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 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.

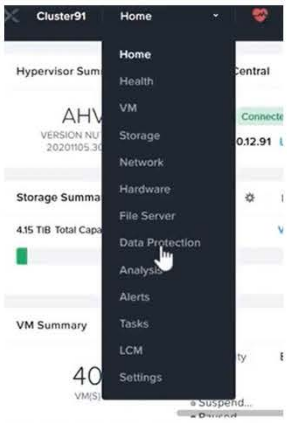
Cluster91
Home
Home
Health
VM
Storage
Network
Hardware
File Server
Data Protection
Alerts
Tasks
LCM
Settings
VM Summary


+ Protection Domain
Async DR

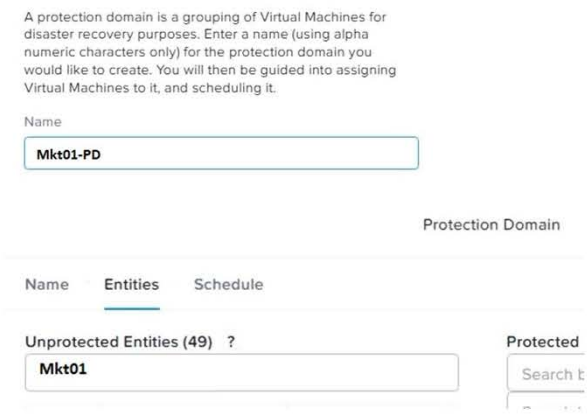
A protection domain is a grouping of Virtual Machines for disaster recovery purposes. Enter a name (using alpha numeric characters only) for the protection domain you would like to create. You will then be guided into assigning Virtual Machines to it, and scheduling it.
Name
Mkt01-PD

Protection Domain
Name
Entities
Schedule
Unprotected Entities (49) ?
Mkt01
Protected
Search t
Auto protect related entities. ?
Protect Selected Entities (1)
Previous
Next


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

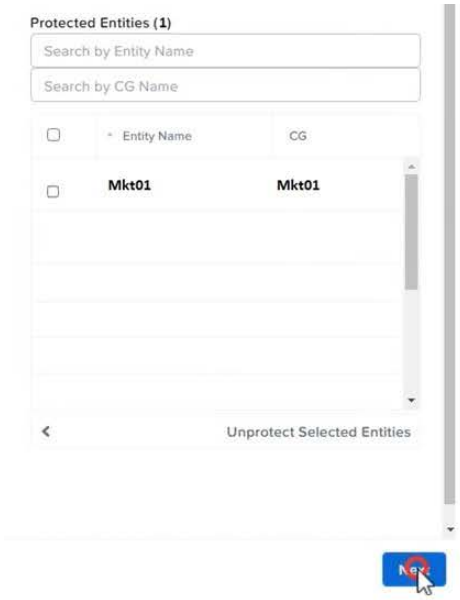
Protected Entities (1)
Search by Entity Name
Search by CG Name
Entity Name
CG
Mkt01
Mkt01
Unprotect Selected Entities
Next


