

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 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 12

An administrator needs to create a report named VMs_Power_State that lists the VMs in the cluster and their basic details including the power state for the last month.

No other entities should be included in the report.

The report should run monthly and should send an email toadmin@syberdyne.netwhen it runs.

Generate an instance of the report named VMs_Power_State as a CSV and save the zip file as Desktop\Files\VMs_Power_state.zip

Note: Make sure the report and zip file are named correctly. The SMTP server will not be configured.

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

Correct Answer: A

To create a report named VMs_Power_State that lists the VMs in the cluster and their basic details including the power state for the last month, 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 VMs_Power_State 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 VM. Click Next. Under the Custom Columns option, add the following variables: Name, Cluster Name, vCPUs, Memory, Power State. Click Next.

Under the Time Period option, select Last Month. Click Next. Under the Report Settings option, select Monthly from the Schedule drop-down menu. Enter admin@syberdyne.net as the Email Recipient. Select CSV as the Report Output

Format. Click Next.

Review the report details and click Finish.

To generate an instance of the report named VMs_Power_State as a CSV and save the zip file as Desktop\Files\VMs_Power_state.zip, you can follow these steps:

Log in to Prism Central and click on Operations on the left menu. Select Reports from the drop-down menu and find the VMs_Power_State report from the list. Click on Run Now.

Wait for the report to be generated and click on Download Report. Save the file as Desktop\Files\VMs_Power_state.zip.

1.Open the Report section on Prism Central (Operations > Reports) 2.Click on the New Report button to start the creation of your custom report 3.Under the Custom Views section, select Data Table 4.Provide a title to your custom report, as

well as a description if required.

5.Under the Entity Type option, select VM

6.This report can include all as well as a selection of the VMs 7.Click on the Custom Columns option and add the below variables:

a.Name - Name of the listed Virtual Machine

b.vCPUs - A combination of the vCores and vCPU's assigned to the Virtual Machine c.Memory - Amount of memory assigned to the Virtual Machine d.Disk Capacity - The total amount of assigned virtual disk capacity e.Disk Usage - The total

used virtual disk capacity f.Snapshot Usage - The total amount of capacity used by snapshots (Excluding Protection Domain snapshots)

8.Under the Aggregation option for Memory and Disk Usage accept the default Average option

Columns

FOCUS

Custom Columns

Custom
⌵

Column Name	Aggregation
Name	-
vCPUs	-
Memory	Average ▾
Disk Capacity	-
Disk Usage	Average ▾
Snapshot Usage	-

9. Click on the Add button to add this custom selection to your report 10. Next click on the Save and Run Now button on the bottom right of the screen 11. Provide the relevant details on this screen for your custom report:

12. You can leave the Time Period For Report variable at the default of Last 24 Hours 13. Specify a report output of preference (PDF or CSV) and if required Additional Recipients for this report to be mailed to. The report can also simply be

