

Nutanix NCM-MCI-6.5 Exam Question & Answers
Nutanix Certified Master - Multicloud Infrastructure
(NCM-MCI) v6.5 Exam

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 toggled 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 duster 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 lo 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

Question: 1

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: 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.

a. 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.

Question: 2

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: 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](#).