New Schedule

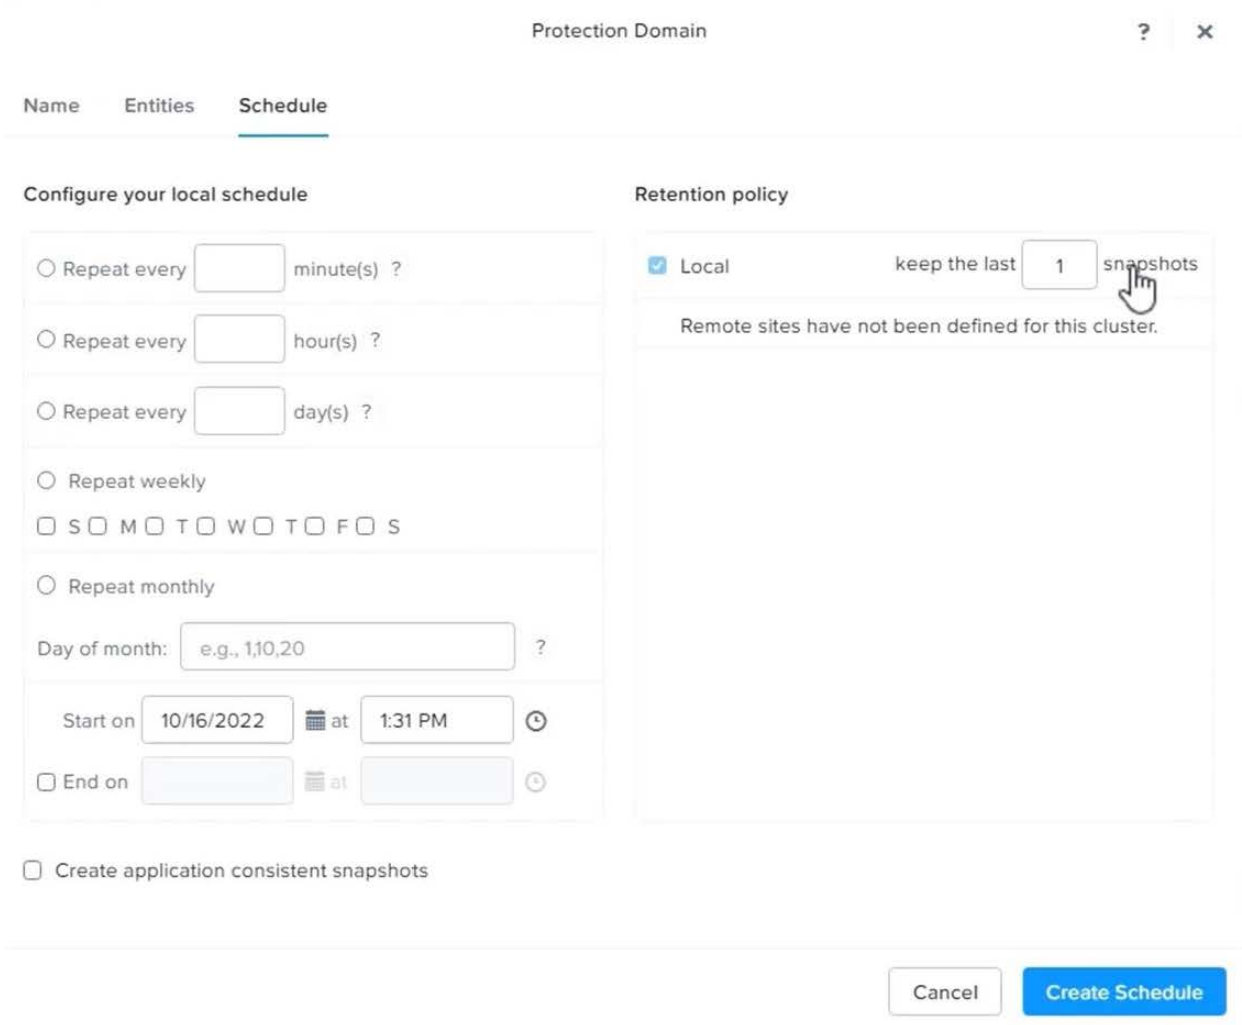
Protection Domain
Name
Entities
Schedule
Configure your local schedule
Repeat every minute(s) ?
Repeat every hour(s) ?
Repeat every day(s) ?
Repeat weekly
S M T W T F S
Repeat monthly
Day of month: e.g., 1,10,20 ?
Start on 10/16/2022 at 1:31 PM
End on
Retention policy
Local keep the last 1 snapshots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

QUESTION 2

CORRECT TEXT

Task 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.

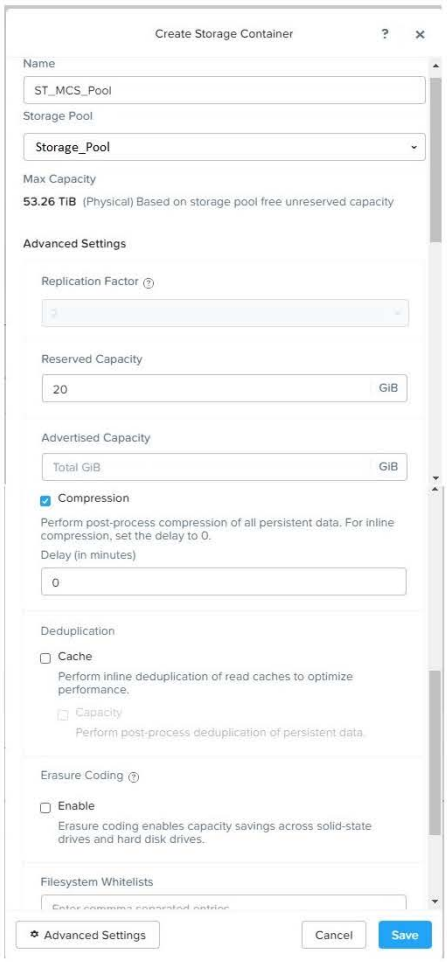
Create Storage Container
Name
ST_MCS_Pool
Storage Pool
Storage_Pool
Max Capacity
53.26 TiB (Physical) Based on storage pool free unreserved capacity
Advanced Settings
Replication Factor
Reserved Capacity
20
GiB
Advertised Capacity
Total GiB
GiB
Compression
Perform post-process compression of all persistent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

