



Free Questions for NCP-MCI by actualtestdumps

Shared by Rivas on 03-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Refer to Exhibit:

I/O working set size | 3 Anomalies



Behavioral Anomaly ⓘ

Every time there is an anomaly, alert

Critical

Ignore all anomalies between

0 MB and 4000 MB

Static Threshold ⓘ

Alert Critical if

<= MB or >= MB

Alert Warning if

< MB or > MB

Trigger alert if conditions persist for

10 Minutes

Which statement is true?

Options:

- A- A critical alert will be triggered if I/O working set size goes over 6000 MB.
- B- A critical alert will be triggered when there is an anomaly above 4000 MB.
- C- A warning alert will be triggered after 3 anomalies have been catch.
- D- A warning alert will be triggered if I/O working set size goes over the blue band.

Answer:

B

Explanation:

A) This statement is incorrect because there is no static threshold set to trigger a critical alert at 6000 MB. The graph shows a peak that goes above 6000 MB, but the alert configuration below does not specify a static threshold at this value.

B) This is the correct statement. The configuration under 'Behavioral Anomaly' is set to alert every time there is an anomaly, with a critical level alert set to trigger when the I/O working set size is between 0 MB and 4000 MB. The graph illustrates that the anomalies (highlighted in pink) occur when the working set size exceeds the normal range (blue band). Therefore, any anomaly detected above 4000 MB would trigger a critical alert.

C) This statement is incorrect because there is no indication that a warning alert is configured to trigger after 3 anomalies. The exhibit does not show any configuration that specifies an alert based on the number of anomalies.

D) This statement is incorrect as there's no indication that a warning alert will be triggered based on the I/O working set size exceeding the blue band. The alert settings are configured to ignore anomalies below 4000 MB and to trigger a critical alert for anomalies above this threshold.

The settings displayed in the exhibit are typically part of Nutanix's Prism infrastructure management platform, which can set various thresholds for performance metrics and trigger alerts based on those thresholds. The behavior is defined in the Prism documentation where the alert configuration is outlined.

Question 2

Question Type: MultipleChoice

Refer to the Exhibit:

```
admin@NTNX:~$ manage_ovs show_uplinks
Bridge: br0
Bond: br0-up
    bond_mode: balance-tcp
    interfaces: eth3 eth2 eth1 eth0
    lacp: active
    lacp-fallback: false
    lacp_speed: fast
admin@NTNX:~$
```

An administrator is adding a new node to a cluster. The node has been imaged to the same versions of AHV and AOS that the cluster is running, configured with appropriate IP addresses, and br0-up has been configured in the same manner as the existing uplink bonds. When attempting to add the node to the cluster with the Expand Cluster function in Prism, the cluster is unable to find the new node. Based on the above output from the new node, what is most likely the cause of this issue?

Options:

- A- There is a firewall blocking the discovery traffic from the cluster.
- B- The ports on the upstream switch are not configured for LACP.

- C- The existing cluster and the expansion node are on different VLANs.
- D- LACP configuration must be completed after cluster expansion.

Answer:

B

Explanation:

The output in the exhibit indicates that the node's network interfaces (eth0-eth3) are bonded together using LACP (Link Aggregation Control Protocol) with 'balance-tcp' as the bonding mode and LACP speed set to 'fast'. For LACP to function correctly, the switch ports to which the node is connected must also be configured to support LACP. If the ports on the upstream switch are not configured for LACP, the bond will not be able to establish properly, and the node will not communicate effectively on the network, making it undiscoverable when attempting to expand the cluster.

The absence of an operational LACP configuration could prevent the new node from joining the existing cluster as the node's network interfaces would not be able to pass traffic correctly. This can be verified by checking the switch configuration to ensure that the ports are set to participate in an LACP bond.

The other options, such as a firewall blocking discovery traffic (Option A) or the node being on different VLANs (Option C), are possible causes for a node not being discovered, but given the specific command output provided, the most likely cause is related to the switch port configuration for LACP. Option D, regarding completing LACP configuration after cluster expansion, is not correct because LACP needs to be operational for the node to communicate with the cluster during the expansion process.

