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.certbus.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 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

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

Correct Answer: A

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.

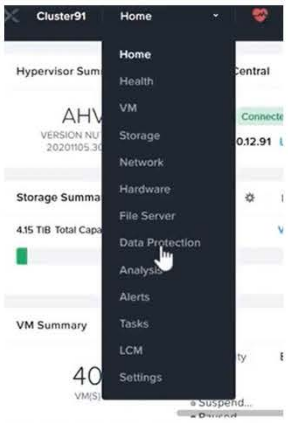
Cluster91
Home
Home
Health
VM
Storage
Network
Hardware
File Server
Data Protection
Alerts
Tasks
LCM
Settings
VM Summary


+ Protection Domain
Async DR

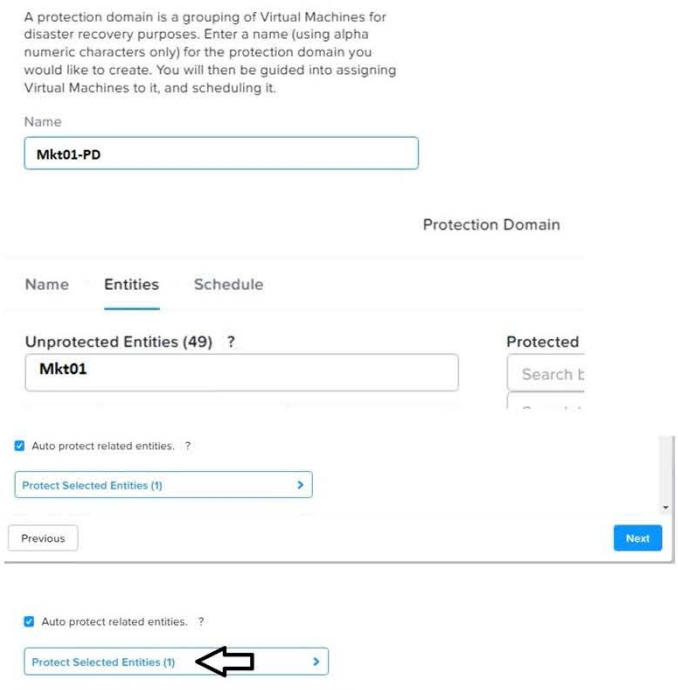
A protection domain is a grouping of Virtual Machines for disaster recovery purposes. Enter a name (using alpha numeric characters only) for the protection domain you would like to create. You will then be guided into assigning Virtual Machines to it, and scheduling it.
Name
Mkt01-PD
Protection Domain
Name
Entities
Schedule
Unprotected Entities (49) ?
Mkt01
Protected
Auto protect related entities. ?
Protect Selected Entities (1)
Previous
Next

Auto protect related entities. ?
Protect Selected Entities (1)

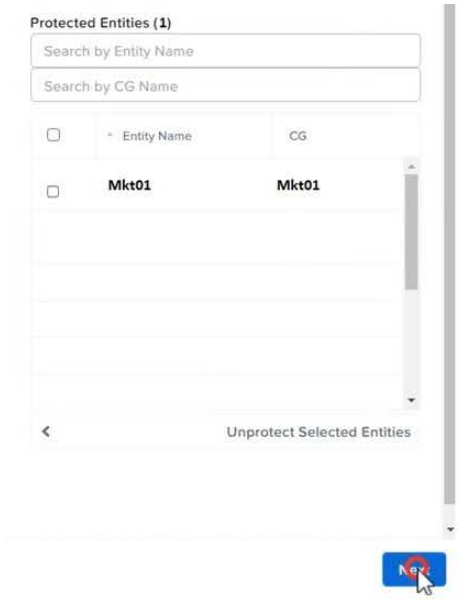
Protected Entities (1)
Search by Entity Name
Search by CG Name
Entity Name
CG
Mkt01
Mkt01
Unprotect Selected Entities
Next