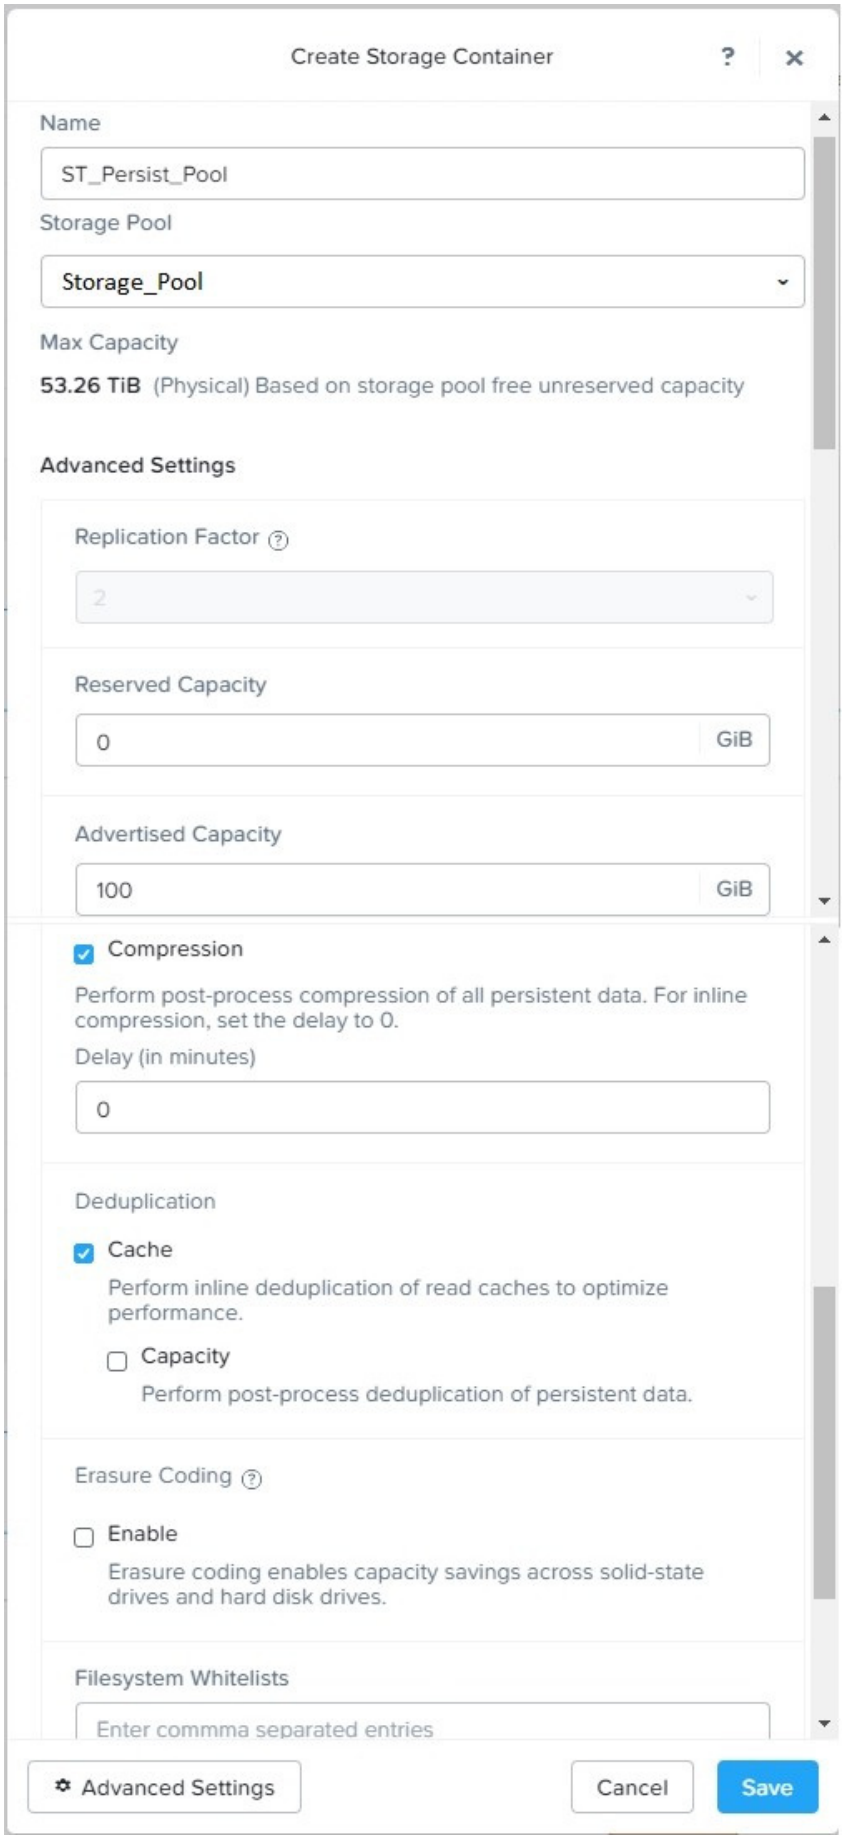
Create Storage Container
Name
ST_Persist_Pool
Storage Pool
Storage_Pool
Max Capacity
53.26 TiB (Physical) Based on storage pool free unreserved capacity
Advanced Settings
Replication Factor
2
Reserved Capacity
0
GiB
Advertised Capacity
100
GiB
Compression
Perform post-process compression of all persistent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

