

주제 5

- Advanced Configuration and Troubleshooting: This topic covers sub-topics of executing API calls, configuring third-party integrations, analyzing AOS security posture, and translating business needs into technical solutions. Lastly, it discusses troubleshooting Nutanix services as well.

>> NCM-MCI덤프문제 <<

인기자격증 NCM-MCI덤프문제 시험덤프공부

어떻게 Nutanix인증 NCM-MCI시험을 패스하느냐에는 여러 가지 방법이 있습니다. 하지만 여러분의 선택에 따라 보장도 또한 틀립니다. 우리 Itcertkr에서는 아주 완벽한 학습가이드를 제공하며, Nutanix인증 NCM-MCI시험은 아주 간편하게 패스하실 수 있습니다. Itcertkr에서 제공되는 문제와 답은 모두 실제 Nutanix인증 NCM-MCI시험에서나 오는 문제들입니다. 일종의 기출문제입니다. 때문에 우리 Itcertkr덤프의 보장 도와 정확도는 안심하셔도 좋습니다. 무조건 Nutanix인증 NCM-MCI시험을 통과하게 만듭니다. 우리 Itcertkr 또한 끈임 없는 덤프갱신으로 페펙트한 Nutanix인증 NCM-MCI시험자료를 여러분들한테 선사하겠습니다.

최신 Master Level NCM-MCI 무료샘플문제 (Q15-Q20):

질문 # 15

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 of SQL VMs.

Note: Do not power on any VMs

정답:

설명:

See the Explanation for step by step solution

Explanation:

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

Name	VMs	Hosts	VM Compliance Status	Modified By	Last Modified
bugtestaffinity	1	2	2 Non-Compliant	admin	Nov 25, 2022, 07:49 PM
2 Non-Compliant	2	2	2 Non-Compliant		

질문 # 16

Topic 1, Performance Based Questions

Environment

You have been provisioned a dedicated environment for your assessment which includes the following:

Workstation

- * windows Server 2019

- * All software/tools/etc to perform the required tasks

- * Nutanix Documentation and whitepapers can be found in desktop\files\Documentation

- * Note that the workstation is the system you are currently logged into Nutanix Cluster

- * There are three clusters provided. The connection information for the relevant cluster will be displayed to the right of the question

Please make sure you are working on the correct cluster for each item. Please ignore any licensing violations

- * Cluster A is a 3-node cluster with Prism Central 2022.6 where most questions will be performed

- * Cluster B is a one-node cluster and has one syslog item and one security item to perform

- * Cluster D is a one-node cluster with Prism Central 5.17 and has a security policy item to perform Important Notes

- * If the text is too small and hard to read, or you cannot see all of the GUI, you can increase/decrease the zoom of the browser with CTRL + and CTRL - (the plus and minus keys). You will be given 3 hours to complete the scenarios for Nutanix NCMMC. Once you click the start button below, you will be provided with:

- A Windows desktop A browser page with the scenarios and credentials (Desktop\instructions). Notes for this exam delivery:

The browser can be scaled to improve visibility and fit all the content on the screen.

- Copy and paste hot-keys will not work. Use your mouse for copy and paste.

- The Notes and Feedback tabs for each scenario are to leave notes for yourself or feedback for

- Make sure you are performing tasks on the correct components.

- Changing security or network settings on the wrong component may result in a failing grade.

- Do not change credentials on any component unless you are instructed to.

- All necessary documentation is contained in the Desktop\Files\Documentation directory. 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.

정답:

설명:

See the Explanation for step by step solution

Explanation:

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.

질문 #17

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 file 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

정답:

설명:

See the Explanation for step by step solution

Explanation:

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/<filename> to remove the file causing the storage bloat. Replace <filename> 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

질문 #18

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 to appdev@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.

