

NCM-MCI認證考試資料庫



VCESoft 考題大師的擬真試題覆蓋了真實的考試真題，已經成為考生通過 Nutanix NCM-MCI 考試的首選學習資料。NCM-MCI 考試主要用於具有較高水準的實施顧問能力，獲取證書，以確保考生有一個堅實的專業基礎知識，有利於他們將此能力企業專業化。準備 Nutanix 的 NCM-MCI 考試的考生，需要熟練了解我們的擬真試題，快速完成測試，就能順利通過考試。

Nutanix NCM-MCI 考試大綱：

主題	簡介
主題 1	<ul style="list-style-type: none">• Advanced Configuration and Troubleshooting: This topic covers sub-topics of executing API calls, configuring third-party integrations, analyzing AOS security posture, and translate business needs into technical solutions. Lastly, it discusses troubleshooting Nutanix services as well.
主題 2	<ul style="list-style-type: none">• Analyze and Optimize VM Performance: Manipulation of VM configuration for resource utilization is discussed in this topic. It also explains interpreting VM, node, and cluster metrics.
主題 3	<ul style="list-style-type: none">• Analyze and Optimize Storage Performance: It covers storage settings, workload requirements, and storage internals.
主題 4	<ul style="list-style-type: none">• Analyze and Optimize Network Performance: Focal points of this topic are overlay networking, physical networks, virtual networks, network configurations, and flow policies. Moreover, questions about configurations also appear.
主題 5	<ul style="list-style-type: none">• Business Continuity: The topic of business continuity measures knowledge about analyzing BCDR plans for compliance and evaluating BCDR plans for specific workloads.

>> NCM-MCI證照指南 <<

NCM-MCI熱門證照，新版NCM-MCI題庫

在你還在猶豫選擇我們VCESoft之前，你可以先嘗試在我們VCESoft免費下載我們為你提供的關於Nutanix NCM-MCI認證考試的部分考題及答案。這樣，你就可以知道我們VCESoft的可靠性。我們VCESoft也會是你通過Nutanix NCM-MCI認證考試最好的選擇，我們VCESoft是你通過Nutanix NCM-MCI認證考試最好的保證。你選擇了我們VCESoft，就等於選擇了成功。

最新的 Master Level NCM-MCI 免費考試真題 (Q16-Q21):

問題 #16

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

答案:

解題說明:

See the Explanation for step by step solution

Explanation:

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.



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

Mkt01-PD

Protection Domain

Name Entities **TM** Schedule

Unprotected Entities (49) ?

Mkt01

Protected

Search b

Auto protect related entities. ?

Protect Selected Entities (1) >

Previous

Next

NUTANIX

Auto protect related entities. ?

Protect Selected Entities (1) <

Protected Entities (1)

Search by Entity Name

Search by CG Name

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

Unprotect Selected Entities <

NUTANIX

New Schedule

NUTANIX Protection Domain

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

Create application consistent snapshots

Retention policy

Local keep the last snapshots

Remote sites have not been defined for this cluster.

Cancel Create Schedule

問題 #17

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 logged 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 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.

答案:

解題說明:

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. 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.

問題 #18

Task4

An administrator will be deploying Flow Networking and needs to validate that the environment, specifically switch vs1, is appropriately configured. Only VPC traffic should be carried by the switch.

Four versions each of two possible commands have been placed in Desktop\Files\Network\flow.txt. Remove the hash mark (#) from the front of correct First command and correct Second command and save the file.

Only one hash mark should be removed from each section. Do not delete or copy lines, do not add additional lines. Any changes other than removing two hash marks (#) will result in no credit.

Also, SSH directly to any AHV node (not a CVM) in the cluster and from the command line display an overview of the Open vSwitch configuration. Copy and paste this to a new text file named Desktop\Files\Network\AHVswitch.txt.

Note: You will not be able to use the 192.168.5.0 network in this environment.

First command

```
#net.update_vpc_traffic_config virtual_switch=vs0
net.update_vpc_traffic_config virtual_switch=vs1
#net.update_vpc_east_west_traffic_config virtual_switch=vs0
#net.update_vpc_east_west_traffic_config virtual_switch=vs1
```

Second command

```
#net.update_vpc_east_west_traffic_config permit_all_traffic=true
net.update_vpc_east_west_traffic_config permit_vpc_traffic=true
#net.update_vpc_east_west_traffic_config permit_all_traffic=false
#net.update_vpc_east_west_traffic_config permit_vpc_traffic=false
```

答案:

解題說明:

See the Explanation for step by step solution

Explanation:

First, you need to open the Prism Central CLI from the Windows Server 2019 workstation. You can do this by clicking on the Start menu and typing "Prism Central CLI". Then, you need to log in with the credentials provided to you.

Second, you need to run the two commands that I have already given you in Desktop\Files\Network\flow.txt. These commands are: `net.update_vpc_traffic_config virtual_switch=vs1` `net.update_vpc_east_west_traffic_config permit_vpc_traffic=true` These commands will update the virtual switch that carries the VPC traffic to vs1, and update the VPC east-west traffic configuration to allow only VPC traffic. You can verify that these commands have been executed successfully by running the command:

```
net.get_vpc_traffic_config
```

This command will show you the current settings of the virtual switch and the VPC east-west traffic configuration.

Third, you need to SSH directly to any AHV node (not a CVM) in the cluster and run the command:

```
ovs-vsctl show
```

This command will display an overview of the Open vSwitch configuration on the AHV node. You can copy and paste the output of this command to a new text file named Desktop\Files\Network\AHVswitch.txt.

