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.geekcert.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: BGb

*

Disk Size: 50Gb

*

Cluster: Cluster A

*

Network: default- net

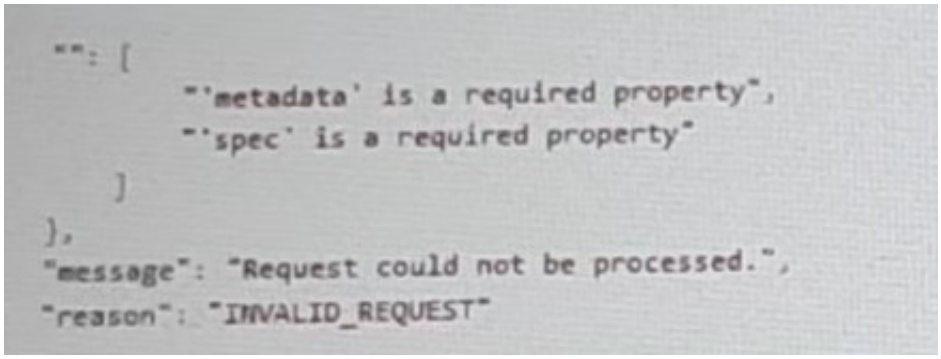

The API call is falling, indicating an issue with the payload:

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=kA00e000000LLEzCAO
https://jsonformatter.curiousconcept.com/#

acli net.list(uuid network defult_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 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.

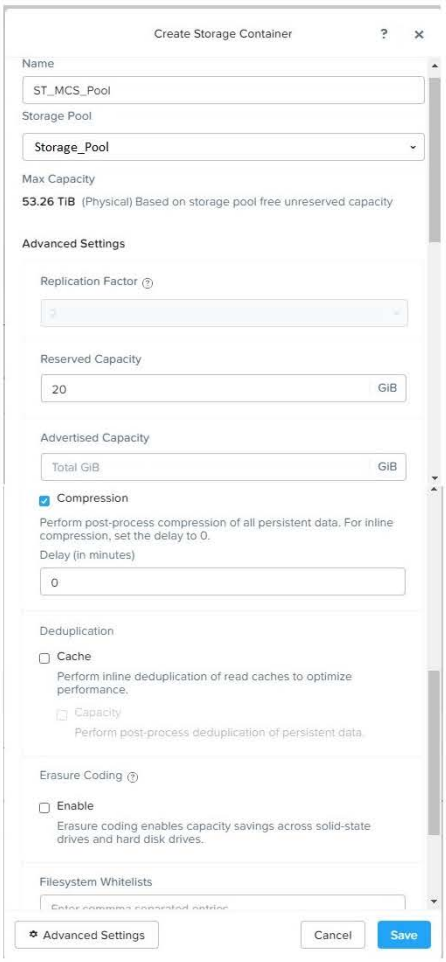
Create Storage Container
Name
ST_MCS_Pool
Storage Pool
Storage_Pool
Max Capacity
53.26 TiB (Physical) Based on storage pool free unreserved capacity
Advanced Settings
Replication Factor
Reserved Capacity
20
GiB
Advertised Capacity
Total GiB
GiB
Compression
Perform post-process compression of all persistent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