New Schedule

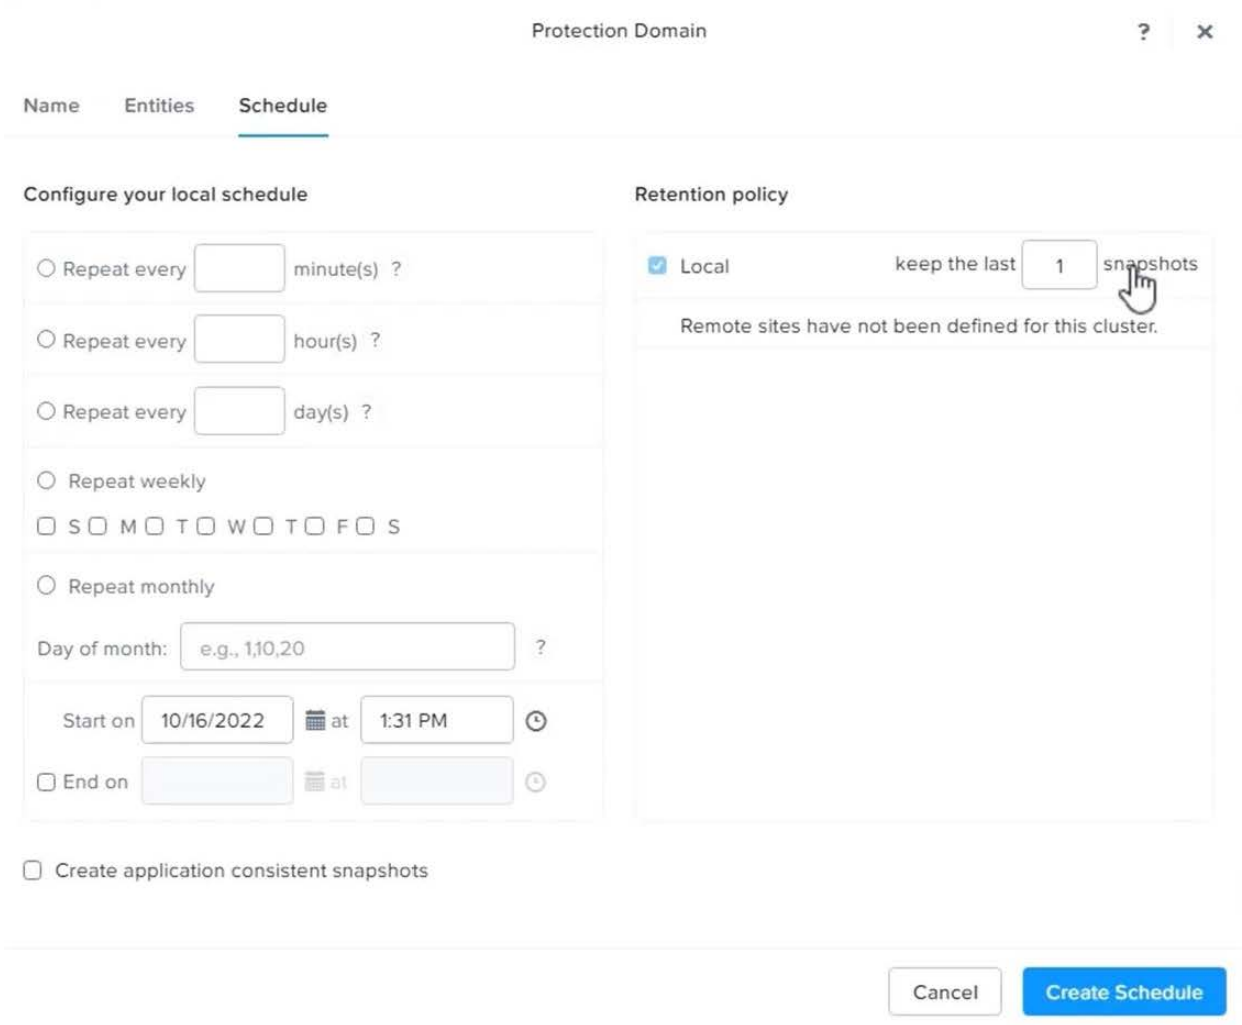
Protection Domain
Name
Entities
Schedule
Configure your local schedule
Repeat every minute(s) ?
Repeat every hour(s) ?
Repeat every day(s) ?
Repeat weekly
S M T W T F S
Repeat monthly
Day of month: e.g., 1,10,20 ?
Start on 10/16/2022 at 1:31 PM
End on
Retention policy
Local keep the last 1 snapshots
Remote sites have not been defined for this cluster.
Create application consistent snapshots
Cancel
Create Schedule

QUESTION 2

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 3

CORRECT TEXT Task 5 An administrator has been informed that a new workload requires a logically segmented network to meet security requirements. Network configuration: VLAN: 667 Network: 192.168.0.0 Subnet Mask: 255.255.255.0 DNS server: 34.82.231.220 Default Gateway: 192.168.0.1 Domain: cyberdyne.net IP Pool: 192.168.9.100-200 DHCP Server IP: 192.168.0.2 Configure the cluster to meet the requirements for the new workload if new objects are required, start the name with 667.