Proper LACP configuration is critical for network communication in a Nutanix AHV cluster, and this is covered in detail in the Nutanix AHV and Networking documentation. It outlines the steps for configuring network bonds and LACP on both the AHV hosts and the

connecting network infrastructure.

Question 3

Question Type: MultipleChoice

After the initial configuration and upgrade of NCC, the administrator notices these critical alerts:

- . IPMI 10.7.133.33 is using default password
- . Host 10.7.133.25 is using default password
- . CVM 10.7.133.31 is using default password

Which two initial cluster configuration tasks were missed during the deployment process? (Choose two.)

Options:

- A-** CVM password changes
- B-** BIOS password changes

C- Host password changes

D- Password policy changes

Answer:

A, C

Explanation:

The critical alerts listed are indicating that the default passwords are still in use for IPMI, the host, and the Controller Virtual Machine (CVM). This suggests that the passwords for these components were not changed from the default during the initial cluster configuration and deployment process, which is a critical security practice.

A) CVM password changes: The alert for the CVM using the default password indicates that the CVM password has not been changed. It is a standard security measure to change default passwords to prevent unauthorized access.

C) Host password changes: Similarly, the alert for the host using the default password indicates that the default password for the host has not been updated. This applies to the passwords used to access the hypervisor host directly.

Changing default passwords is a critical step in securing the Nutanix environment. This is highlighted in Nutanix's best practices and security guidelines, which recommend changing default passwords as part of the initial configuration to ensure that the environment is not left vulnerable to unauthorized access due to known default credentials. This process is typically part of the initial setup procedures outlined in the Nutanix documentation for cluster deployment and security configuration.

The IPMI alert also points to the need for changing default passwords, but since IPMI (Intelligent Platform Management Interface) is not specifically mentioned in the provided options, it falls under the broader category of host-level password changes, which would be covered by option C.

BIOS password changes (Option B) and Password policy changes (Option D) are also important but were not directly flagged by the alerts mentioned. BIOS password changes are usually a separate task and not indicated by the alerts given, while password policy changes are related to the policies governing password complexity and rotation rather than the initial password setup.

Question 4

Question Type: Hotspot

An administrator has created several custom alert policies, which are applied to the same entities. Prism Central displays a message that a similar policy exists.

In what order of precedence are overlapping policies evaluated?

Preference

First Preference

Correct Sequence

Select



Select

Policy is applied to a specific entity

Policy is applied to an entity type in a category

Policy is applied to an entity type in a cluster

Policy is applied to all entities of an entity type

Second Preference

Select



Select

Policy is applied to a specific entity

Policy is applied to an entity type in a category

Policy is applied to an entity type in a cluster

Policy is applied to all entities of an entity type

Third Preference

Select



Select

Policy is applied to a specific entity

Policy is applied to an entity type in a category

Policy is applied to an entity type in a cluster

Policy is applied to all entities of an entity type

Answer:

Question 5

Question Type: MultipleChoice

An administrator is implementing a VDI solution. The workload will be a series of persistent desktops in a dedicated storage container within a four-node cluster. Storage optimizations should be set on the dedicated storage container to give optimal performance including during a node failure event.

Which storage optimizations should the administrator set to meet the requirements?

Options:

- A-** Compression only
- B-** Deduplication and Erasure Coding
- C-** Compression and Deduplication
- D-** Compression, Deduplication, and Erasure Coding

Answer:

D

Explanation:

According to the web search results, three storage optimizations that should be set on a dedicated storage container for a VDI solution with persistent desktops are compression, deduplication and erasure coding. Compression is a technique that reduces the size of data by removing redundant information. Deduplication is a technique that eliminates duplicate copies of data blocks and stores only unique blocks. Erasure coding is a technique that increases usable capacity by reducing the replication factor of data blocks and using parity information instead⁷. These three optimizations can improve performance and save storage space for VDI workloads, especially during a node failure event when data needs to be rebuilt from parity information⁸. The administrator can enable these optimizations by using Prism Element web console or ncli command-line interface⁹.

Question 6

Question Type: MultipleChoice

An administrator has a Custom backup application that requires a 2TB disk and runs in Windows. Throughput is considerably lower than expected.

