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 11

An administrator has noticed that after a host failure, the SQL03 VM was not powered back on from another host within the cluster. The Other SQL VMs (SQL01, SQL02) have recovered properly in the past.

Resolve the issue and configure the environment to ensure any single host failure affects a minimal number os SQL VMs.

Note: Do not power on any VMs

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

Correct Answer: A

One possible reason why the SQL03 VM was not powered back on after a host failure is that the cluster was configured with the default (best effort) VM high availability mode, which does not guarantee the availability of VMs in case of

insufficient resources on the remaining hosts. To resolve this issue, I suggest changing the VM high availability mode to guarantee (reserved segments), which reserves some memory on each host for failover of VMs from a failed host. This

way, the SQL03 VM will have a higher chance of being restarted on another host in case of a host failure. To change the VM high availability mode to guarantee (reserved segments), you can follow these steps:

Log in to Prism Central and select the cluster where the SQL VMs are running. Click on the gear icon on the top right corner and select Cluster Settings. Under Cluster Services, click on Virtual Machine High Availability. Select Guarantee

(Reserved Segments) from the drop-down menu and click Save. To configure the environment to ensure any single host failure affects a minimal number of SQL VMs, I suggest using anti-affinity rules, which prevent VMs that belong to the

same group from running on the same host. This way, if one host fails, only one SQL VM will be affected and the other SQL VMs will continue running on different hosts. To create an anti-affinity rule for the SQL VMs, 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 Group. Enter a name for the group, such as SQL Group, and click Next. Select the SQL VMs (SQL01,

SQL02, SQL03) from the list and click Next. Select Anti-Affinity from the drop-down menu and click Next.

Review the group details and click Finish.

I hope this helps. How else can I help?

https://portal.nutanix.com/page/documents/details?targetId=AHV-Admin-Guide-v6_5:ahv- affinity-policies-c.html

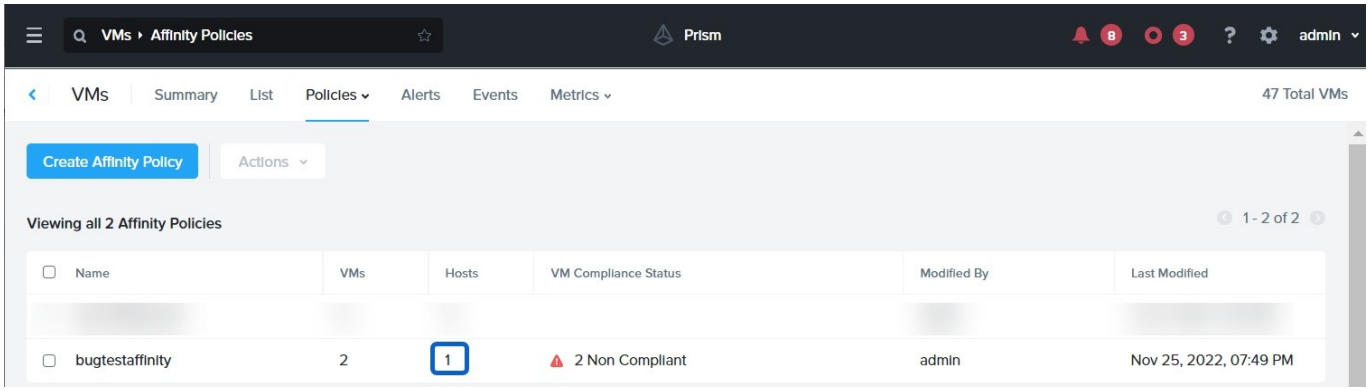

A screenshot of a computer

Description automatically generated with medium confidence

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.