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

Correct Answer: A

To configure the cluster to meet the requirements for the new workload, you need to do the following steps: Create a new VLAN with ID 667 on the cluster. You can do this by logging in to Prism Element and going to Network Configuration > VLANs > Create VLAN. Enter 667 as the VLAN ID and a name for the VLAN, such as 667_VLAN. Create a new network segment with the network details provided. You can do this by logging in to Prism Central and going to Network > Network Segments > Create Network Segment. Enter a name for the network segment, such as 667_Network_Segment, and select 667_VLAN as the VLAN. Enter 192.168.0.0 as the Network Address and 255.255.255.0 as the Subnet Mask. Enter 192.168.0.1 as the Default Gateway and 34.82.231.220 as the DNS Server. Enter cyberdyne.net as the Domain Name. Create a new IP pool with the IP range provided. You can do this by logging in to Prism Central and going to Network > IP Pools > Create IP Pool. Enter a name for the IP pool, such as 667_IP_Pool, and select 667_Network_Segment as the Network Segment. Enter 192.168.9.100 as the Starting IP Address and 192.168.9.200 as the Ending IP Address. Configure the DHCP server with the IP address provided. You can do this by logging in to Prism Central and going to Network > DHCP Servers > Create DHCP Server. Enter a name for the DHCP server, such as 667_DHCP_Server, and select 667_Network_Segment as the Network Segment. Enter 192.168.0.2 as the IP Address and select 667_IP_Pool as the IP Pool.

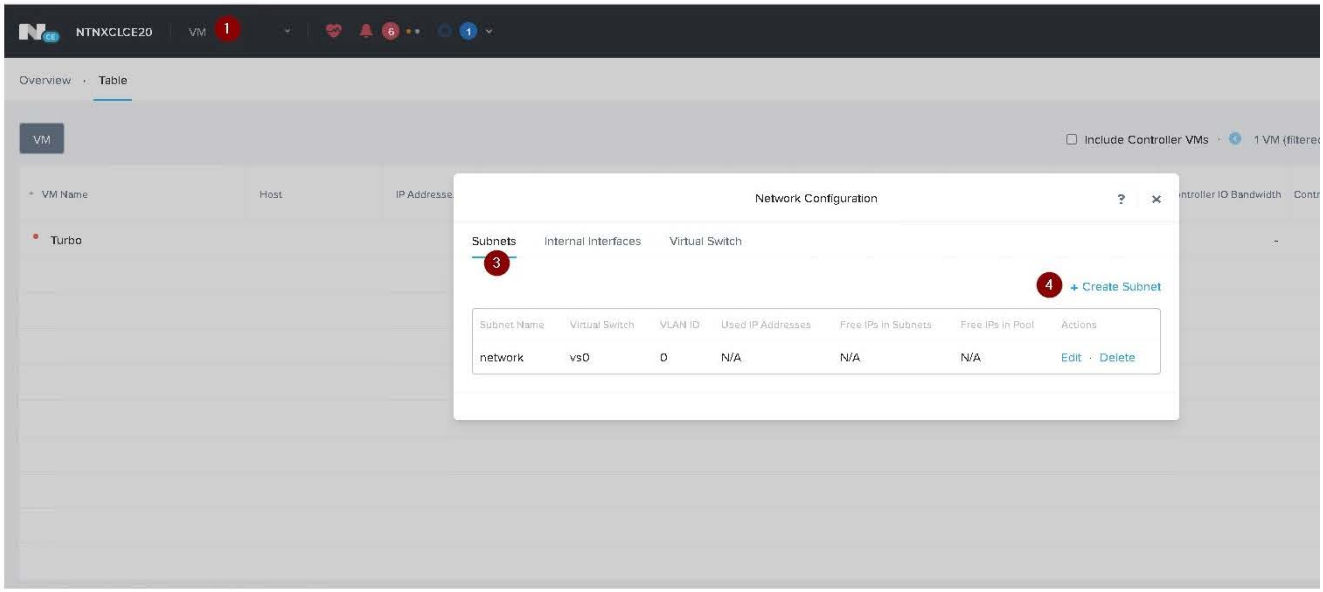
NTNXCLCE20
VM
1
Overview
Table
VM
Include Controller VMs
VM Name
Host
IP Address
Turbo
Network Configuration
Subnets
Internal Interfaces
Virtual Switch
3
4
+ Create Subnet
Subnet Name
Virtual Switch
VLAN ID
Used IP Addresses
Free IPs in Subnets
Free IPs in Pool
Actions
network
vs0
0
N/A
N/A
N/A
Edit
Delete

