Nutanix.NCP-MCI-6.10.v2025-09-05.q55

Exam Code:	NCP-MCI-6.10
Exam Name:	Nutanix Certified Professional - Multicloud Infrastructure (NCP-MCI v6.10)
Certification Provider:	Nutanix
Free Question Number:	55
Version:	v2025-09-05
# of views:	118
# of Questions views:	550
https://www.exam-tests.com/NCP-MCI-6.10-exam/Nutanix.NCP-MCI-6.10.v2025-09-05.q55.html	

NEW QUESTION: 1
An administrator has been asked tocalculate baseline Capacity Runwayon anewly registered AHV cluster.
The cluster has been running for16 days, butno runway projections are displayed.
Why are no Capacity Runway projections being displayed?
A. Capacity Planning requires at least 30 days of data.
B. Capacity Planning requires at least 21 days of data.
C. Capacity Planning requires at least 3 months of data.
D. Capacity Planning requires at least 6 months of data.
Answer: B (LEAVE A REPLY)
Nutanix Prism Central requires at least 21 days of usage data to generate accurate Capacity Runway projections.
* Option B (21 days) is correct:
* Until 21 days of data is collected, no runway analysis is available.
* Option A (30 days) is incorrect:
* 30 daysis recommended for long-term accuracy, butnot required for initial projections.
* Option C (3 months) and Option D (6 months) are incorrect:
* Extended data collection helps trend accuracy, but runway calculationsbegin after 21 days.
References:
* Nutanix Prism Central Guide #Understanding Capacity Runway Calculations
* Nutanix KB #Why No Capacity Runway Data is Displayed for New Clusters

NEW QUESTION: 2
An administrator receives analert in Prismstating:

"Storage container <container_name> on cluster <cluster_name> will run out of storage resources in approximately 1 day." However, the cluster hasplenty of available space remaining. What configuration setting is causing the container to run out of space while the cluster has space remaining?

A. Advertised Capacity is set too low.

B. Reserved Capacity is set too high.

C. Compression is set too low.

D. Replication Factor is set too high.

Answer: B (LEAVE A REPLY)

Reserved Capacity settings define how much storage is exclusively allocated for a specific container.

* Option B (Reserved Capacity is too high) is correct:

* If too much space is reserved for a container, it can report "out of space" while the cluster still has free capacity.

* Options A, C, and D are incorrect:

* Advertised Capacity, Compression, and RF settings do not directly cause storage exhaustion unless misconfigured with Reserved Capacity.

References:

* Nutanix Storage Best Practices#Configuring Reserved and Advertised Capacity

* Nutanix KB#Troubleshooting Storage Container Out-of-Space Alerts

NEW QUESTION: 3

Due toapplication requirements, an administrator needs to support amulticast configurationin anAHV cluster.

Which AHV feature can be used to optimize network traffic so that multicast traffic is only forwarded to the VMs that need to receive it?

A. LACP

B. UDP

C. IGMP Snooping

D. Network Segmentation

Answer: C (LEAVE A REPLY)

Multicast traffic can generate unnecessary overhead if it is not properly managed.IGMP Snooping (Option C)ensures thatmulticast packets are only sent to VMs that have requested them, rather than broadcasting to all VMs.

* Option C (IGMP Snooping) is correct:

* Itreduces unnecessary multicast trafficby ensuring that only subscribed VMs receive the packets.

* It is supportednatively in AHV networking.

* Option A (LACP) is incorrect:

* Link Aggregation Control Protocol (LACP)improves bandwidth and redundancy but doesnot control multicast traffic.

* Option B (UDP) is incorrect:
* UDP (User Datagram Protocol) is a transport protocol, not a network optimization feature.
* Option D (Network Segmentation) is incorrect:
* Segmentation (VLANs, VPCs) isolates networksbut does not optimizemulticast traffic specifically.
References:
* Nutanix AHV Networking Guide #Enabling IGMP Snooping
* Nutanix Bible #Network Traffic Optimization in AHV
* Nutanix KB #Best Practices for Multicast Traffic in AHV

NEW QUESTION: 4
What happens if an agent VM is powered off and then manually started on another host?
A. Agent VM become unresponsive.
B. Agent VM cannot be migrated back to the original host.
C. Agent VM migrates back to the original host once it's powered on.
D. Agent VM migrates to another host automatically
Answer: (SHOW ANSWER)
Agent VMs, such as CVMs (Controller VMs) or Witness VMs, have strict affinity and anti-affinity rules to ensure they remain on specific hosts and maintain data consistency and high availability. If an agent VM is powered off and then manually started on another host, it becomes unresponsive because it breaks these rules.
From theNutanix Enterprise Cloud Administration (ECA)course materials:
"Agent VMs have specific configuration and affinity constraints. Manually starting them on another host violates these constraints, resulting in the agent VM becoming unresponsive to the cluster."
Further clarification:
"The cluster expects the agent VM to be on a particular host. Moving it manually to another host breaks this expectation and causes the VM to be unable to properly join the cluster services, leading to an unresponsive state." Therefore, it is essential to avoid manually starting agent VMs on different hosts, as doing so can disrupt cluster services.

NEW QUESTION: 5
A consultant is configuringsyslog monitoringand wants to receiveCRITICAL logsfrom theAudit module.
Which severity level setting should be configured to get the desired output?
A. 0
B. 2
C. 5
D. 7
Answer: B (LEAVE A REPLY)
Syslog severity levels follow a standard numerical system, wherelower numbers indicate higher severity.

* Option B (Severity Level 2) is correct:
* Level 2 = CRITICAL# Used forserious conditions requiring immediate attention.
* Audit logs often generate Critical logsrelated tosecurity, access violations, or major failures.
A screenshot of a computer Description automatically generated

Syslog Severity Level	Meaning
0	Emergency (System is unusable)
1	Alert (Action must be taken immediately)
2	Critical (Severe conditions, failures, or security issues)
3	Error (General errors, software failures)
4	Warning (Potential issues)
5	Notice (Normal but significant events)
6	Informational (General system messages)
7	Debug (Detailed debugging information)

* Options A (0), C (5), and D (7) are incorrect:
* 0 = Emergency (Too severe, only used for full system failures).
* 5 = Notice (Not critical, used for general system events).
* 7 = Debug (Not relevant for critical logs).
References:
Nutanix Security Guide#Syslog Integration and Severity Levels
Nutanix KB#Configuring Syslog for Prism Central Audit Logs

NEW QUESTION: 6
An administrator receives complaints about VM performance.