downloaded after this creation and initial run if required

14. Below is an example of this report in a CSV format:

QUESTION 3

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 4

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

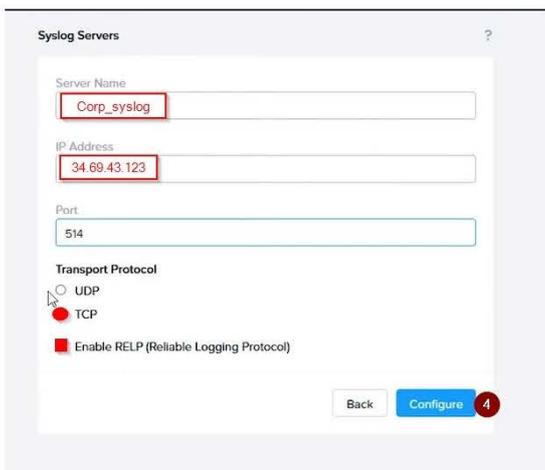
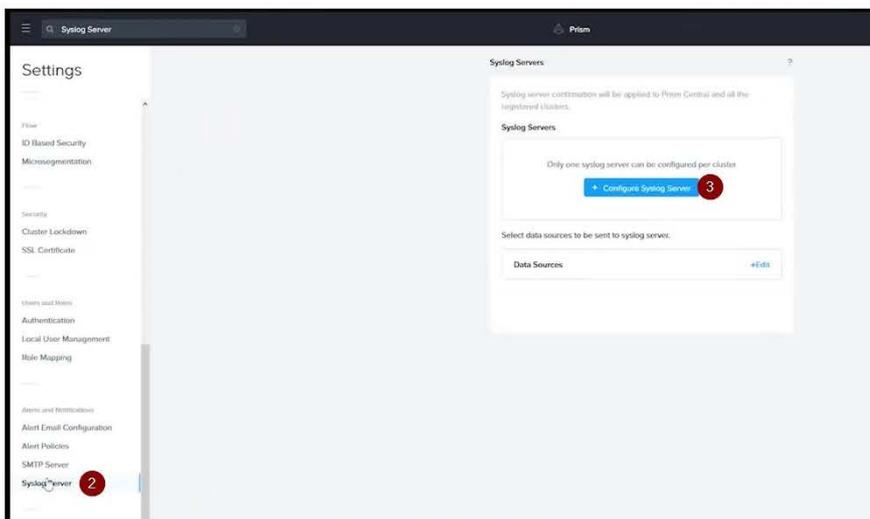
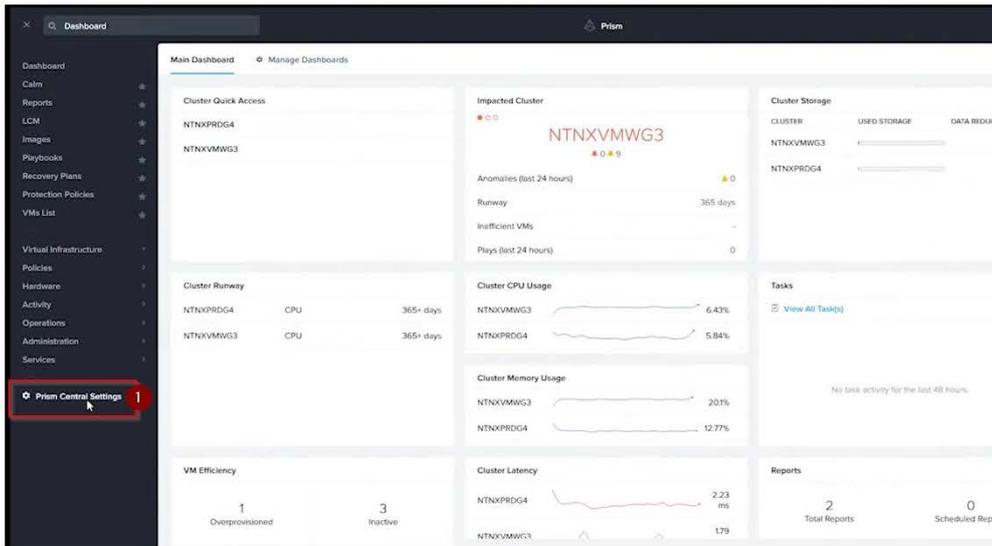
To configure the cluster to meet the requirements of the infosec team, you need to do the following steps:

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 0 - Emergency: system is unusable 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
<input checked="" type="checkbox"/> Audit	
<input checked="" type="checkbox"/> Flow	

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>

QUESTION 5

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.