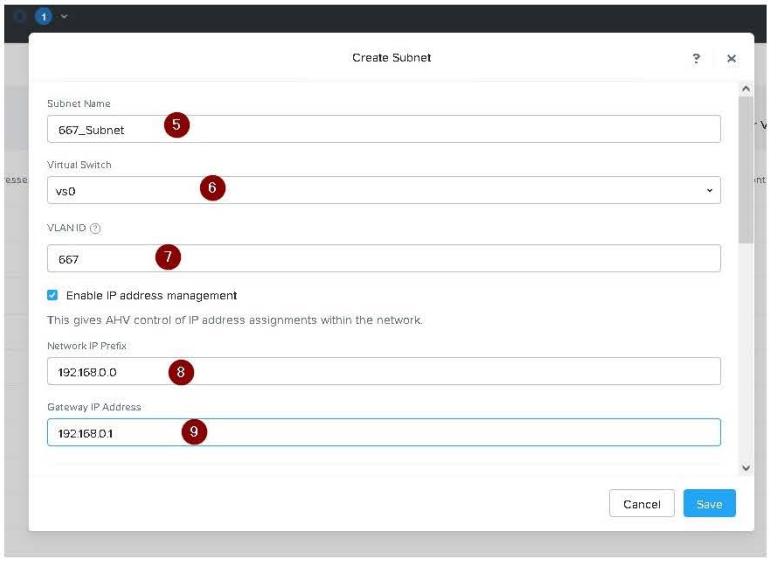
Create Subnet
Subnet Name
667_Subnet
5
Virtual Switch
vs0
6
VLAN ID
667
7
Enable IP address management
This gives AHV control of IP address assignments within the network.
Network IP Prefix
192.168.0.0
8
Gateway IP Address
192.168.0.1
9
Cancel
Save


Create Subnet
DHCP Settings
Domain Name Servers (Comma Separated)
34.82.231.220
10
Domain Search (Comma Separated)
cyberdyne.net
11
Domain Name
cyberdyne
12
TFTP Server Name
Boot File Name
Cancel
Save