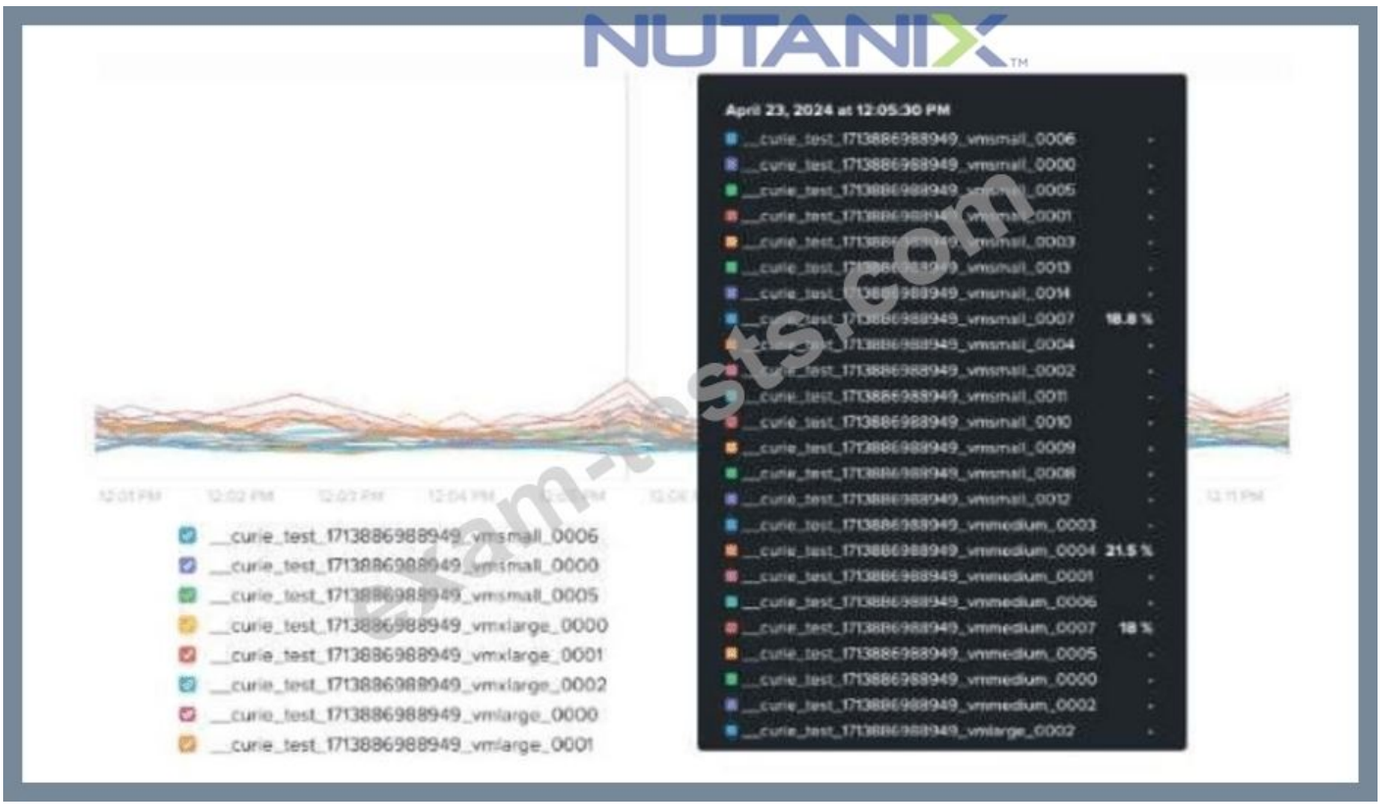

After reviewing theVM's CPU Ready Timedata shown in the exhibit, which step should the administrator take to diagnose the issue further?

A. Check the number of vCPUs assigned to each CVM.

B. Review host CPU utilization.

C. Assess cluster SSD capacity.

D. Enable VM memory oversubscription.

Answer: (SHOW ANSWER)

Understanding the Issue

The administrator is investigating VM performance complaints and is analyzing CPU Ready Time data.

* CPU Ready Time is a crucial metric in Nutanix and virtualization environments (AHV, ESXi, or Hyper- V).
* It measures the amount of time a VM is waiting for CPU scheduling due to resource contention.
* High CPU Ready Time indicates that VMs are ready to run but are waiting because the host lacks available CPU resources.

Analysis of the Exhibit

* The graph shows CPU Ready Time spikes for multiple VMs.
* Some VMs have CPU Ready Time exceeding 18% to 21.5%, which is very high.
* A healthy CPU Ready Time should be below 5%.
* Values above 10% indicate CPU contention, and anything above 20% is critical and requires immediate troubleshooting.

Evaluating the Answer Choices

#(A) Check the number of vCPUs assigned to each CVM. (Incorrect)

* CVMs (Controller VMs) have fixed CPU allocation, and modifying their vCPU count is not recommended unless advised by Nutanix Support.
* The issue is related to VM CPU contention, not CVM configuration.
#(B) Review host CPU utilization. (Correct Answer)
* High CPU Ready Time suggests CPU overcommitment or host saturation.
* The administrator should check host CPU usage in Prism Central to determine if the cluster is overloaded.
* If host CPU usage is consistently above 85-90%, VMs are competing for CPU resources, leading to high CPU Ready Time.
#(C) Assess cluster SSD capacity. (Incorrect)
* SSD capacity impacts storage performance (latency, read/write speeds) but does not affect CPU Ready Time.
* High CPU Ready Time is a CPU scheduling issue, not a storage bottleneck.
#(D) Enable VM memory oversubscription. (Incorrect)
* Memory oversubscription does not impact CPU scheduling.
* Enabling memory oversubscription affects RAM allocation, but CPU Ready Time is strictly related to CPU contention.
Next Steps to Diagnose & Resolve the Issue
* Review Host CPU Utilization:
* Navigate to Prism Central # Analysis # CPU Usage per Host.
* Identify hosts experiencing high CPU load.
* Check VM vCPU Allocation:
* Ensure that VMs do not have excessive vCPUs assigned, which can lead to scheduling inefficiencies.
* Overprovisioning vCPUs can cause unnecessary contention.
* Balance Workload Across Hosts:
* Use Nutanix AHV DRS (Dynamic Scheduling) or VMware DRS to redistribute VMs across hosts.
* Check if certain hosts are overloaded while others have spare CPU capacity.
* Consider Scaling Out the Cluster:
* If CPU usage is consistently high, adding more nodes may be required to reduce CPU contention.
Multicloud Infrastructure References & Best Practices
* CPU Ready Time Best Practices:
* Keep CPU Ready Time below 5%.
* Avoid overcommitting vCPUs on heavily loaded hosts.
* Monitor Prism Central Runway Metrics to predict future CPU resource needs.
* Nutanix AHV CPU Scheduling Optimization:
* Ensure proper VM sizing (avoid excessive vCPU allocation).
* Balance workloads using Nutanix AHV DRS.
References:
Nutanix Prism Central: Performance Analysis and CPU Metrics

Nutanix Bible: VM Performance and Resource Management
Nutanix KB: Troubleshooting High CPU Ready Time in AHV

NEW QUESTION: 7
The customer expects to maintain a cluster runway of 9 months. The customer doesn't have a budget for 6 months but they want to add new workloads to the existing cluster.

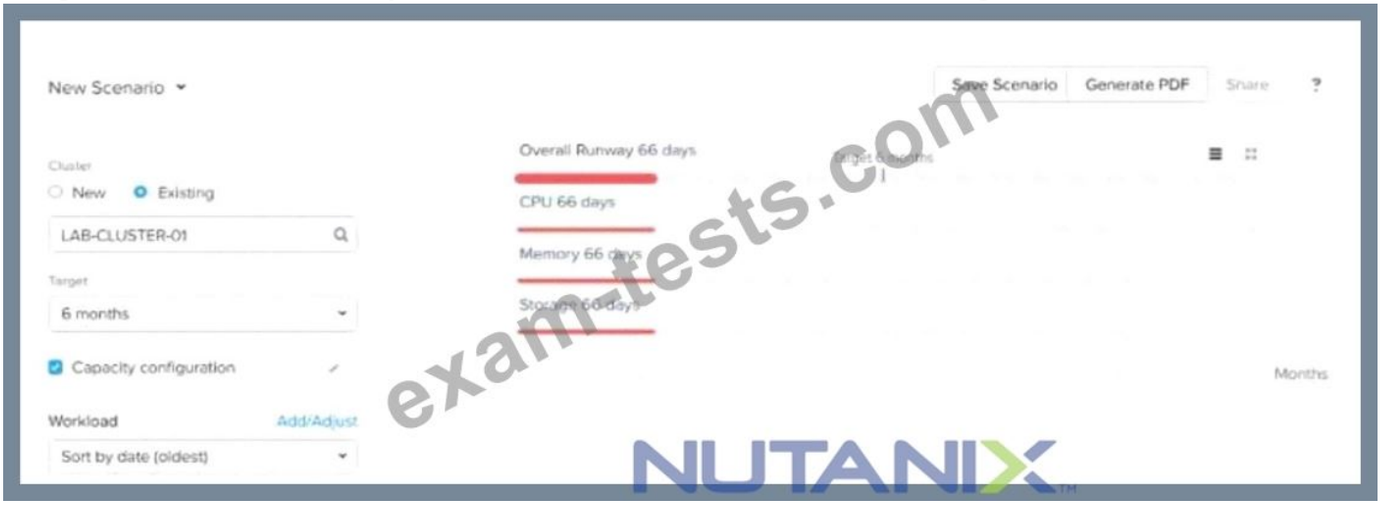

Based on the exhibit, what is required to meet the customer's budgetary timeframe?
A. Add resources to the cluster.
B. Postpone the start of new workloads.
C. Delete workloads running on the cluster.
D. Change the target to 9 months.
Answer: A (LEAVE A REPLY)
The exhibit shows that theoverall runway is only 66 days, meaning that the current clusterdoes not have enough capacity to sustain workloads for 6 months, let alone 9 months.
* The best solution is toadd resources to the cluster(Option A), such as CPU, memory, or storage, to extend the runway.
* Postponing new workloads(Option B) may help in the short term but does not align with the business need to continue adding workloads.
* Deleting workloads(Option C) is not a viable option because the customer wants to add more, not remove them.
* Changing the target to 9 months(Option D) does not change the actual resource constraints; it only alters the target timeframe.
References:
* Nutanix Prism Central # Capacity Planning and Runway Analysis
* Nutanix Bible # Cluster Resource Management and Scaling
* Nutanix Support KB # How to Extend Cluster Runway with Resource Scaling

NEW QUESTION: 8
A new employee has inherited a partially configured Disaster Recovery (DR) schema. Source workloads have been identified and Nutanix Guest Tools has been installed.