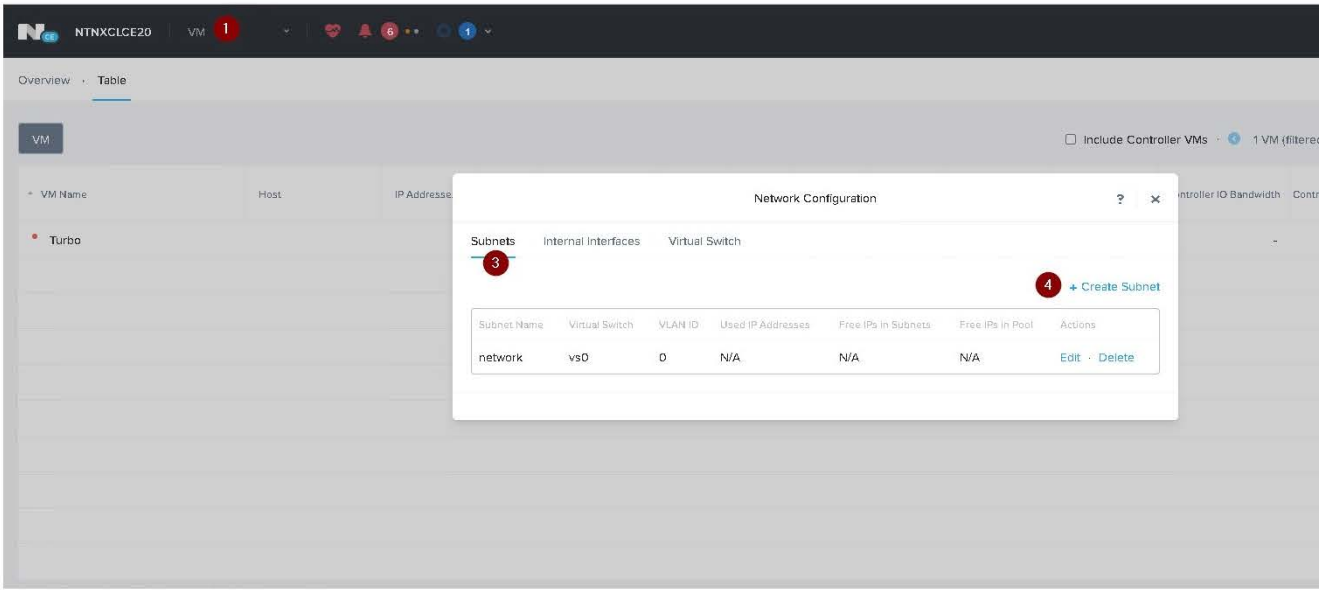
NTNXCLCE20
VM
1
Overview
Table
VM
Include Controller VMs
VM Name
Host
IP Address
Turbo
Network Configuration
Subnets
Internal Interfaces
Virtual Switch
3
4
+ Create Subnet
Subnet Name
Virtual Switch
VLAN ID
Used IP Addresses
Free IPs in Subnets
Free IPs in Pool
Actions
network
vs0
0
N/A
N/A
N/A
Edit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

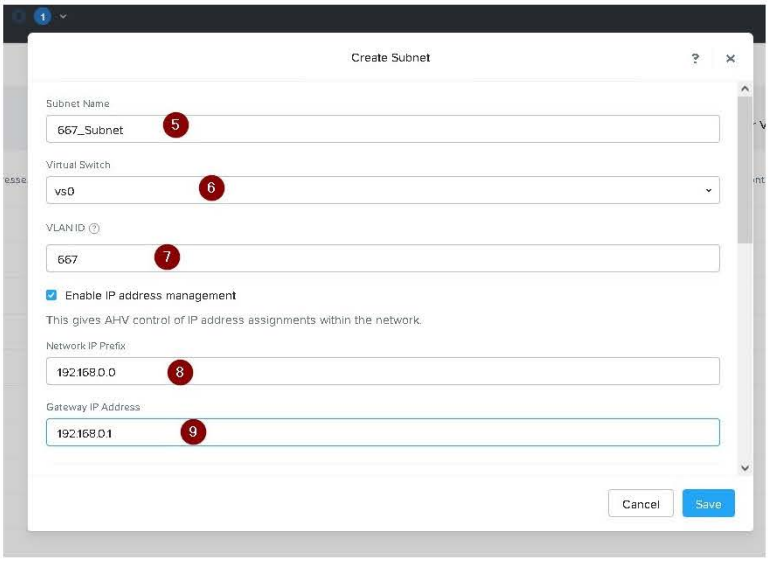
Create Subnet
Subnet Name
667_Subnet
5
Virtual Switch
vs0
6
VLAN ID
667
7
Enable IP address management
This gives AHV control of IP address assignments within the network.
Network IP Prefix
192.168.0.0
8
Gateway IP Address
192.168.0.1
9
Cancel
Save


