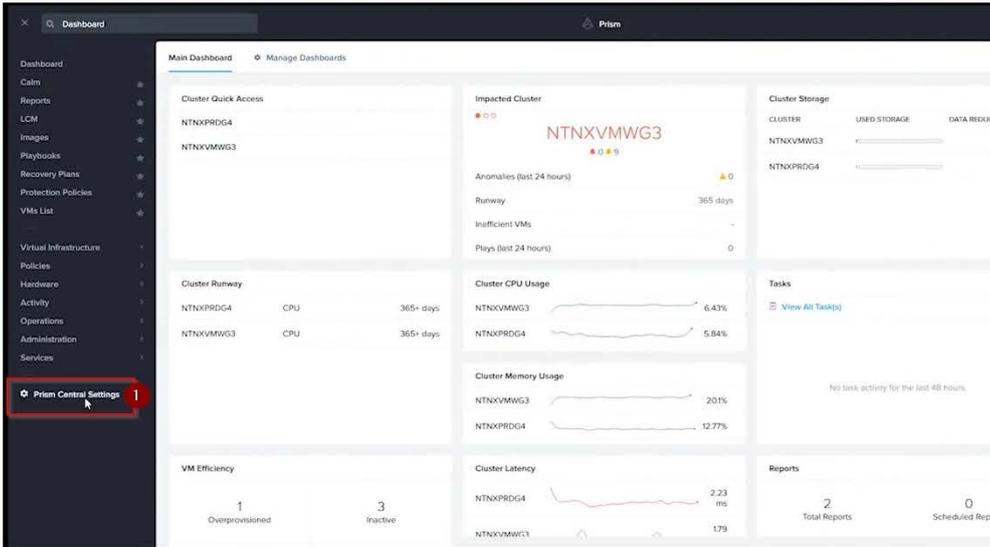
To configure the cluster to meet the requirements of the infosec team, you need to do the following steps:

Log in to Prism Central and go to Network > Syslog Servers > Configure Syslog Server. Enter Corp_syslog as the Server Name, 34.69.43.123 as the IP Address, and 514 as the Port. Select TCP as the Transport Protocol and enable RELP

(Reliable Logging Protocol). This will create a syslog server with the highest reliability possible. Click Edit against Data Sources and select Cluster B as the cluster. Select API Requests and Replication as the data sources and set the log level

to CRITICAL for both of them. This will enable audit logs for API requests and replication capabilities for the top 4 severity levels (EMERGENCY, ALERT, CRITICAL, and ERROR) and push them to the syslog server. Click Save.

Repeat step 2 for any other clusters that you want to configure with the same requirements.



Syslog Servers

Syslog server confirmation will be applied to Prism Central and all the registered clusters.

Syslog Servers +Configure Syslog Server

Name	Server IP
Corp_syslog	34.69.43.123

Select data sources to be sent to syslog server.

Data Sources +Edit **5**

Syslog Servers

Data Sources and Respective Severity Level

<input checked="" type="checkbox"/> Module Name	Severity Level
<input checked="" type="checkbox"/> API Audit	Select Severity Level
<input checked="" type="checkbox"/> Audit	0 - Emergency: system is unusable
<input checked="" type="checkbox"/> Flow	1 - Alert: action must be taken immediately
	2 - Critical: critical conditions
	3 - Error: error conditions
	4 - Warning: warning conditions
	5 - Notice: normal but significant condition
	6 - Informational: informational messages
	7 - Debug: debug-level messages

To configure the Nutanix clusters to enable audit logs for API Requests and replication capabilities, and push them to the syslog system with the highest reliability possible, you can follow these steps:

Log in to the Nutanix Prism web console using your administrator credentials. Navigate to the "Settings" section or the configuration settings interface within Prism. Locate the "Syslog Configuration" or "Logging" option and click on it.

Configure the syslog settings as follows:

Syslog Name: Enter "Corp_syslog" as the name for the syslog configuration. **Syslog IP:** Set the IP address to "34.69.43.123", which is the IP address of the syslog system.

Port: Set the port to "514", which is the default port for syslog. Enable the option for highest reliability or persistent logging, if available. This ensures that logs are sent reliably and not lost in case of network interruptions.

Save the syslog configuration.

Enable Audit Logs for API Requests:

In the Nutanix Prism web console, navigate to the "Cluster" section or the cluster management interface.

Select the desired cluster where you want to enable audit logs. Locate the "Audit Configuration" or "Security Configuration" option and click on it. Look for the settings related to audit logs and API requests. Enable the audit logging feature and

select the top 4 severity levels to be logged.

Save the audit configuration.

Enable Audit Logs for Replication Capabilities:

In the Nutanix Prism web console, navigate to the "Cluster" section or the cluster management interface.

Select the desired cluster where you want to enable audit logs. Locate the "Audit Configuration" or "Security Configuration" option and click on it. Look for the settings related to audit logs and replication capabilities. Enable the audit logging

feature and select the top 4 severity levels to be logged.

Save the audit configuration.

After completing these steps, the Nutanix clusters will be configured to enable audit logs for API Requests and replication capabilities. The logs will be sent to the specified syslog system with the highest reliability possible.

ncli

```
rsyslog-config set-status enable=false
```

```
rsyslog-config add-server name=Corp_Syslog ip-address=34.69.43.123 port=514 network-protocol=tdp relp-enabled=false
```

```
rsyslog-config add-module server-name= Corp_Syslog module-name=APLOS level=INFO
```

```
rsyslog-config add-module server-name= Corp_Syslog module-name=CEREBRO level=INFO
```

```
rsyslog-config set-status enable=true
```

<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e0000009CEECA2>

[NCM-MCI-6.5 PDF Dumps](#)

[NCM-MCI-6.5 Exam
Questions](#)

[NCM-MCI-6.5 Braindumps](#)