There are two Protection Polices in place, one with an asynchronous schedule with a 1-hour RPO and a second policy utilizing synchronous replication. All of these workloads need to be recovered at a DR location and this will be orchestrated by Prism Central Recovery Plans.
What is the best way to setup this recovery orchestration?
A. Identify the workload startup order and create Recovery Plans corresponding to the startup order.
B. Setup two Recovery Plans, one for the asynchronous replication and one for the synchronous replication.
C. Setup a single Recovery Plan utilizing stages of recovery delays as needed.
D. Setup a Recovery Plan for the asynchronous replication and convert the synchronous replication to a Protection Domain.
Answer: (SHOW ANSWER)

NEW QUESTION: 9
Refer to the Exhibit:

Host 1 (128 GB)			
VMs	VM Memory (GB)	Utilized Memory (GB)	Unutilized Memory (GB)
VM1	64 GB	48 GB	16 GB
VM2	32 GB	20 GB	12 GB
VM3	32 GB	24 GB	8 GB
Total	128 GB	92 GB	36 GB

An administrator needs to create two virtual machines: VM4 and VM5 that leverage the memory over-commit feature.
Once VM4 is created and running, the administrator notices that it uses only 28GB of RAM.
What will be the maximum RAM that can be allocated to VM5 so that it can be powered on?
A. 4GB
B. 8GB
C. 16GB
D. 32GB
Answer: B (LEAVE A REPLY)
Understanding the Exhibit & Memory Allocation
* Thehost has 128GB of physical RAM.
* Thecurrent memory allocationacrossthree VMs (VM1, VM2, VM3) is 128GB, but only92GB is actually utilized.
* This means there is36GB of unutilized memory available for allocation.
Step-by-Step Breakdown
* Existing Memory Usage Before Adding VM4
* Total Physical RAM:128GB
* Used by running VMs (VM1, VM2, VM3):92GB
* Unutilized Memory Available:36GB
* After Creating and Running VM4

* VM4 is allocatedmemory but only utilizes 28GB.
* The table does not show VM4's allocated RAM, but assuming it was given a reasonable allocation, it must have been taken from the36GB unutilized memory pool.
* IfVM4 uses 28GB, theremaining unutilized memory is now (36GB - 28GB) = 8GB.
* Maximum Memory Allocation for VM5
* Sinceonly 8GB remains unutilized, the maximum memory VM5 can be allocated while still allowing it to power on is8GB.
Evaluating the Answer Choices
* (A) 4GB#(Incorrect)
* More memory (8GB) is available, so limiting to 4GB is unnecessary.
* (B) 8GB#(Correct)
* Theremaining unutilized memory after VM4 is 8GB, so VM5 can be allocated up to8GBwhile ensuring it can power on.
* (C) 16GB#(Incorrect)
* Only8GB is left, so 16GB isnot possible.
* (D) 32GB#(Incorrect)
* There isnot enough unutilized memoryto allocate 32GB.
Key Concept: Nutanix Memory Overcommit
* Nutanix AHV supportsmemory overcommit, meaning VMs can be allocated more memory than physically available usingmemory ballooning and swapping.
* However,to power on VM5 without impacting performance, it must fit within the available unutilized memory, which is8GB.

NEW QUESTION: 10
What is required to create a category in Nutanix?
A. A name and a value
B. A policy and an entity
C. A service and a scope
D. A catalog and a template
Answer: (SHOW ANSWER)
Categories in Nutanix are used to group resources and require only a name and a value for definition.
* Option A (A name and a value) is correct:
* Categoriesrequire a name (e.g., "Production VMs")and a value (e.g., "Tier 1").
* These are then applied to VMs, storage, and other resources forpolicy-based management.
* Option B (Policy and Entity) is incorrect:
* Policiesuse categories but are not required to define a category.
* Option C (Service and Scope) is incorrect:
* Categoriesdo not require servicesor a defined scope.
* Option D (Catalog and Template) is incorrect:
* These apply toself-service provisioning, not categories.

References:
* Nutanix Prism Central Guide #Creating and Managing Categories
* Nutanix KB #Using Categories for RBAC and VM Grouping

NEW QUESTION: 11
Due toapplication requirements, an administrator needs to support amulticast configurationin anAHV cluster.
Which AHV feature can be used to optimize network traffic so that multicast traffic is only forwarded to the VMs that need to receive it?
A. IGMP Snooping
B. UDP
C. Network Segmentation
D. LACP
Answer: A (LEAVE A REPLY)

NEW QUESTION: 12
An administrator is tasked with protection of a business critical application. The application is running on a Linux VM and is using a custom DB that require application consistent snapshots for data integrity.
An administrator has written a pre_freeze and post_thaw scripts and placed them under /usr/local/sbin/.
During protection domain scheduled run an alert is generated:
Execution of the PostThaw Script Failed
Which two resolution steps could an administrator conduct to fix the issue? (Choose two.)
A. Ensure that scripts have nutanix user ownership and admin access.
B. Review the NGT logs under /usr/local/sbin/post_thaw.
C. Ensure NGT service is up and running.
D. Execute scripts manually and ensure they succeed
Answer: (SHOW ANSWER)
To resolve issues with application-consistent snapshots (script execution failures), two critical actions are needed:
* NGT Service Status:From the ECA materials:
"The Nutanix Guest Tools (NGT) service is responsible for executing pre-freeze and post-thaw scripts for application-consistent snapshots. If NGT is not running, these scripts will not execute."
* Script Validation:
"Before relying on scheduled snapshot runs, execute pre-freeze and post-thaw scripts manually to ensure they complete successfully. This helps to rule out script logic or permission issues."
Ownership or log review of the script path itself (A and B) are secondary and typically not root-cause resolution steps.

NEW QUESTION: 13

An administrator want to create a VM with memory overcommit features enabled in Nutanix environment.
Which statement best describes how the administrator will perform this VM creation?
A. Memory overcommit can only be updated using the Prism Element Web Console once VM created.
B. Memory overcommit can be enabled while creating a VM using Prism Element Web Console.
C. Memory overcommit can not be enabled for VM from the Prism Central console.
D. Memory overcommit can only be updated using the Prism Central console.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 14
An administrator started anLCM upgrade of the AHV hostsbut realized that theupgrade would exceed the planned maintenance window.
Which feature should be leveraged to prevent additional updates from occurring?
A. Cancel the LCM tasks via theErgon command line (ecli).
B. Run thelcm_task_cleanup.pyscript.
C. RestartGenesison the cluster to restart the LCM service.
D. Use theStop Update feature in LCM.
Answer: D (LEAVE A REPLY)
When performing aLife Cycle Manager (LCM) upgrade, therecommended way to stop the process is to use the "Stop Update" feature in LCM (Option D).
* Option A (Cancel via Ergon ecli)is not a recommended approach since manually interfering with running tasks can cause inconsistencies.
* Option B (lcm_task_cleanup.py script)is used for post-upgrade cleanup but does not stop ongoing updates.
* Option C (Restarting Genesis)does not stop an LCM upgrade and can cause instability.
References:
* Nutanix Life Cycle Manager (LCM) User Guide
* Nutanix KB:Best Practices for Stopping and Restarting LCM Tasks
* Nutanix Prism Central #LCM Feature Documentation

NEW QUESTION: 15
Which storage attributes do Storage Policies manage?
A. Storage Containers and Volume Groups
B. Replication Factor and Encryption
C. Shares and Object Stores
D. Data Protection and Security
Answer: (SHOW ANSWER)
Storage Policies in Nutanix allow administrators to configuredata protection and performance settingsat the storage container level.

* Replication Factor (RF)definesthe number of copies of datastored across nodes for fault tolerance.
* Encryptionensures thatdata at rest is protectedvia Nutanix-native encryption methods.
* Option A (Storage Containers and Volume Groups)refers tostorage organization, not policies.
* Option C (Shares and Object Stores)applies tofile and object storage services, not VM storage policies.
* Option D (Data Protection and Security)is a broad term but does not definespecific policy attributes.
References:
* Nutanix Prism Element #Storage Policies and Replication Factor (RF)
* Nutanix Bible #Storage Fabric and Data Resiliency
* Nutanix KB #Enabling Encryption in Storage Policies

NEW QUESTION: 16
An administrator is preparing for afirmware upgradeon a host and wants to manuallymigrate VMs before executing the LCM upgrade. However,one VM is unable to migratewhile others migrate successfully.
Which action would fix the issue?
A. EnableAcropolis Dynamic Scheduling (ADS)at the cluster level.
B. UpdateLink Layer Discovery Protocol (LLDP).
C. DisableAgent VMwithin the VM configuration options.
D. Configurebackplane port groupsthat are assigned to the CVM.
Answer: C (LEAVE A REPLY)
If aVM is unable to migrate, the most likely cause is that it is anAgent VM(such as aNutanix Witness VM or a VM with special dependencies).
* Option C (Disable Agent VM) is correct:
* SomeAgent VMs are configured to prevent migrationdue to critical roles (e.g., a Witness VM for Metro Availability).
* DisablingAgent VM restrictionsallows it to migrate before a host enters maintenance mode.
* Option A (Enable ADS) is incorrect:
* Acropolis Dynamic Scheduling (ADS)helps with VM placementafter migration, but it does not force an unmigratable VM to move.
* Option B (Update LLDP) is incorrect:
* LLDP is used fornetwork discovery, but it does not impact VM migration behavior.
* Option D (Configure backplane port groups)is incorrect:
* Backplane settingsimpact CVM communication, notVM migration.
References:
* Nutanix Prism Element Guide #Managing Agent VM Settings
* Nutanix Bible #Host Maintenance and VM Live Migration
* Nutanix KB #Troubleshooting VM Migration Failures in AHV

