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.pass4itsure.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 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 fromNutanix Portal. Power on the VM and install the latest Nutanix VirtIO drivers for Windows 10. You can

follow the instructions fromNutanix 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

Replaceeth0with the name of your network interface and8with the number of multiqueues that you want to enable. You can useipconfig /allto 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 2

CORRECT TEXT

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.

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

Correct Answer: A

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 VMand go to More -> Guest Shutdown

Go to Snapshots tab and revert to latest snapshot available power on vm and verify if ping is working

QUESTION 3

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 Replacle 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

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 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, seeCitrix Virtual Apps and Desktops on NutanixorNutanix

virtualization environments.

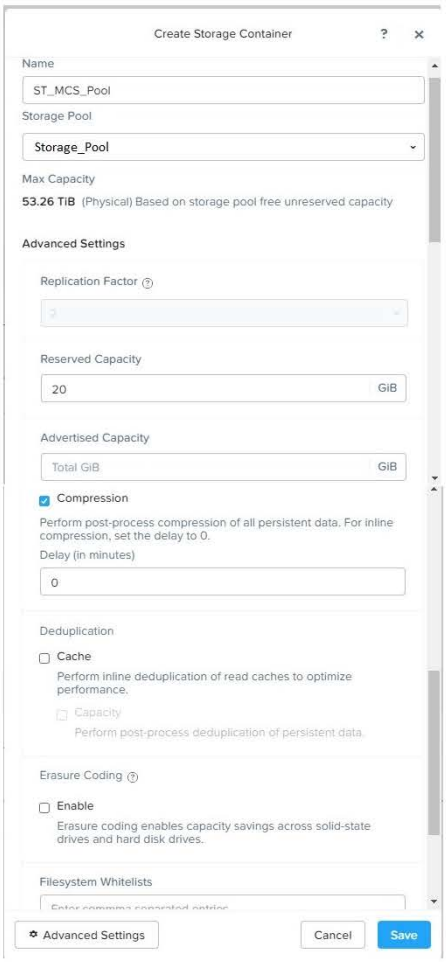
Create Storage Container
Name
ST_MCS_Pool
Storage Pool
Storage_Pool
Max Capacity
53.26 TiB (Physical) Based on storage pool free unreserved capacity
Advanced Settings
Replication Factor
Reserved Capacity
20
GiB
Advertised Capacity
Total GiB
GiB
Compression
Perform post-process compression of all persistent 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
Advanced Settings
Cancel
Save