The application was installed on a VM with the following configuration:

* FOU vCPUs with one core/vCPU

* 4GB of Memory

* One 50GB vDisk for the Windows installation

* One 2TB vDisk for the application

What is the recommended configuration change to improve throughput?

Options:

A- Add 4GB of memory to the VM

B- Increase the vCPUs assigned to the VM

C- Span the 2TB disk across four vDisks

D- Increase the number of cores per vCPU

Answer:

C

Explanation:

According to the web search results, one recommended configuration change to improve throughput for a custom backup application that requires a 2TB disk and runs on Windows is to span the 2TB disk across four vDisks. Spanning is a technique that allows you to create a single logical disk from multiple physical disks. Spanning can improve throughput by distributing I/O requests across multiple

disks and reducing contention⁵. To span a disk across four vDisks, the administrator should create four vDisks of equal size (500 GB each) and attach them to the VM. Then, in Windows Disk Management, create a spanned volume from the four vDisks and format it as NTFS⁶.

Question 7

Question Type: MultipleChoice

An administrator is commissioning a Nutanix Enterprise Cloud. Once the user VMs have been deployed and are running, the administrator finds that VMs on the same host are able to communicate, but are unable to communicate between hosts.

What must be changed to enable full inter-VM communications?

Options:

- A-** Change the spanning-tree port type on the switch.
- B-** Change the network the VMs are connected to
- C-** Update the switch to specifically allow VLAN 15
- D-** The VMs need to have static IP addresses.

Answer:

C

Explanation:

According to the web search results, one possible cause of inter-VM communication failure between hosts is that the switch is not configured to allow VLAN 15 traffic. VLAN 15 is the default VLAN ID used by AHV for internal communication between CVMs and hosts. If the switch blocks or drops VLAN 15 packets, it will prevent inter-VM communication across hosts. To enable full inter-VM communication, the administrator should update the switch to specifically allow VLAN 15 traffic on the ports connected to the AHV hosts.

Question 8

Question Type: MultipleChoice

An administrator needs to relocate an AHV cluster to a new datacenter during a maintenance window. The cluster will use the same IPs in the new datacenter.

Which two steps should be taken to prepare for this task? (Choose two.)

Options:

- A- Reconfigure IPMI for the new datacenter
- B- Shut down all user VMs in the cluster
- C- Relocate the linked LDAP server
- D- Stop all Nutanix Files clusters

Answer:

B, D

Explanation:

According to the web search results, two steps that should be taken to prepare for relocating an AHV cluster to a new datacenter during a maintenance window are:

Shut down all user VMs in the cluster: This step is necessary to ensure that there is no data loss or corruption during the relocation process. The user VMs can be shut down either individually or in bulk by using the Prism Element web console or the acli command-line interface¹.

Stop all Nutanix Files clusters: If the AHV cluster hosts any Nutanix Files clusters, they should be stopped before relocating the cluster. Nutanix Files clusters are composed of one or more virtual machines that provide file services to clients. Stopping a Nutanix Files cluster will stop all the file server VMs and release the resources they consume². The Nutanix Files clusters can be stopped by using the Prism Element web console or the ncli command-line interface³.

Question 9

Question Type: MultipleChoice

An administrator responsible for a VDI environment needs to investigate reports of slow logins. The administrator finds that increasing the number of vCPUs from 2 to 4 will reduce the login times. Production workloads are consuming 75% of the host CPU on the cluster. The administrator increases the vCPU count on all of the VDI VMs.

What are two impacts on the cluster? (Choose two)

Options:

- A- Increasing CPU counts will decrease memory utilization
- B- Increase memory utilization%
- C- Increase CPU utilization%
- D- Increase CPU ready%

Answer:

C, D

Explanation:

According to the web search results, the two impacts on the cluster that will result from increasing the vCPU count on all of the VDI VMs are:

Increase CPU utilization%: CPU utilization is the percentage of time that a CPU is busy executing instructions⁵. By increasing the vCPU count on all of the VDI VMs, the administrator will increase the demand for CPU resources on the cluster, which will increase the CPU utilization percentage⁶.