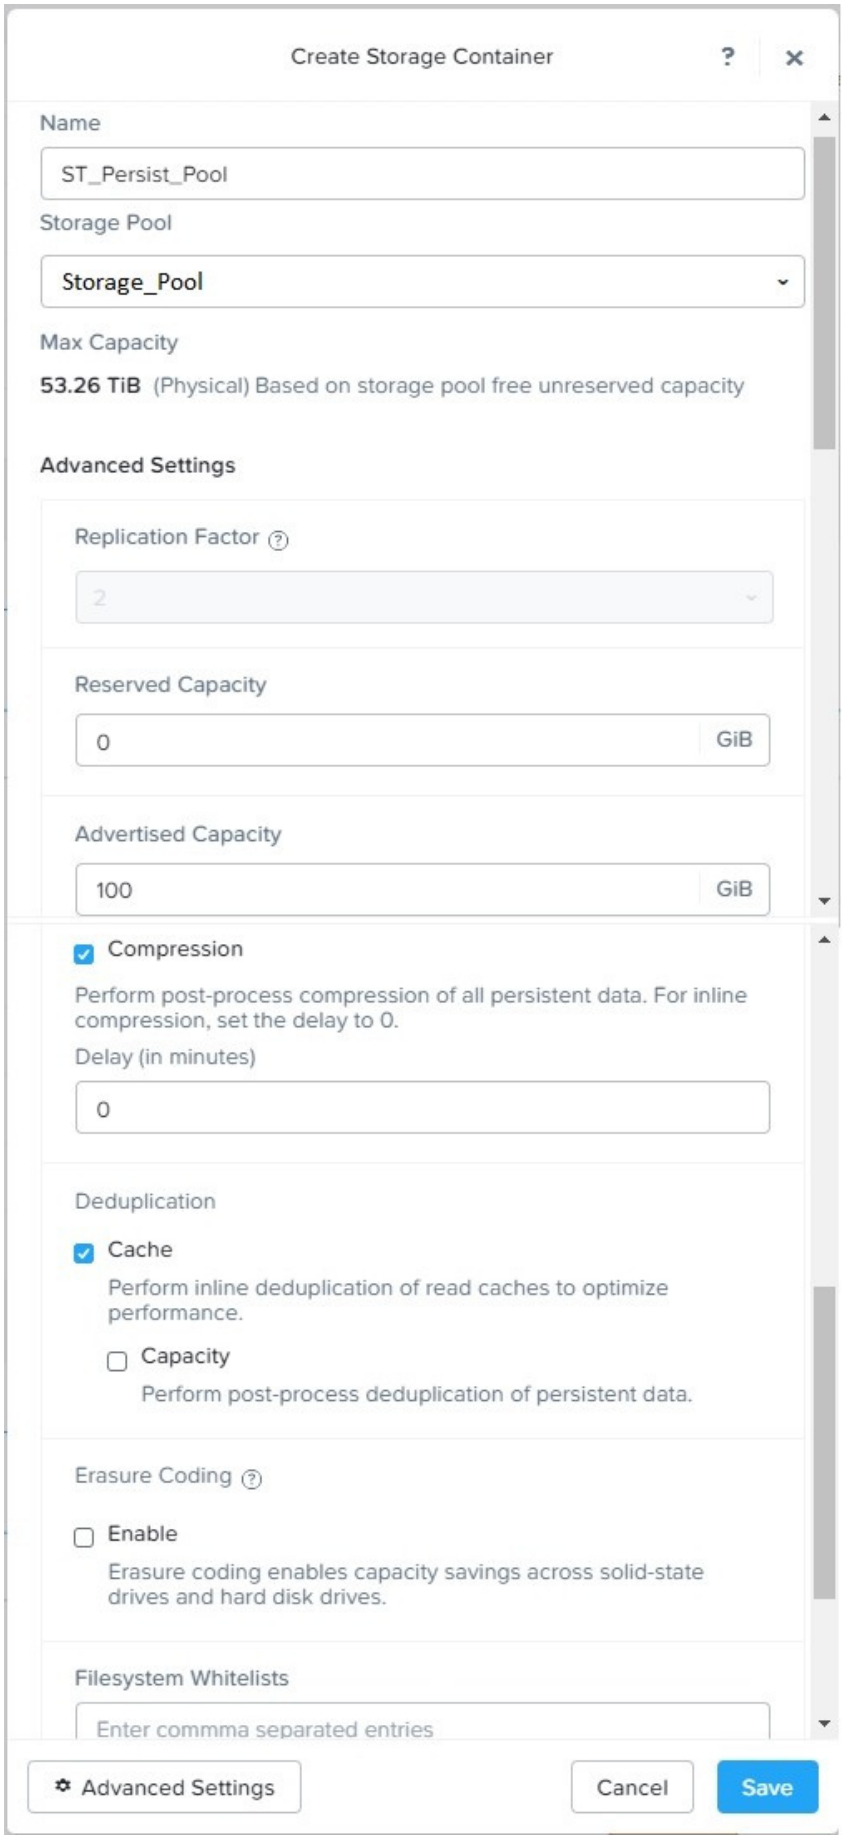
Create Storage Container
Name
ST_Persist_Pool
Storage Pool
Storage_Pool
Max Capacity
53.26 TiB (Physical) Based on storage pool free unreserved capacity
Advanced Settings
Replication Factor
2
Reserved Capacity
0
GiB
Advertised Capacity
100
GiB
Compression
Perform post-process compression of all persistent 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 3

CORRECT TEXT

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

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

Correct Answer: A

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

QUESTION 4

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 or 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 toappdev@cyberdyne.netwhen competed.