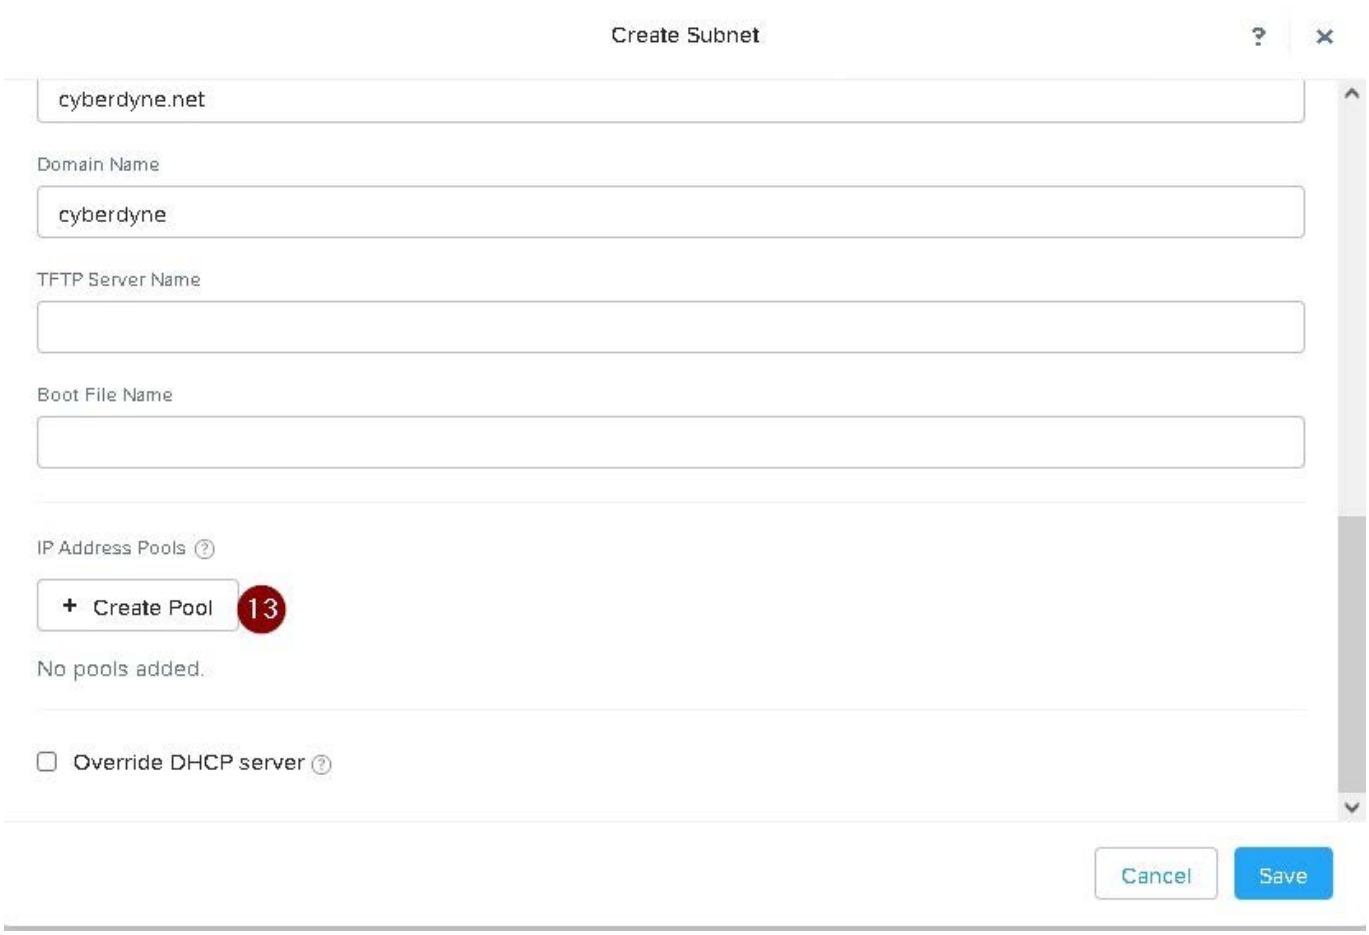
Create Subnet
cyberdyne.net
Domain Name
cyberdyne
TFTP Server Name
Boot File Name
IP Address Pools
+ Create Pool
13
No pools added.
Override DHCP server
Cancel
Save

QUESTION 4

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 5

CORRECT TEXT

Task 16

Running NCC on a cluster prior to an upgrade results in the following output

FAIL: CVM System Partition /home usage at 93% (greater than threshold, 90%) Identify the CVM with the issue, remove the fil causing the storage bloat, and check the health again by running the individual disk usage health check only on the problematic CVM do not run NCC health check

Note: Make sure only the individual health check is executed from the affected node

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

Correct Answer: A

To identify the CVM with the issue, remove the file causing the storage bloat, and check the health again, 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 find the NCC health check output file from the list. You can use the date and time information to locate the file. The file name

should be something like ncc-output-YYYY-MM-DD-HH-MM-SS.log. Open the file and look for the line that says FAIL: CVM System Partition /home usage at 93% (greater than threshold, 90%). Note down the IP address of the CVM that has

this issue. It should be something like X.X.X.X.

Log in to the CVM using SSH or console with the username and password provided. Run the command du -sh /home/* to see the disk usage of each file and directory under /home. Identify the file that is taking up most of the space. It could

be a log file, a backup file, or a temporary file. Make sure it is not a system file or a configuration file that is needed by the CVM.

Run the command rm -f /home/ to remove the file causing the storage bloat. Replace with the actual name of the file. Run the command ncc health_checks hardware_checks disk_checks disk_usage_check -cvm_list=X.X.X.X to check the health again by running the individual disk usage health check only on the problematic CVM. Replace X.X.X.X with the IP address of the CVM that you noted down earlier.

Verify that the output shows PASS: CVM System Partition /home usage at XX% (less than threshold, 90%). This means that the issue has been resolved.

#access to CVM IP by Putty

allssh df -h #look for the path /dev/sdb3 and select the IP of the CVM ssh CVM_IP

ls

cd software_downloads

ls

cd nos

ls -l -h

rm files_name

df -h

ncc health_checks hardware_checks disk_checks disk_usage_check