Increase CPU ready%: CPU ready is the percentage of time that a vCPU is ready to run but is waiting for a physical CPU to become available⁵. By increasing the vCPU count on all of the VDI VMs, the administrator will increase the contention for physical CPU resources on the cluster, which will increase the CPU ready percentage⁶. A high CPU ready percentage can indicate performance issues such as latency or slowdowns⁵.

Question 10

Question Type: MultipleChoice

A system administrator needs to add more VMs to their Nutanix cluster.

Which two actions should the administrator perform to determine if the current cluster can accommodate the new VMs? (Choose two)

Options:

- A- Enable Deduplication and Ensure Coding
- B- Utilize Optimize Resources for VM efficiency
- C- Determine utilization with Cluster Runway
- D- Perform an inventory with Life Cycle Management

Answer:

B, C

Explanation:

According to the web search results, the two actions that the administrator should perform to determine if the current cluster can accommodate the new VMs are:

[Utilize Optimize Resources for VM efficiency](#): Optimize Resources is a feature in Prism Central that helps the administrator improve the efficiency and performance of their VMs by identifying and resolving issues such as overprovisioning, inactivity, constraints, or bullying³. By using Optimize Resources, the administrator can reclaim unused resources and optimize the resource allocation for their VMs³.

[Determine utilization with Cluster Runway](#): Cluster Runway is a feature in Prism Central that helps the administrator estimate how long their cluster can continue to run normally based on the current consumption rate of CPU, memory, and storage resources⁴. By using Cluster Runway, the administrator can see how adding more VMs will affect the resource utilization and capacity of their cluster⁴.

Question 11

Question Type: MultipleChoice

Microsegmentation was recently enabled in a Nutanix environment. The administrator wants to leverage Prism Central to create a policy that will block all traffic regardless of direction, between two groups of VMs identified by their category.

Which policy should be used to meet this requirement?

Options:

- A- An Application Security Policy
- B- A Quarantine Policy
- C- A Whitelist-Based Policy
- D- An Isolation Environment Policy

Answer:

D

Explanation:

According to the web search results, the policy that should be used to meet this requirement is an Isolation Environment Policy. An Isolation Environment Policy is a type of security policy that can be created in Prism Central using Flow Network Security, which is a feature that provides microsegmentation and network security for Nutanix environments¹. An Isolation Environment Policy allows the administrator to isolate a group of VMs from another group of VMs based on their categories, and block all traffic between them regardless of direction². This policy can be useful for creating isolated environments for testing, development, or compliance purposes².

Question 12

Question Type: MultipleChoice

An administrator needs to deploy an application with a large amount of data connected via Nutanix volumes.

Which two actions should the administrator take when designing the Volume Group? (Choose two.)

Options:

- A- Distribute workload across multiple virtual disks
- B- Enable RSS (Receive Side Scaling)

- C- Use multiple subnets for iSCSI traffic
- D- Enable thick provisioning on the Volume Group(s)

Answer:

A, B

Explanation:

[According to the Nutanix Volumes - Recommendations And Best Practices web search result3](#), two actions that the administrator should take when designing the Volume Group are:

Distribute workload across multiple virtual disks: Use multiple disks rather than a single large disk for an application. Consider using a minimum of one disk per Nutanix node to distribute the workload across all nodes in a cluster. Multiple disks per Nutanix node may also improve an application's performance. For performance-intensive environments, we recommend using between four and eight disks per CVM for a given workload.

Enable RSS (Receive Side Scaling): Receive-side scaling (RSS) allows the system to use multiple CPUs for network activity. With RSS enabled, multiple CPU cores process network traffic, preventing a single CPU core from becoming a bottleneck. Enabling RSS within hosts can be beneficial for heavy iSCSI workloads. For VMs running in ESXi environments, RSS requires VMXNET3 VNICs. For Hyper-V environments, enable VMQ to take full advantage of Virtual RSS.

To Get Premium Files for NCP-MCI Visit

<https://www.p2pexams.com/products/ncp-mci>

For More Free Questions Visit

<https://www.p2pexams.com/nutanix/pdf/ncp-mci>