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

QUESTION 4

CORRECT TEXT

Task 15

An administrator found a CentOS VM, Cent_Down, on the cluster with a corrupted network stack. To correct the issue, the VM will need to be restored from a previous snapshot to become reachable on the network again.

VM credentials:

Username: root

Password: nutanix/4u

Restore the VM and ensure it is reachable on the network by pinging 172.31.0.1 from the VM.

Power off the VM before proceeding.

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

Correct Answer: A

To restore the VM and ensure it is reachable on the network, you can follow these steps:

Log in to the Web Console of the cluster where the VM is running. Click on Virtual Machines on the left menu and find Cent_Down from the list. Click on the power icon to power off the VM.

Click on the snapshot icon next to the power icon to open the Snapshot Management window.

Select a snapshot from the list that was taken before the network stack was corrupted. You can use the date and time information to choose a suitable snapshot. Click on Restore VM and confirm the action in the dialog box. Wait for the

restore process to complete.

Click on the power icon again to power on the VM. Log in to the VM using SSH or console with the username and password provided. Run the command ping 172.31.0.1 to verify that the VM is reachable on the network. You should see a

reply from the destination IP address.

Go to VMS from the prism central gui

Select the VMand go to More -> Guest Shutdown

Go to Snapshots tab and revert to latest snapshot available power on vm and verify if ping is working

QUESTION 5

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

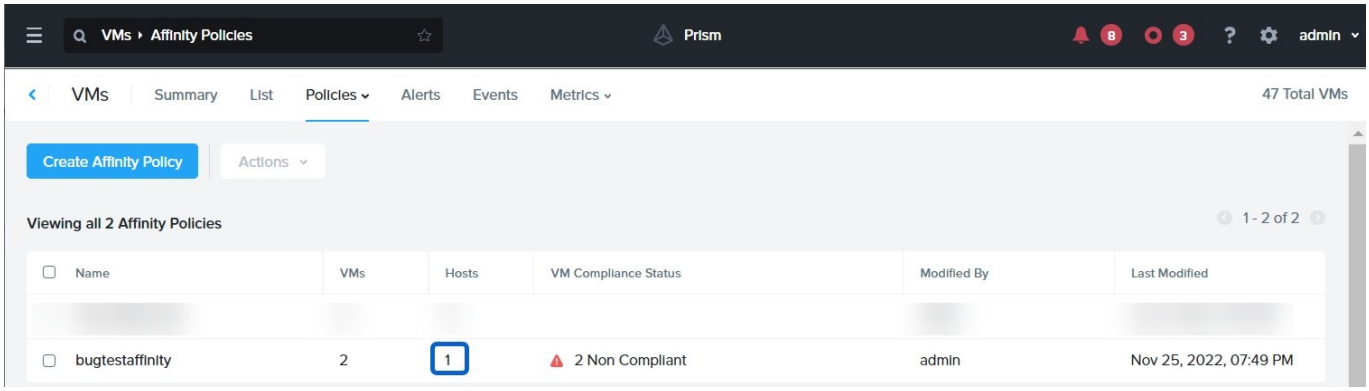

A screenshot of a computer

Description automatically generated with medium confidence