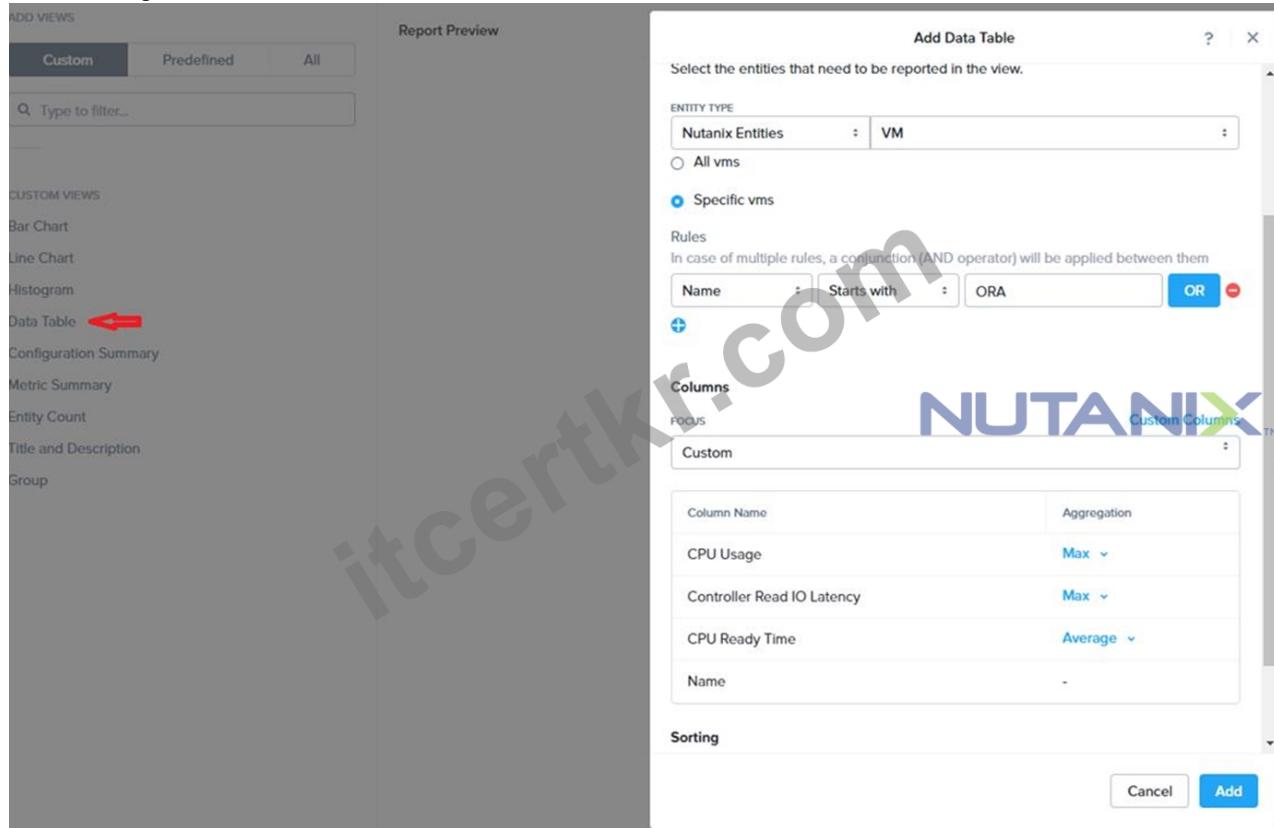
A: Click Next.

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.



정답:

설명:

See the Explanation for step by step solution

Explanation:

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 OR

질문 # 19

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.

정답:

설명:

See the Explanation for step by step solution

Explanation:

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 VM and go to More -> Guest Shutdown

Go to Snapshots tab and revert to latest snapshot available

power on vm and verify if ping is working

질문 # 20

.....

Nutanix NCM-MCI 덤프의 PDF 버전과 Software 버전의 내용은 동일합니다. PDF버전은 프린트 가능한 버전으로서 단독구매하셔도 됩니다. Software 버전은 테스트용으로 PDF 버전 공부를 마친후 시험전에 실력테스트 가능합니다. Software 버전은 PDF버전의 보조용이기에 단독 판매하지 않습니다. 소프트웨어버전까지 필요하신 분은 PDF버전을 구입하실때 공동구매하셔야 합니다.

NCM-MCI 높은 통과율 시험자료 : https://www.itcertkr.com/NCM-MCI_exam.html

- 최신 업데이트버전 NCM-MCI덤프문제 덤프문제공부 □ 무료 다운로드를 위해▶ NCM-MCI ◀를 검색하려면
▶ www.koreadumps.com◀(를) 입력하십시오NCM-MCI유효한 인증시험덤프
- NCM-MCI시험패스 가능 덤프공부 □ NCM-MCI합격보장 가능 덤프 □ NCM-MCI합격보장 가능 덤프 □
□ www.itdumpskr.com□은▶ NCM-MCI □무료 다운로드를 받을 수 있는 최고의 사이트입니다NCM-MCI최신
인증시험 공부자료
- 높은 통과율 NCM-MCI덤프문제 시험덤프로 시험패스가능 □ 시험 자료를 무료로 다운로드하려면[
kr.fast2test.com]을 통해[NCM-MCI]를 검색하십시오NCM-MCI시험패스 덤프공부자료
- NCM-MCI덤프문제 덤프 최신버전 □ ★ www.itdumpskr.com ★□을 통해 쉽게▶ NCM-MCI □무료 다운로드 받기NCM-MCI최고합격덤프
- NCM-MCI덤프문제 최신 시험 최신 덤프자료 □ 오픈 웹 사이트□ www.koreadumps.com□검색 「 NCM-MCI 」 무료 다운로드NCM-MCI퍼펙트 최신 덤프자료
- NCM-MCI시험패스 인증공부자료 □ NCM-MCI시험대비 공부문제 □ NCM-MCI테스트자료 □ ⇒
www.itdumpskr.com⇐에서▶ NCM-MCI □□□를 검색하고 무료로 다운로드하세요NCM-MCI시험패스 덤프공부자료