NEW QUESTION: 17
Which predefined view should be leveraged in Prism Central Intelligent Operations to determine which VM is consuming too many resources and causing other VMs to starve?
A. Constrained VMs List
B. Bully VMs List
C. Inactive VMs List
D. Overprovisioned VMs List
Answer: (SHOW ANSWER)
TheBully VMs Listview in Prism Central'sIntelligent Operationsidentifies VMs that are consumingexcessive resources(CPU, memory, or storage IO) and causing resource starvation for other VMs.
From theNutanix ECA documentation:
"Bully VMs List identifies virtual machines with high resource consumption, enabling administrators to investigate and take action to prevent performance degradation in other VMs."
The other options:
* A (Constrained VMs): Shows VMs that are being starved of resources.
* C (Inactive VMs): Shows VMs that are using little to no resources.
* D (Overprovisioned VMs): Shows VMs with allocated resources beyond actual usage.

NEW QUESTION: 18
Which two entities can be categorized? (Choose two.)
A. Storage Containers
B. Alerts
C. Virtual Machines
D. ISO Images
Answer: (SHOW ANSWER)
In Nutanix Prism Central,categoriesallow administrators togroup and organize entitiesfor management, automation, and policy enforcement.
* Alerts (Option B) can be categorizedto group similar system events and create filtering rules.
* Virtual Machines (Option C) can be categorizedto apply security policies, automation tasks, and resource allocation rules.

* Option A (Storage Containers) cannot be categorizedin Prism Central. Storage policies apply at the container level but are not managed via categories.
* Option D (ISO Images) cannot be categorizedbecause ISOs arestatic objects, not active entities.
References:
* Nutanix Prism Central Guide #Working with Categories
* Nutanix Bible #Category-Based Management and Security Policies
* Nutanix KB #Using Categories for VM Management in Prism Central

NEW QUESTION: 19
An administrator needs to compare the performance of two VMs that are running on separate Nutanix clusters.
When creating an Analysis chart in Prism Central Intelligent Operations, the administrator discovers that it is not possible to add VM metrics for IOPS and CPU utilization to the same chart.
How should the administrator resolve this issue?
A. Create separate charts for metrics with different units of measurement
B. Migrate one of the VMs so that both VMs are running on the same cluster.
C. Create an Entity chart instead of a Metric chart.
D. Ensure that both VMs have the same number of vCPUs provisioned.
Answer: (SHOW ANSWER)

NEW QUESTION: 20
Refer to Exhibit:

An administrator sees the alert shown in the exhibit.
What should the administrator do to ensure the nutanix user can no longer SSH to a CVM using a password?
A. Rename the nutanix user.
B. Block port 22 on the CVM firewall.
C. Enable Cluster Lockdown.
D. Delete the nutanix user.
Answer: (SHOW ANSWER)
Understanding the Exhibit & the Alert
The alert states:

* "The cluster is using password-based SSH access for the CVM."
* "Password-based remote login is enabled on the cluster."
* "It is recommended to use key-based SSH access instead of password-based SSH access for better security." This means that the nutanix user can log in to Controller VMs (CVMs) using a password, which is a security risk.
Corrective Action: Enabling Cluster Lockdown
#(C) Enable Cluster Lockdown. (Correct Answer)
* Cluster Lockdown Mode restricts password-based SSH access and forces key-based authentication.
* This prevents users from logging into CVMs using passwords, enhancing cluster security.
* To enable Cluster Lockdown:
* Go to Prism Central or Prism Element.
* Navigate to Settings # Security # Cluster Lockdown.
* Enable Cluster Lockdown Mode.
Evaluating the Other Answer Choices
#(A) Rename the nutanix user. (Incorrect)
* The nutanix user is a built-in system account required for cluster operations.
* Renaming the user will not prevent SSH access via password.
#(B) Block port 22 on the CVM firewall. (Incorrect)
* Blocking port 22 (SSH) will completely disable SSH access, including key-based authentication.
* This may break cluster management and troubleshooting operations.
#(D) Delete the nutanix user. (Incorrect)
* The nutanix user is a critical system account required for cluster functionality.
* Deleting the account will cause serious issues with cluster management.
Multicloud Infrastructure References & Best Practices
* Nutanix Security Best Practices:
* Always use key-based SSH authentication instead of password-based logins.
* Enable Cluster Lockdown Mode to enforce security policies.
* Regularly audit user access to ensure security compliance.
* Cluster Lockdown Benefits:
* Prevents unauthorized SSH access via passwords.
* Enforces public key authentication, reducing brute-force attack risks.
* Strengthens CVM security against potential exploits.
References:
* Nutanix Security Guide #Enabling Cluster Lockdown for SSH Security
* Nutanix KB #Securing SSH Access on Nutanix Clusters

NEW QUESTION: 21
An administrator is trying toconfigure Metro Availabilitybetween NutanixESXi-based clusters. However, theCompatible Remote Sitesscreen does not list all required storage containers.

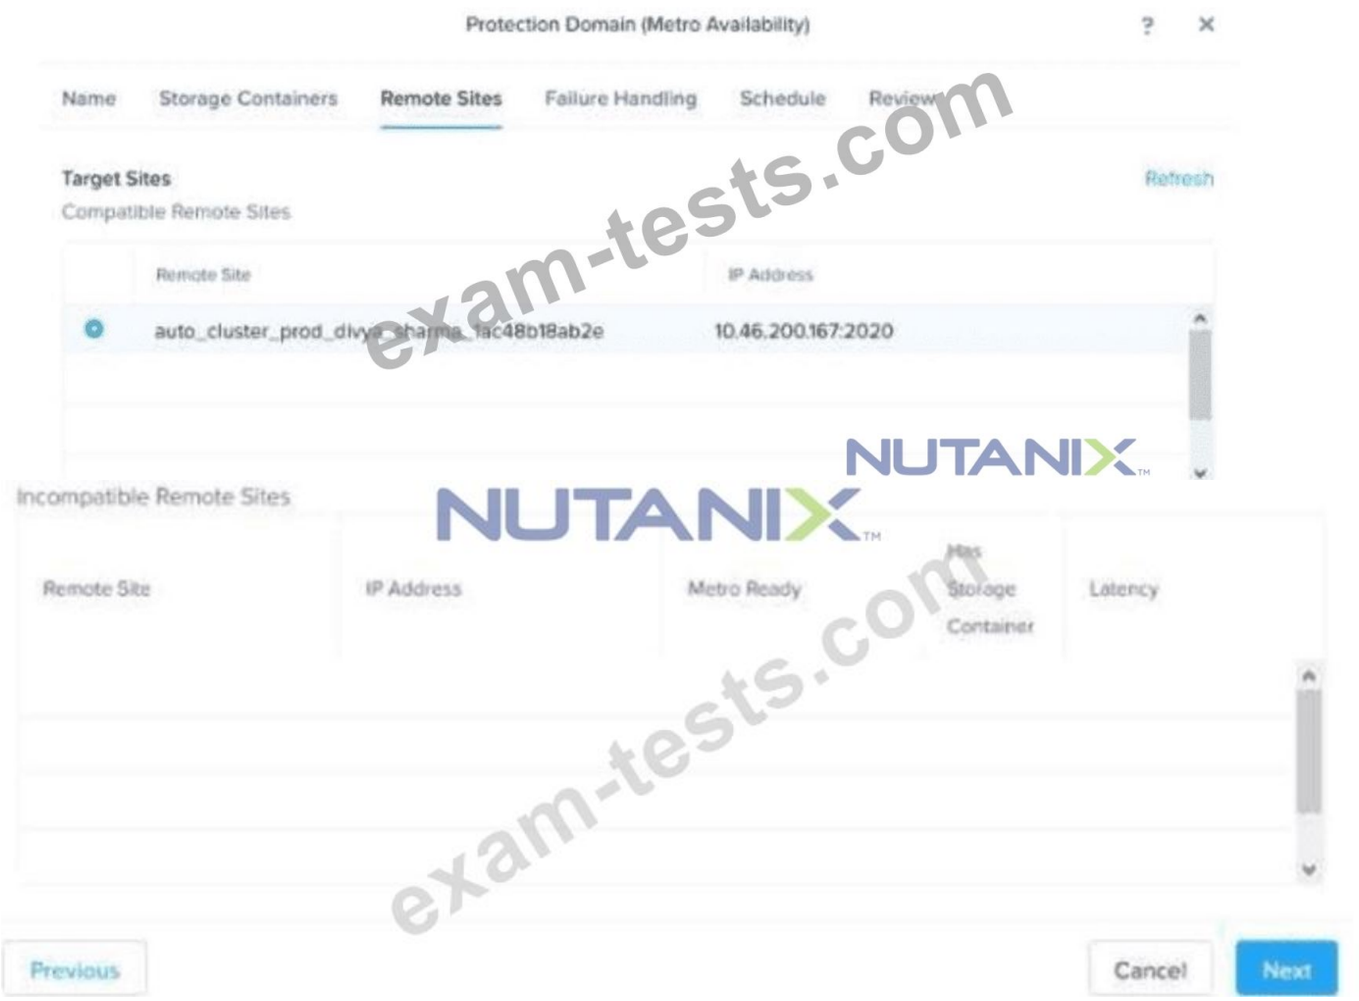

Which two reasons could be a cause for this issue? (Choose two.)