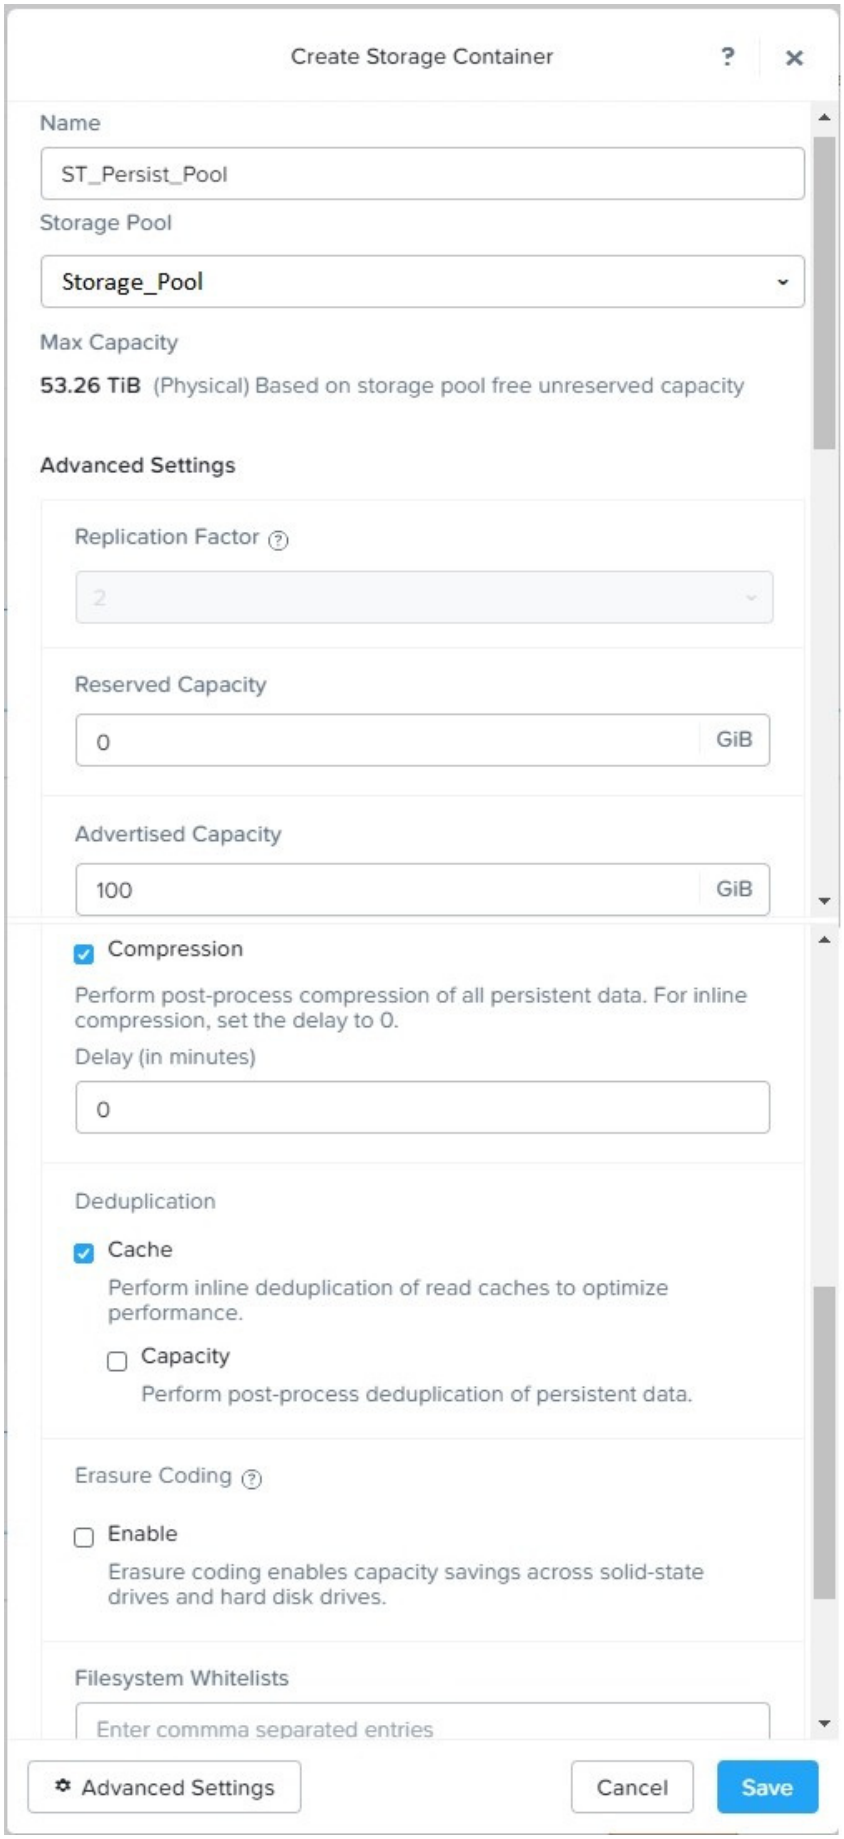
Create Storage Container
Name
ST_Persist_Pool
Storage Pool
Storage_Pool
Max Capacity
53.26 TiB (Physical) Based on storage pool free unreserved capacity
Advanced Settings
Replication Factor
2
Reserved Capacity
0
GiB
Advertised Capacity
100
GiB
Compression
Perform post-process compression of all persistent 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 commma separated entries
Advanced Settings
Cancel
Save

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 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.

NCM-MCI-6.5 PDF Dumps NCM-MCI-6.5 VCE Dumps NCM-MCI-6.5 Exam Questions