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.

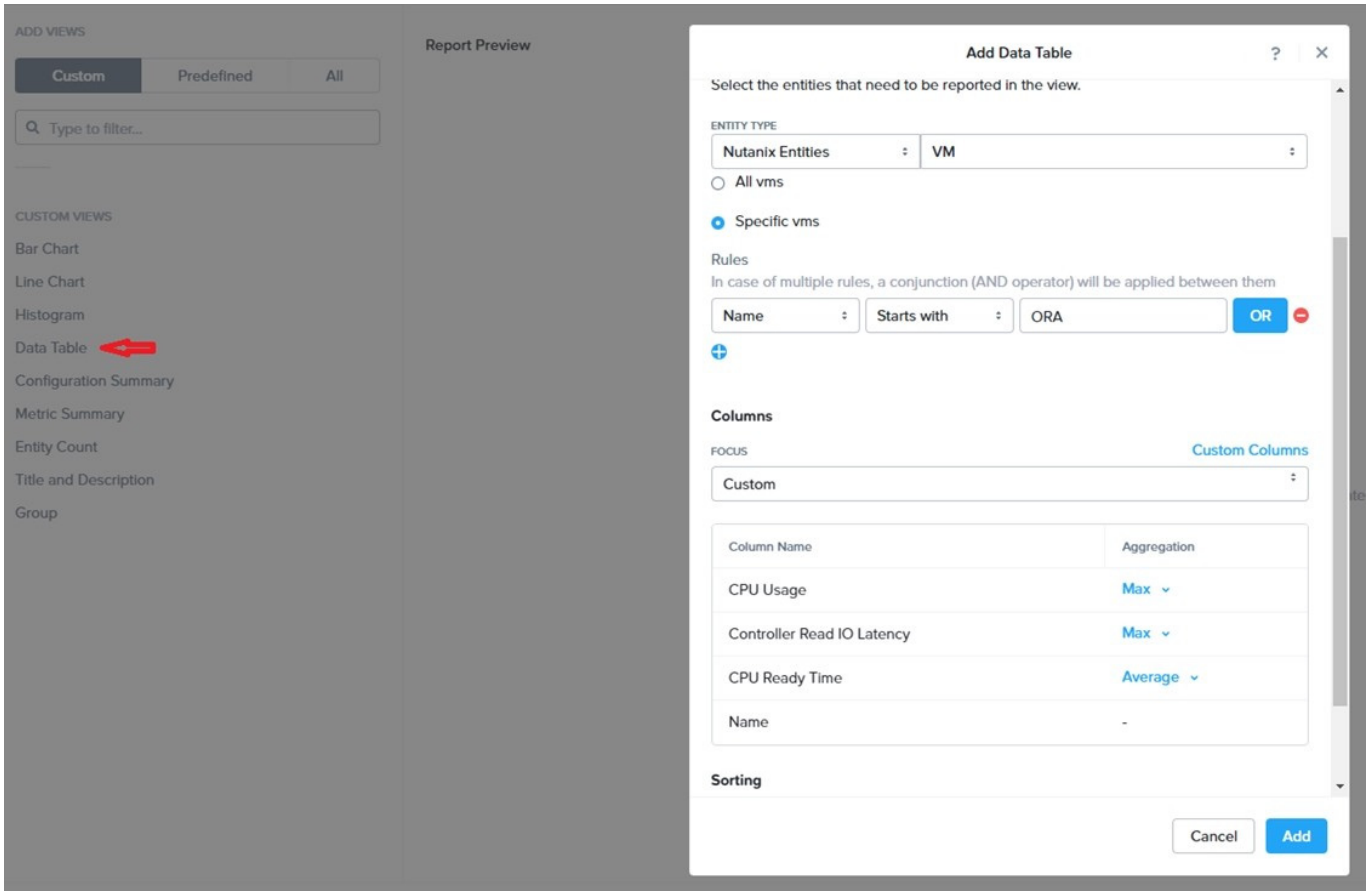

QUESTION 5

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 xxxxxxxxxxx --interfaces ethX,ethX update_uplinks manage_ovs --bridge_name brX-up --interfaces ethX,ethX --bond_name bond1 -- bond_mode xxxxxxxxxxx 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