A. Source and destination hardware are from different vendors.

B. The remote site storage container has compression enabled.

C. The destination storage container is not empty.

D. Both storage containers must have the same name.

Answer: C,D (LEAVE A REPLY)

Metro Availability in Nutanix requires that the primary and secondary storage containers be configured identicallyto ensuredata replication consistency.

* Option C (The destination storage container is not empty) is correct:
* Theremote storage container must be emptybefore Metro Availability can be enabled.
* Existing data can cause conflictsand prevent it from appearing in the "Compatible Remote Sites" list.
* Option D (Both storage containers must have the same name) is correct:
* Metro Availability requires thatstorage containers have identical names across clusters.
* If namesdo not match, thestorage container will not be listed as compatible.
* Option A is incorrect: Metro Availability works regardless of hardware vendor differences.
* Option B is incorrect:Compression does not affect compatibilitybut may impact performance.

References:

Nutanix Metro Availability Deployment Guide

Nutanix Best Practices forConfiguring Remote Sites for Metro Availability Nutanix KB #Troubleshooting Storage Container Issues in Metro Availability

NEW QUESTION: 22
An administrator would like to plan for new project-related growth.
New project workload requirements have been included for a cluster named ClusterXYZ:
2 Medium Sized SQL Servers
10 VMs with 16Gb RAM, 4 vCPU, 100GB Storage
Which two additional information items should be added to the capacity planning scenario to provide a proper capacity runway expectation? (Choose two.)
A. Date(s) workload(s) will be added
B. Storage compression ratio(s) for new workload
C. Change in Demand percentage
D. Existing cluster hardware specifications
Answer: (SHOW ANSWER)

NEW QUESTION: 23
An administrator wants toenable application discoveryon aNutanix clusterto monitor applications.
APrism Central instanceis alreadyconfigured with sufficient CPU and memory.
What other prerequisites must be met before enabling application discovery? (Choose two.)
A. Sufficient Prism Central VM resources
B. Internet connection
C. API key and key ID
D. Network controller is enabled
Answer: A,B (LEAVE A REPLY)
Application discovery in Prism Central requires sufficient Prism Central resources and an active internet connection to retrieve application signatures.
* Option A (Sufficient Prism Central VM resources) is correct:
* Prism Central needsadequate CPU and memoryto process application signatures and discovery data.
* Option B (Internet connection) is correct:
* Application discovery depends on Nutanix'scloud databaseto recognize application patterns.
* Option C (API key and key ID) is incorrect:
* API keys are not requiredfor automatic application discovery.
* Option D (Network controller enabled) is incorrect:
* Application discoverydoes not depend on a network controller feature.
References:
* Nutanix Prism Central Guide #Enabling and Using Application Discovery
* Nutanix KB #Requirements for Application Discovery in Prism Central

NEW QUESTION: 24

In afive-node cluster, an administrator noticed thatthree VMsare consuming too many resources on asingle host.
Acropolis Dynamic Scheduling (ADS) is not able to migrate these VMs.
What is the most likely reason preventing ADS from migrating these VMs?
A. VMs use a Volume Group.
B. VMs use GPU pass-through.
C. VM-VM anti-affinity policy is set.
D. VMs use external Network Attached Storage.
Answer: B (LEAVE A REPLY)
VMs using GPU pass-through cannot be live-migrated because they are directly tied to a physical GPU on a specific host.
* Option B (VMs use GPU pass-through) is correct:
* Pass-through devices (such as GPUs) are directly assigned to VMs, makingmigration impossibleunless the VM is powered off first.
* Option A (VMs use a Volume Group) is incorrect:
* Volume Groupssupport live migrationunless they are configured incorrectly.
* Option C (VM-VM anti-affinity) is incorrect:
* Anti-affinity rules prevent two specific VMs from running together, but do not prevent migration.
* Option D (VMs use external NAS) is incorrect:
* Using NAS does not block VM migration, as Nutanixsupports shared storage across hosts.
References:
Nutanix AHV Best Practices#GPU Pass-through and VM Migration Limitations Nutanix KB#Why Can't I Live Migrate a VM with GPU Passthrough?

NEW QUESTION: 25
Refer to Exhibit:

NUTANIX™

After adding new workloads, why isOverall Runway below 365 daysandthe scenario still shows the cluster is in good shape?

A. Because Storage Runway is still good.

B. Because new workloads are sustainable.

C. Because there are recommended resources.

D. Because the Target is 1 month.

Answer: (SHOW ANSWER)

InNutanix Capacity Planning,Overall Runwayrepresents how long the cluster can support current and new workloads before resources are exhausted.

* Even if the runway is below 365 days, the system considers the cluster to be in good shape if new workloads are sustainable (Option B).
* Option A is incorrect: Storage runway alone is not the only factor; CPU and memory are equally important.
* Option C is incorrect: The presence of recommended resources does not mean the cluster is in good shape.
* Option D is incorrect: The target of 1 month affects projections but does not explain why the cluster is in good shape.

References:

* Nutanix Prism Central # Capacity Runway and Planning
* Nutanix Bible # Workload Placement and Cluster Sizing
* Nutanix Support KB # Capacity Planning Best Practices

NEW QUESTION: 26

What happens when a VM is associated with multiple VM-Host affinity policies?

A. All policies are applied simultaneously.

B. The oldest policy is applied

C. The VM is automatically removed from all policies.

D. The newest policy takes precedence.

Answer: (SHOW ANSWER)

NEW QUESTION: 27

An administrator needs tooptimize a VM's storageby leveragingcompression features. The VM'svDisksare currently stored in adefault storage containerwithno optimizations enabled.

How should the administrator proceed?

A. Migrate vDisks to the Production storage container.

B. Recreate the VM in the Production storage container and copy data.

C. Migrate the VM to the Production storage container.

D. Recreate the vDisk in the Production storage container and copy data.

Answer: (SHOW ANSWER)

Moving vDisks to a storage container with compression enabled ensures better data efficiency without downtime.

* Option A (Migrate vDisks) is correct:
* vDisk migration is non-disruptiveand allows compression settings to be applied dynamically.
* Option B (Recreate the VM) is incorrect:
* Rebuilding the VM is unnecessaryand would causedowntime.
* Option C (Migrate the VM) is incorrect:
* VM migration does not guarantee that only vDisks move, and it may disrupt performance.
* Option D (Recreate vDisk) is incorrect:
* This method ismanual and time-consuming, while Nutanix provides anautomated approach.
References:
* Nutanix Storage Optimization Guide#Enabling Compression on Existing vDisks
* Nutanix KB#Migrate vDisks Between Storage Containers for Optimization

NEW QUESTION: 28
Which task should be performed first when upgrading host memory?
A. Gracefully stop the host by using the out of band management interface.
B. Remove node from the cluster.
C. Execute "shutdown -h now" from the AHV command line interface.
D. Place node into the maintenance mode
Answer: (SHOW ANSWER)
The Nutanix ECA course provides detailed procedures for performing hardware upgrades, such as adding host memory, to ensure cluster stability and data availability. Upgrading host memory requires safely preparing the node to avoid disrupting running VMs or cluster operations.
Extract from Nutanix Enterprise Cloud Administration (ECA) Course Documents:
* Module: Cluster Management, Section: Hardware Upgrades"Before performing hardware upgrades, such as adding host memory, the node must be placed into maintenance mode. This ensures that all VMs are migrated to other nodes and the host is safely isolated from cluster operations."
* Module: Host Maintenance, Section: Maintenance Mode"Placing a node into maintenance mode is the first step for hardware upgrades. Maintenance mode migrates all VMs to other nodes, stops the Controller VM (CVM), and prepares the host for safe shutdown or hardware changes."
Explanation of Options:
* A. Gracefully stop the host by using the out of band management interfaceThis is incorrect. Stopping the host via the out-of-band management interface (e.g., IPMI or iLO) without first entering maintenance mode risks disrupting running VMs and cluster services. The ECA course warns:"Shutting down a host without maintenance mode can cause VM crashes and data unavailability, as VMs are not migrated."
* B. Remove node from the clusterThis is incorrect. Removing a node from the cluster is a permanent action that detaches it from the cluster's metadata and storage pool, requiring re-imaging to rejoin. It is not appropriate for a temporary hardware upgrade like adding memory. The ECA course states:" Removing a node is not required for hardware upgrades and should be avoided, as it disrupts cluster configuration."