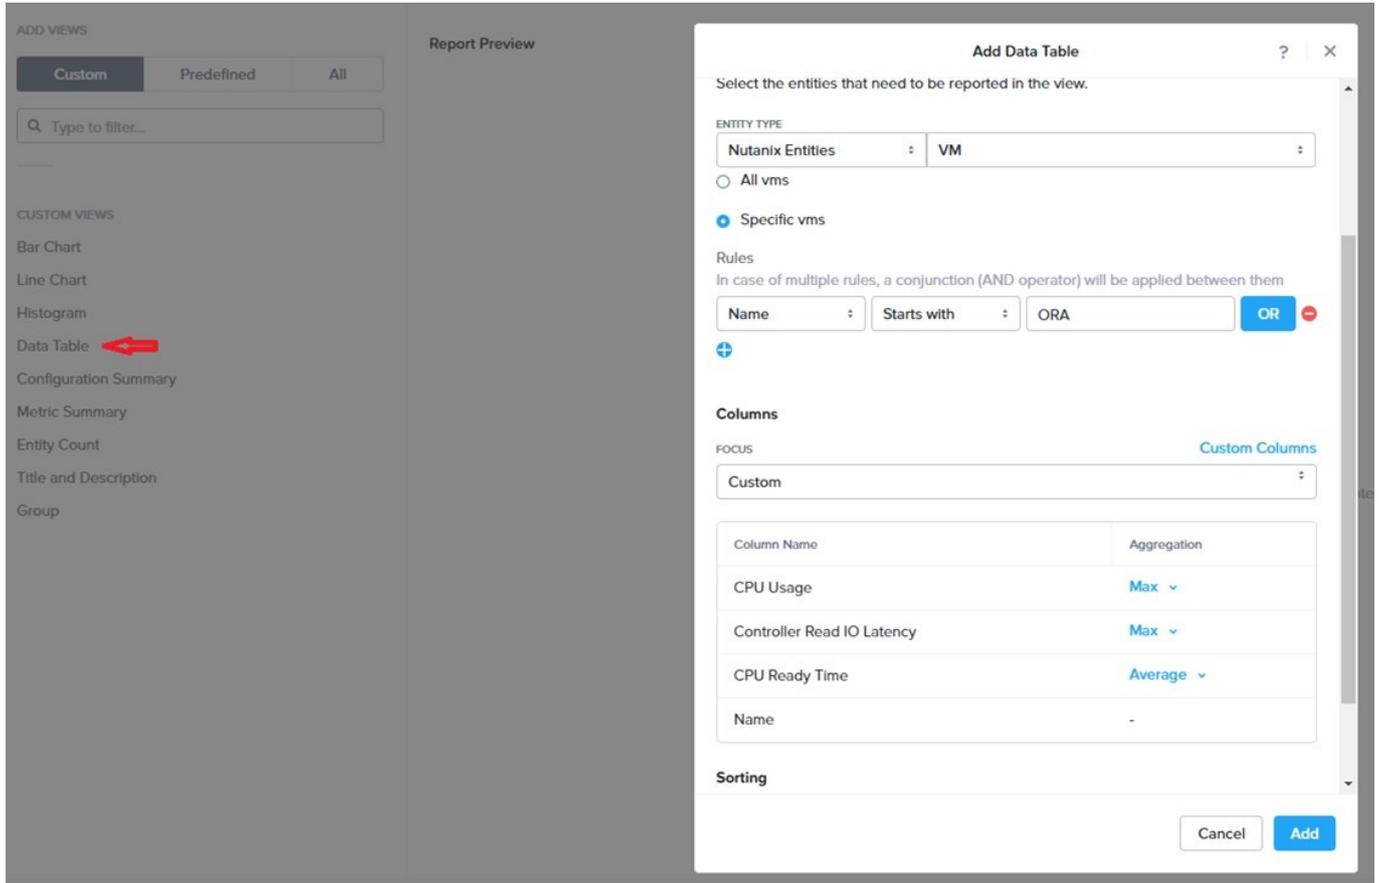
Under the Filter option, enter ORA* in the Name field. This will include any future VM with the prefix ORA. Click Next.

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.



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