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

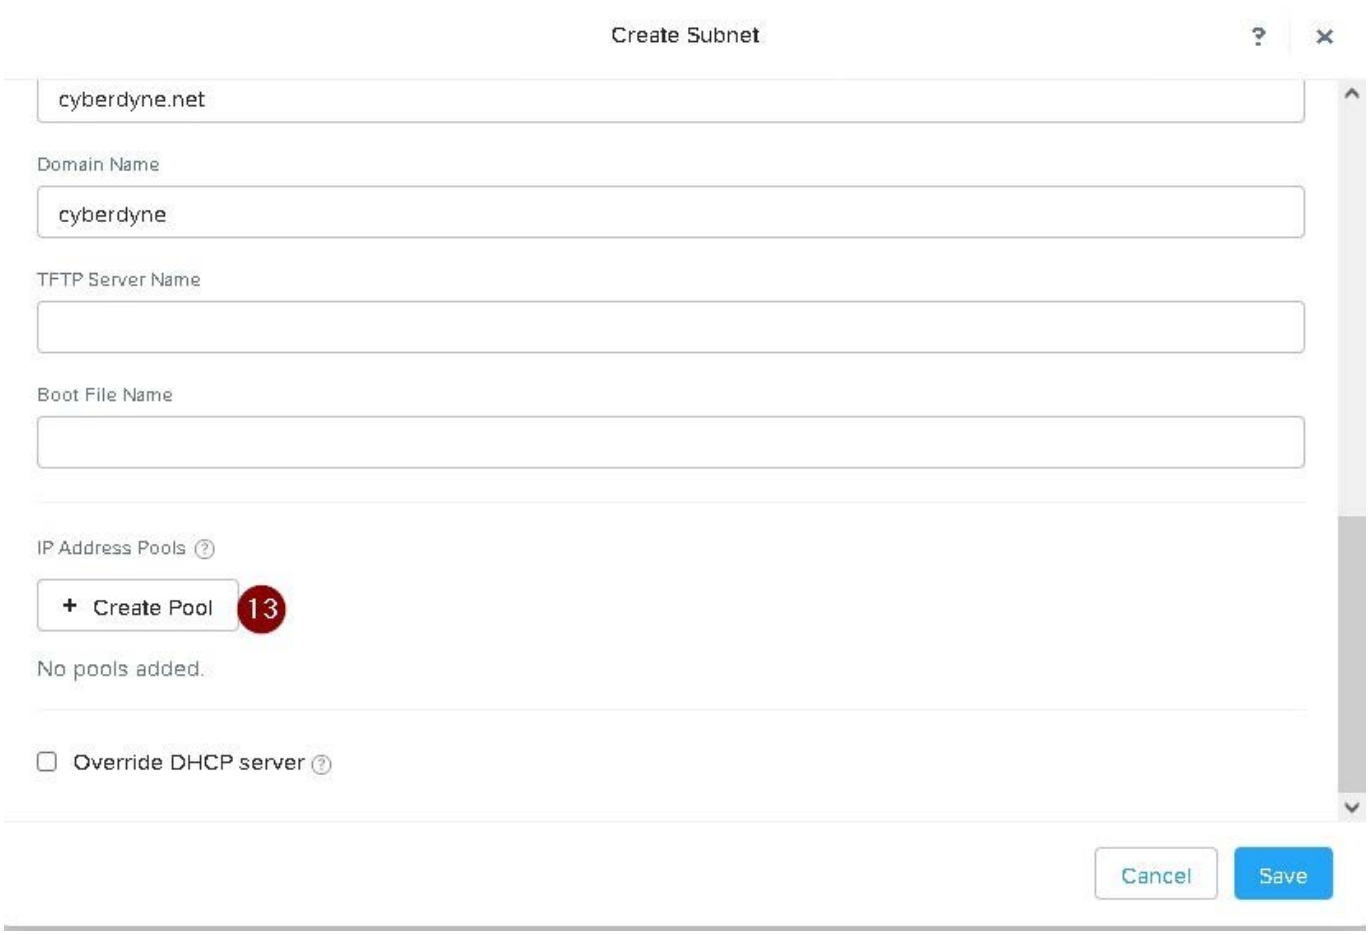
Create Subnet
cyberdyne.net
Domain Name
cyberdyne
TFTP Server Name
Boot File Name
IP Address Pools
+ Create Pool
13
No pools added.
Override DHCP server
Cancel
Save

QUESTION 4

CORRECT TEXT

Task 7

An administrator has environment that will soon be upgraded to 6.5. In the meantime, they need to implement log and apply a security policy named Staging_Production, such that not VM in the Staging Environment can communicate with any

VM in the production Environment,

Configure the environment to satisfy this requirement.

Note: All other configurations not indicated must be left at their default values.

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

Correct Answer: A

To configure the environment to satisfy the requirement of implementing a security policy named Staging_Production, such that no VM in the Staging Environment can communicate with any VM in the production Environment, you need to do the following steps: Log in to Prism Central and go to Network > Security Policies > Create Security Policy. Enter Staging_Production as the name of the security policy and select Cluster A as the cluster. In the Scope section, select

VMs as the entity type and add the VMs that belong to the Staging Environment and the Production Environment as the entities. You can use tags or categories to filter the VMs based on their environment. In the Rules section, create a new rule with the following settings: Direction: Bidirectional Protocol: Any Source: Staging Environment Destination: Production Environment Action: Deny Save the security policy and apply it to the cluster. This will create a security policy that will block any traffic between the VMs in the Staging Environment and the VMs in the Production Environment. You can verify that the security policy is working by trying to ping or access any VM in the Production Environment from any VM in the Staging Environment, or vice versa. You should not be able to do so.