* C. Execute "shutdown -h now" from the AHV command line interfaceThis is incorrect. Running shutdown -h now on the AHV host without entering maintenance mode will abruptly stop the host, potentially crashing VMs and disrupting cluster operations. The ECA course notes:"Directly shutting down a host via CLI without maintenance mode risks data loss and service disruption."
* D. Place node into maintenance modeThis is the correct answer. Placing the node into maintenance mode is the first step for hardware upgrades, as it safely migrates all VMs to other nodes, stops the CVM, and prepares the host for shutdown or hardware changes. The ECA course emphasizes that maintenance mode ensures cluster stability during upgrades.
* Supporting Extract:"To upgrade host memory, place the node into maintenance mode using Prism Element or the CLI command ncli host maintenance_mode. This ensures safe VM migration and host isolation." Additional Context from ECA:
* Maintenance Mode Process: In Prism Element, maintenance mode can be enabled underHardware > Host > Enter Maintenance Mode. The process automatically migrates VMs using live migration, stops the CVM, and isolates the host. For AHV, the CLI command is ncli host maintenance_mode id=<host_id> enable=true.
* Memory Upgrade: After entering maintenance mode, the host can be safely powered off to add memory, then powered back on and exited from maintenance mode.
Supporting Reference from Web Results:
The Nutanix Bible (https://www.nutanix.com/go/the-nutanix-bible) confirms:"Maintenance mode is the required first step for host hardware upgrades, ensuring VMs are migrated and the node is isolated before changes like memory upgrades."

NEW QUESTION: 29
An administrator needs tooptimize a VM's storageby leveragingcompression features. The VM'svDisksare currently stored in adefault storage containerwithno optimizations enabled.
How should the administrator proceed?
A. Migrate vDisks to the Production storage container.
B. Recreate the VM in the Production storage container and copy data.
C. Migrate the VM to the Production storage container.
D. Recreate the vDisk in the Production storage container and copy data.
Answer: A (LEAVE A REPLY)
Moving vDisks to a storage container with compression enabled ensures better data efficiency without downtime.
* Option A (Migrate vDisks) is correct:
* vDisk migration is non-disruptiveand allows compression settings to be applied dynamically.
* Option B (Recreate the VM) is incorrect:
* Rebuilding the VM is unnecessaryand would causedowntime.
* Option C (Migrate the VM) is incorrect:
* VM migration does not guarantee that only vDisks move, and it may disrupt performance.
* Option D (Recreate vDisk) is incorrect:
* This method ismanual and time-consuming, while Nutanix provides anautomated approach.

References:
Nutanix Storage Optimization Guide#Enabling Compression on Existing vDisks Nutanix KB#Migrate vDisks Between Storage Containers for Optimization

NEW QUESTION: 30
Refer to Exhibit:

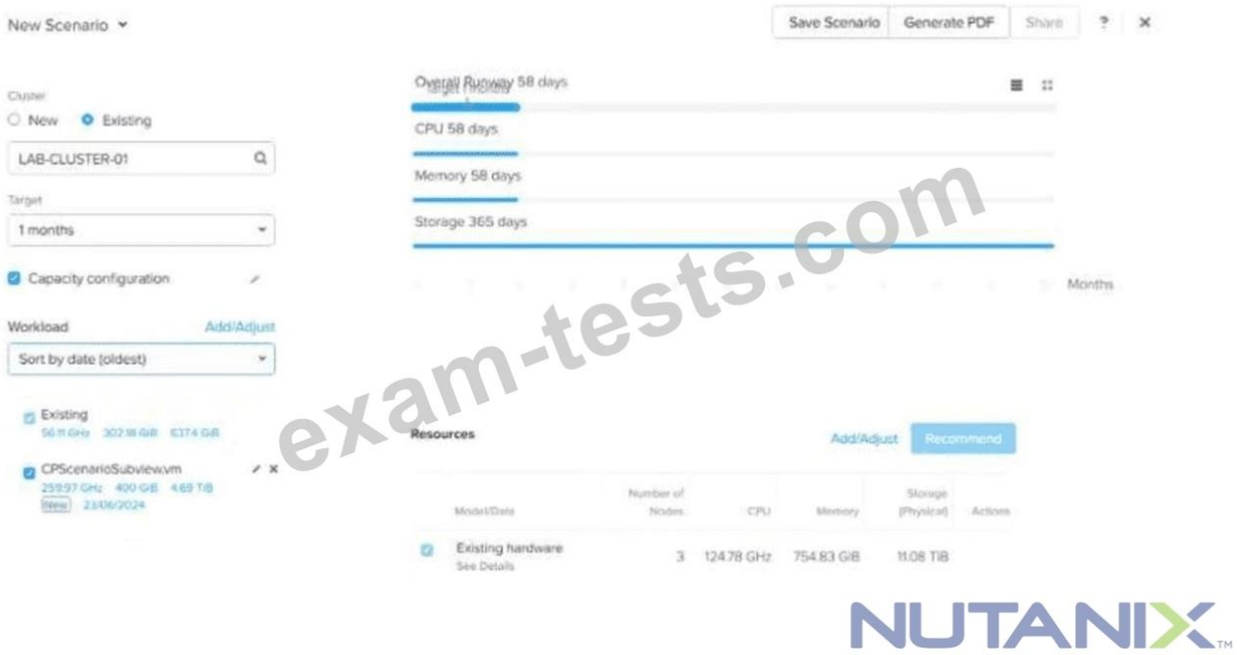

After adding new workloads, why isOverall Runway below 365 daysand the scenario still shows the cluster is in good shape?
A. Because Storage Runway is still good.
B. Because new workloads are sustainable.
C. Because there are recommended resources.
D. Because the Target is 1 month.
Answer: B (LEAVE A REPLY)
InNutanix Capacity Planning,Overall Runwayrepresents how long the cluster can support current and new workloads before resources are exhausted.
* Even if the runway is below 365 days, the system considers the cluster to be in good shape if new workloads are sustainable (Option B).
* Option A is incorrect: Storage runway alone is not the only factor; CPU and memory are equally important.
* Option C is incorrect: The presence of recommended resources does not mean the cluster is in good shape.
* Option D is incorrect: The target of 1 month affects projections but does not explain why the cluster is in good shape.
References:
Nutanix Prism Central # Capacity Runway and Planning
Nutanix Bible # Workload Placement and Cluster Sizing

Nutanix Support KB # Capacity Planning Best Practices

NEW QUESTION: 31
An administrator has configured Metro Availability but a few hours later got an NCC warning:
Node x.x.X.X:
WARN: Break replication timeout of Metro protection domain 'M1' is below the recommended minimum.
What is a possible resolution for this issue?
A. Update the break_replication_timeout to 10 seconds.
B. Update the break_replication_timeout to 5 milliseconds.
C. Update the break_replication_timeout to 15 milliseconds.
D. Update the break_replication_timeout to 15 seconds
Answer: D (LEAVE A REPLY)
The Nutanix ECA course addresses Metro Availability, a high-availability feature that provides synchronous replication between two Nutanix clusters for zero Recovery Point Objective (RPO) and near-zero Recovery Time Objective (RTO). The NCC warning about thebreak_replication_timeoutbeing below the recommended minimum indicates a configuration issue that could affect the stability of Metro Availability.
Thebreak_replication_timeoutparameter determines how long the Protection Domain (PD) waits before breaking replication if connectivity between the Metro clusters is disrupted.
Extract from Nutanix Enterprise Cloud Administration (ECA) Course Documents:
* Module: Data Protection, Section: Metro Availability Configuration"Metro Availability uses synchronous replication to ensure data consistency between two clusters. The break_replication_timeout parameter defines the timeout period for replication. The recommended minimum value is 15 seconds to prevent premature replication breaks due to transient network issues."
* Module: Nutanix Cluster Check (NCC), Section: Metro Availability Alerts"An NCC warning indicating that the break_replication_timeout for a Metro Protection Domain is below the recommended minimum suggests the timeout is too low, risking unnecessary replication breaks. The recommended setting is 15 seconds to balance stability and responsiveness in Metro Availability setups." Explanation of Options:
* A. Update the break_replication_timeout to 10 secondsThis is incorrect. A timeout of 10 seconds is below the recommended minimum of 15 seconds, as specified in the ECA course. Setting the timeout too low increases the risk of replication breaking due to transient network latency or jitter, which could disrupt Metro Availability and cause unnecessary failovers. The ECA documentation warns:"A break_replication_timeout below 15 seconds may lead to frequent replication breaks, reducing the reliability of Metro Availability."
* B. Update the break_replication_timeout to 5 millisecondsThis is incorrect. A timeout of 5 milliseconds is far too low and impractical for Metro Availability, as even minor network delays would trigger replication breaks. The ECA course does not support millisecond-level timeouts and explicitly recommends 15 seconds as the minimum. Such a low value would destabilize the Metro