You can use any SSH client such as PuTTY or Windows PowerShell to connect to the AHV node. You will need the IP address and the credentials of the AHV node, which you can find in Prism Element or Prism Central.

remove # from greens

On AHV execute:

```
sudo ovs-vsctl show
```

CVM access AHV access command

```
nutanix@NTNX-A-CVM:192.168.10.5:~$ ssh root@192.168.10.2 "ovs-vsctl show" Open AHVswitch.txt and copy paste output
```

問題 #19

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 `appdev@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.

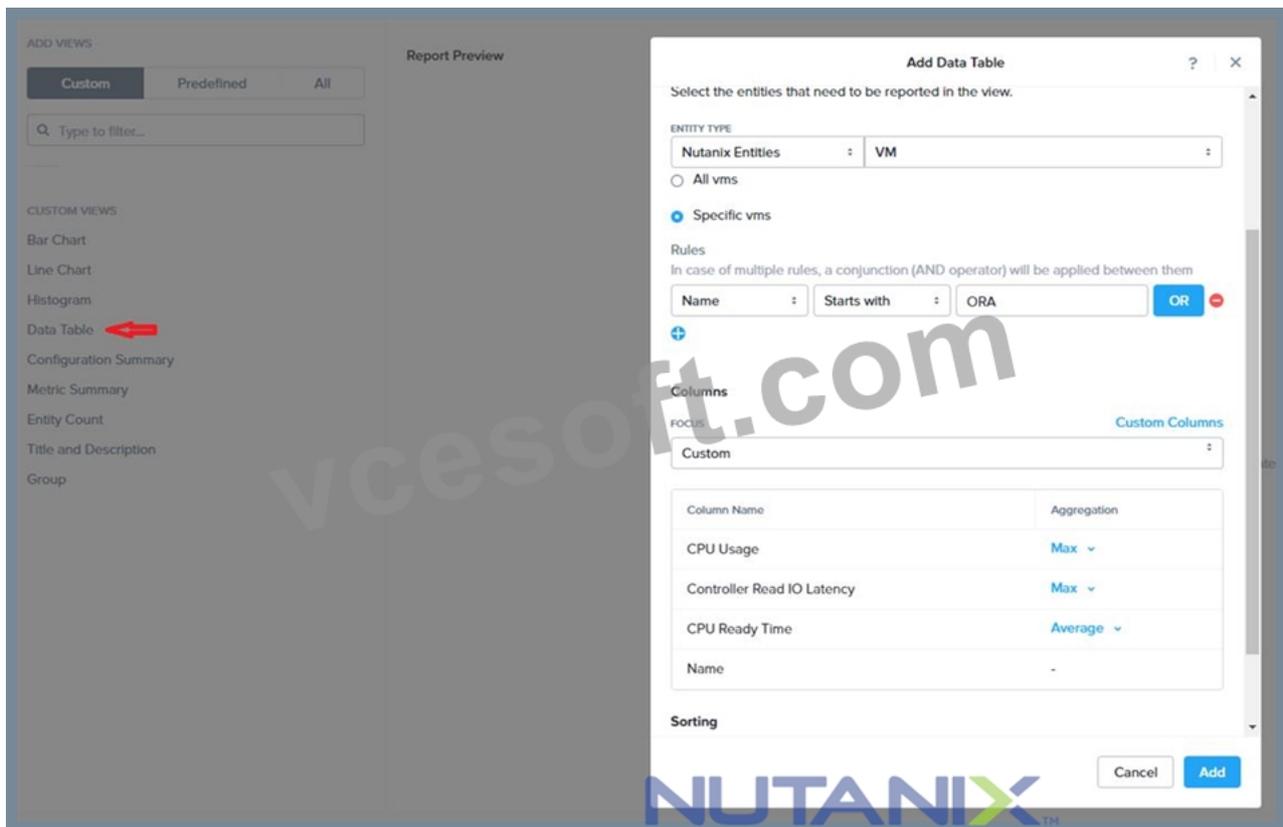
A: Click Next.

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.



答案:

解題說明:

See the Explanation for step by step solution

Explanation:

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 OR

問題 #20

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.

答案：

解題說明：

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.

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

NUTANIX 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.

Erasur Coding ⓘ

Enable

Erasur 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
Erasure coding enables capacity savings across solid-state drives and hard disk drives.

Filesystem Whitelists
Enter comma separated entries

Buttons: 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>

問題 #21

.....

選擇了VCESoft提供的最新最準確的關於Nutanix NCM-MCI考試產品，屬於你的成功就在不遠處。

NCM-MCI熱門證照: <https://www.vcesoft.com/NCM-MCI-pdf.html>

- NCM-MCI考證 □ NCM-MCI最新考古題 □ NCM-MCI題庫下載 □ 打開網站▷ www.kaoguti.com ◁搜索▷ NCM-MCI ◁免費下載最新NCM-MCI題庫資源
- NCM-MCI考證 □ NCM-MCI考證 □ NCM-MCI最新考古題 □ 在【 www.newdumpsdf.com 】網站下載免費▷ NCM-MCI ◁題庫收集NCM-MCI考古題更新
- NCM-MCI測試題庫 □ NCM-MCI認證考試 □ NCM-MCI考試證照綜述 □ 打開網站“www.vcesoft.com”搜索⇒ NCM-MCI ◁免費下載NCM-MCI考試證照綜述
- 受信任的Nutanix NCM-MCI: Nutanix Certified Master - Multicloud Infrastructure v6.10證照指南 - 最新的Newdumpsdf.com NCM-MCI熱門證照 □ ▷ www.newdumpsdf.com ◁上搜索▷ NCM-MCI □輕鬆獲取免費下載NCM-MCI測試題庫
- 最新版的NCM-MCI證照指南，免費下載NCM-MCI學習資料得到妳想要的Nutanix證書 □ 打開網站「www.newdumpsdf.com」搜索“NCM-MCI”免費下載NCM-MCI測試題庫