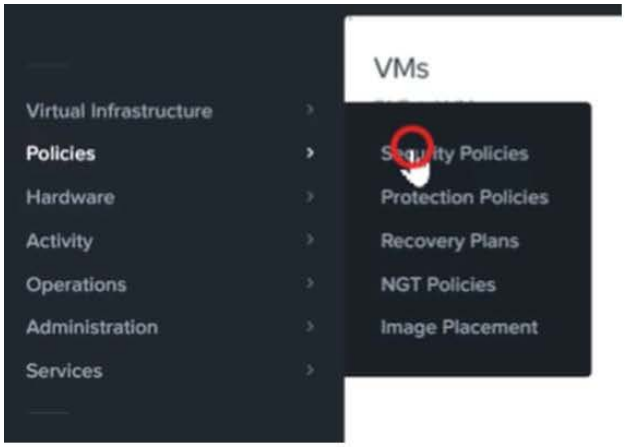
VMs
Virtual Infrastructure
Policies
Hardware
Activity
Operations
Administration
Services
Protection Policies
Recovery Plans
NGT Policies
Image Placement

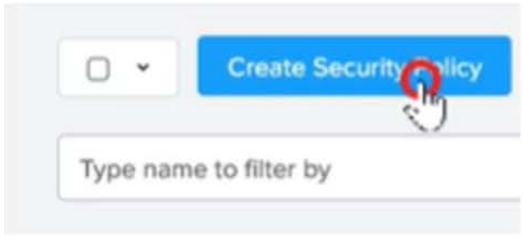
Create Security Policy
Type name to filter by

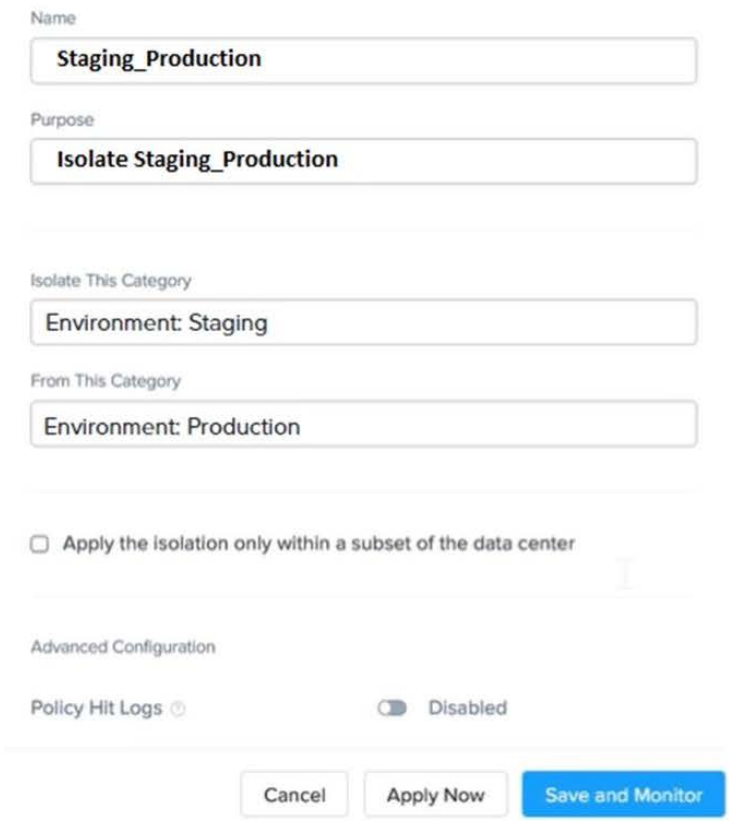
Name
Staging_Production
Purpose
Isolate Staging_Production
Isolate This Category
Environment: Staging
From This Category
Environment: Production
Apply the isolation only within a subset of the data center
Advanced Configuration
Policy Hit Logs
Disabled
Cancel
Apply Now
Save and Monitor


Actions
Create Security Policy
Export & Import
Filters
Update
Delete
To enforce the policy, check the box next to the policy, choose Actions, then Apply.
Staging_Production
Isolate HR from IT
Environment: Staging
Environment: Production
Monitoring
few seconds ago

QUESTION 5

CORRECT TEXT

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.

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

Correct Answer: A

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.