setup, as noted:"Extremely low timeout values are not supported, as they cause replication to break under normal network conditions."
* C. Update the break_replication_timeout to 15 millisecondsThis is incorrect. A timeout of 15 milliseconds is still significantly below the recommended minimum of 15 seconds. Similar to option B, this setting would cause replication to break too quickly, undermining the purpose of Metro Availability. The ECA course clarifies:"Timeouts in milliseconds are not recommended for Metro Availability, as they do not account for typical network latency in synchronous replication setups."
* D. Update the break_replication_timeout to 15 secondsThis is the correct answer. The ECA course explicitly recommends abreak_replication_timeoutof 15 seconds as the minimum to ensure Metro Availability remains stable. This value allows the system to tolerate transient network issues without prematurely breaking replication, maintaining data consistency and availability. The NCC warning indicates the current timeout is below this threshold, and updating it to 15 seconds resolves the issue.
* Supporting Extract:"To resolve NCC warnings about break_replication_timeout, set the value to 15 seconds using the ncli command: ncli pd update-metro-avail-pd name=<PD_NAME> break_replication_timeout=15. This ensures compliance with Nutanix best practices." Additional Context from ECA:
* Metro Availability Overview: Metro Availability synchronously replicates data between two clusters, typically within 100 km, to achieve zero RPO. Thebreak_replication_timeoutis a critical parameter that balances responsiveness to network issues with the need to avoid unnecessary replication breaks. The ECA course notes:"A timeout of 15 seconds is the default and recommended value to handle typical network fluctuations in Metro setups."
* NCC Warning Resolution: The NCC (Nutanix Cluster Check) monitors cluster health and flags configurations that deviate from best practices. The warning aboutbreak_replication_timeoutindicates a risk to Metro Availability stability, and setting it to 15 seconds aligns with Nutanix recommendations.
Supporting Reference from Web Results:
The Nutanix Support Portal (https://portal.nutanix.com) confirms the ECA guidance:"For Metro Availability, the break_replication_timeout should be set to a minimum of 15 seconds to prevent replication breaks due to transient network issues, as flagged by NCC warnings."

NEW QUESTION: 32
An administrator configured aremote site for Protection Domain replication, butnetwork performance and stabilityare impacted.
How can the remote site configuration be adjusted to fix the issue?
A. ConfigureNetwork Address Translation (NAT)between the two Nutanix clusters.
B. Configure theProtection Domain with many-to-many replication.
C. Configure aBandwidth Throttling Policy.
D. Configure theremote Cluster VIP as a proxy.
Answer: (SHOW ANSWER)
Network performance issues during replication can be mitigated using Bandwidth Throttling to control traffic spikes.
* Option C (Configure a Bandwidth Throttling Policy) is correct:
* Bandwidth Throttling ensures that replicationdoes not saturate the network, especially during peak usage hours.
* This is particularly useful forlow-bandwidth connectionsbetween remote sites.
* Option A (Configure NAT) is incorrect:
* NAT is not required forremote site replicationbetween Nutanix clusters.
* Option B (Many-to-Many Replication) is incorrect:
* This does not directly address network performance and mayincrease traffic load.
* Option D (Remote Cluster VIP as Proxy) is incorrect:
* VIP configurations help withload balancing but do not resolve bandwidth issues.
References:
* Nutanix Protection Policies Guide #Bandwidth Throttling for Remote Site Replication
* Nutanix KB #Optimizing Network Performance for Disaster Recovery

NEW QUESTION: 33
An administrator isprotecting an application and its datastored onVolume GroupsusingProtection Domains.
Duringfailover tests, allapplication VMs restore successfully, but theapplication data is completely missing.
How can the Protection Domain configuration be adjusted to avoid this issue in the future? (Choose two.)
A. Select the "Auto protect related entities" checkbox.
B. Manually add Volume Groups to Protected Entities.
C. Place Volume Groups in a separate Protection Domain.
D. Use application-consistent snapshots.
Answer: A,B (LEAVE A REPLY)
Protection Domains (PDs) in Nutanix ensure that entire applications and their associated data are protected during failover.However,Volume Groups (VGs) are not automatically includedunless explicitly configured.
* Option A (Select "Auto protect related entities") is correct:

* This setting ensures thatassociated Volume Groups, networks, and other dependencies are included in the Protection Domainautomatically.
* Without enabling this, only the VM itself would be protected, leading tomissing application data upon failover.
* Option B (Manually add Volume Groups to Protected Entities) is correct:
* If"Auto protect related entities" is not enabled, the administrator mustmanually add Volume Groups to the Protection Domain.
* This ensures that bothVMs and their attached Volume Groupsare replicated and recovered together.
* Option C (Place Volume Groups in a separate Protection Domain) is incorrect:
* Separating Volume Groups into a different PDdoes not guarantee they failover together with VMs.
* It is best practice tokeep related VMs and Volume Groups in the same PD.
* Option D (Use application-consistent snapshots) is incorrect:
* Whileapplication-consistent snapshots improve data integrity, theydo not fix missing Volume Groups in failover scenarios.
References:
Nutanix Disaster Recovery Guide#Protection Domain Configuration and Volume Groups Nutanix KB#Ensuring Volume Groups Are Included in Disaster Recovery Failovers

NEW QUESTION: 34
An administrator is configuring areplication scheduleonmultiple remote locationsdeployed using asingle- node cluster. The goal is to achieve thelowest possible RPO (Recovery Point Objective). How should the administrator configure the replication schedule?
A. Configure NearSync replication.
B. Configure a schedule for 16 minutes up to 59 minutes.
C. Configure Async replication.
D. Configure a schedule for 1 minute up to 15 minutes.
Answer: D (LEAVE A REPLY)
Nutanix NearSync replication provides the lowest RPO (as low as 1 minute) and is the best option for minimizing data loss in DR scenarios.
* Option D (Configure a schedule for 1 minute up to 15 minutes) is correct:
* NearSyncallows anRPO as low as 1 minute, providingnear-continuous data protection.
* This isideal for mission-critical applicationswhere minimal data loss is required.
* Option A (Configure NearSync) is incorrect:
* WhileNearSync is the best choice,just enabling it is not enough-the schedule must be set to1-15 minutes.
* Option B (16 to 59 minutes) is incorrect:
* NearSync operates within a 1-15 minute range. If set above 15 minutes,it defaults to Async replication.
* Option C (Async replication) is incorrect:

* Async replication typically has an RPO of 1 hour or more, which does not meet thelowest RPO requirement.
References:
* Nutanix Protection Policies Guide#NearSync vs. Async Replication
* Nutanix Bible#RPO and RTO in Disaster Recovery
* Nutanix KB#Configuring NearSync Replication for Single-Node Clusters

NEW QUESTION: 35
An administrator is trying to troubleshoot the environment after NCC raised an alert:
Detailed information for remote_site_connectivity_check: Node x.x.x.x:
WARN: Failed to connect to the remote site <remote_site>.
Which two steps should an administrator follow to provide a solution? (Choose two.)
A. Confirm that the remote cluster is reachable, and ports 2009 and 2020 are open between the clusters.
B. Configure Network Address Translation performed by any device in between the two Nutanix clusters.
C. If the remote site has been re-configured and the cluster has a new cluster incarnation ID, re-create the remote site.
D. Check if ping packets with an MTU of 9000 reach the destination cluster.
Answer: A,C (LEAVE A REPLY)
The NCC alert indicates connectivity failure to the remote site. Resolving this involves confirming network connectivity and re-establishing the remote site configuration if necessary.
From theNutanix Enterprise Cloud Administration (ECA)course materials:
"The primary ports used for replication between clusters are 2009 (for Prism Element API) and 2020 (for data replication). Ensuring these ports are open and reachable is critical for remote site connectivity." Also:
"If the remote site has been re-imaged or reconfigured, it may have a new cluster incarnation ID. In such cases, the remote site configuration must be recreated to align with the current cluster information." Steps like checking ping with MTU 9000 (D) are not directly related to remote site connectivity for replication, and NAT configurations (B) are generally not recommended unless explicitly required.

NEW QUESTION: 36
An administrator is managing a4-node clusterwithdifferent hardware generations:
* Two G5 Nodes# 2 CPUs (12 cores), 1 SSD (1.92 TB), 2 HDDs (4 TB).
* Two G7 Nodes# 2 CPUs (16 cores), 2 SSDs (1.92 TB), 4 HDDs (4 TB).
The cluster will bedecommissioned from productionand used forDisaster Recovery (DR) purposeswithan RPO of 1 hour.
What is the best approach when replacing G5 nodes without impacting performance?
A. New node must have at least 2 SSDs.
B. New node must be G7 or G8.

C. New node must have 2 CPUs with 12 cores.
D. New node must be hybrid.
Answer: A (LEAVE A REPLY)
For optimal Disaster Recovery performance, new nodes must match or exceed the storage performance of existing nodes.
* Option A (New node must have at least 2 SSDs) is correct:
* Since the G7 nodes havetwo SSDs, replacing G5 nodes withat least 2 SSDsensuresconsistent SSD cache and performance.
* Option B is incorrect:
* G7 or G8 nodes may help, but storage performance is more critical for DR.
* Option C is incorrect:
* CPU core count does not impact DR storage performance as much as SSD capacity.
* Option D is incorrect:
* Hybrid nodes are already in use, but SSDs must match for performance balance.
References:
* Nutanix Hardware Guide#Choosing Nodes for Hybrid and DR Clusters
* Nutanix KB#Balancing Storage Across Different Hardware Generations

NEW QUESTION: 37
An administrator is configuringNutanix Disaster Recovery (DR)for across-hypervisor setup (ESXi to AHV) but finds thatguest VMs do not recover properly at the DR location.
What is required for a successful cross-hypervisor DR event?
A. Utilize delta disks.
B. Deploy Legacy BIOS boot on hosts within the cluster.
C. Use raw device mappings.
D. Nutanix Guest Tools (NGT) must be installed on source guest VMs.
Answer: D (LEAVE A REPLY)
For cross-hypervisor DR failover (e.g., ESXi to AHV), Nutanix Guest Tools (NGT) must be installed on VMs to ensure proper configuration and recovery.
* Option D (NGT must be installed on source guest VMs) is correct:
* NGTensures correct reconfiguration of VM devices and networking settingsduring failover.
* Ithandles disk and driver reassignmentsbetween ESXi and AHV.
* Option A (Utilize delta disks) is incorrect:
* Delta disks are used insnapshot optimization, not DR failover.
* Option B (Deploy Legacy BIOS boot) is incorrect:
* AHV prefers UEFI boot mode, and Legacy BIOS is not a requirement.
* Option C (Use raw device mappings) is incorrect:
* RDMs areVMware-specificand arenot used in AHV failover scenarios.
References:
Nutanix Disaster Recovery Guide#Cross-Hypervisor Failover Best Practices Nutanix KB#Ensuring VM Compatibility During ESXi to AHV DR

NEW QUESTION: 38
An administrator has been tasked with performingfirmware upgradesfor all Nutanix sites.
When attempting to performfirmware upgrades via Life Cycle Manager (LCM)at aremote site with a single-node cluster, no firmware updates are listed as available. The administrator confirmed that the currently installed firmware isseveral revisions behind.
Why are no firmware upgrades listed in LCM for this cluster?
A. Single-node clusters only support one-disk firmware upgrades.
B. LCM is not supported on single-node clusters.
C. LCM cannot perform firmware upgrades on single-node clusters.
D. LCM does not have connectivity to the internet.
Answer: B (LEAVE A REPLY)
LCM (Life Cycle Manager) does not support automatic firmware upgrades for single-node clusters becausefirmware updates require cluster-wide operations, which are not possible with only one node.
* Option B (LCM is not supported on single-node clusters) is correct:
* Single-node clusterslack failover capability, making firmware upgradesunsafe without manual intervention.
* Option A (Single-node clusters only support one-disk firmware upgrades) is incorrect:
* This limitation does not apply toLCM as a whole.
* Option C (LCM cannot perform firmware upgrades) is incorrect:
* LCMcan perform manual firmware upgrades, butautomatic updates are not supported.
* Option D (LCM lacks internet connectivity) is incorrect:
* Even if the cluster is in adark site(no internet),LCM can use local update bundles.
References:
* Nutanix LCM Guide #Firmware Upgrade Considerations for Single-Node Clusters
* Nutanix KB #Why LCM Updates Are Not Available for Single-Node Deployments

NEW QUESTION: 39
An administrator needs to perform anLCM upgradeon anAHV host with GPUs.
What additional step is required before upgrading the host?
A. Create an agent VM on each host that has GPU drivers installed.
B. Run LCM in dark site mode so it can update AHV independently.
C. Use Direct Uploads to upload appropriate driver bundles.
D. Update NCC to the latest version and re-run Inventory.
Answer: C (LEAVE A REPLY)
Upgrading an AHV host with GPUs requires that the correct GPU drivers be manually uploaded to LCM, as GPU firmware is not updated automatically.
* Option C (Use Direct Uploads to upload appropriate driver bundles) is correct:
* LCM does notautomatically fetch GPU drivers.

* The administrator mustdownload and manually upload the appropriate firmware bundle before upgrading.
* Option A is incorrect:
* Agent VMs are not requiredfor GPU updates.
* Option B is incorrect:
* Running LCM indark site mode does not impact GPU firmware updates.
* Option D is incorrect:
* Updating NCCis a best practice but does not resolve GPU driver issues.
References:
* Nutanix LCM Guide#Manually Uploading GPU Firmware Bundles
* Nutanix KB#Updating AHV Hosts with GPUs

NEW QUESTION: 40
An administrator has successfully configuredMetro Availabilityfor aProtection Domain. However, after a few days, anNCC warningis raised:
"Following VMs are accessing data from remote clusters: VM-1 from remote cluster Remote-ML"
What is the first action an administrator must take to fix the issue?
A. Run the command:
ncli pd list metro-avail=true | egrep "Protection Domain Stretch Role" | grep "ACTIVE"
B. Usemust-affinity rulesto avoid automated VM migration to the standby datastore.
C. Migrate the VM to itsprimary siteand set appropriate rules forDRS and affinity.
D. Run the command:
ncc health_checks metro_availability_checks data_locality_check --cvm_list=X.X.X.20
Answer: (SHOW ANSWER)
Metro Availability requires that VMs always read data from their primary site to maintain optimal performance and prevent remote data access latency.
* Option C (Migrate the VM to its primary site and set appropriate rules) is correct:
* If a VMfails over to the secondary sitebut is still running in the primary site, it will read data remotely, causinghigh latency and performance issues.
* The solution is tomigrate the VM back to the primary siteand configureDRS rulesorhost affinity settingsto prevent unwanted movement.
* Option A is incorrect:
* The command listsactive Metro Availability protection domainsbut doesnot resolvethe issue.
* Option B is incorrect:
* Must-affinity rules can help, but they should be configuredafter migrating the VM back to the primary site.
* Option D is incorrect:
* RunningNCC health checkswill onlydiagnose the issue, not resolve it.
References:
* Nutanix Bible #Metro Availability and Data Locality
* Nutanix Best Practices #VM Affinity Rules for Metro Availability

* Nutanix KB #Troubleshooting Remote Data Access in Metro Availability

NEW QUESTION: 41
Theteam leads of a development environmentwant tolimit developer accessto a specific set ofVMs.
What is the most efficient way to enable the team leads to directly manage these VMs?
A. Create arole mappingfor each team lead and assign appropriately.
B. Create aVPCfor each team lead and give themVPC Admin.
C. Create aProjectfor each team lead and assign access.
D. CreateSecurity Policiesto isolate users.
Answer: C (LEAVE A REPLY)
Themost efficient wayto allow team leads to manage aspecific set of VMsis bycreating a Project (Option C)in Prism Central and assigning the team leads to that Project.
* Nutanix Projects allow administrators to control VM accessbased on groups and permissions, ensuring that users only manage VMs assigned to their project.
* Option A (Role Mapping)applies more broadly to roles but does not restrict access to specific VM groups.
* Option B (VPC Admin)is related to network segmentation, not VM access control.
* Option D (Security Policies)are used for network and firewall rules, not VM access control.
References:
* Nutanix Prism Central #Projects and Role-Based Access Control (RBAC)
* Nutanix Bible #Multi-Tenancy and Project-Based Access Control
* Nutanix KB #Setting Up Role-Based Access Control (RBAC) for Prism Central

NEW QUESTION: 42
Refer to Exhibit:

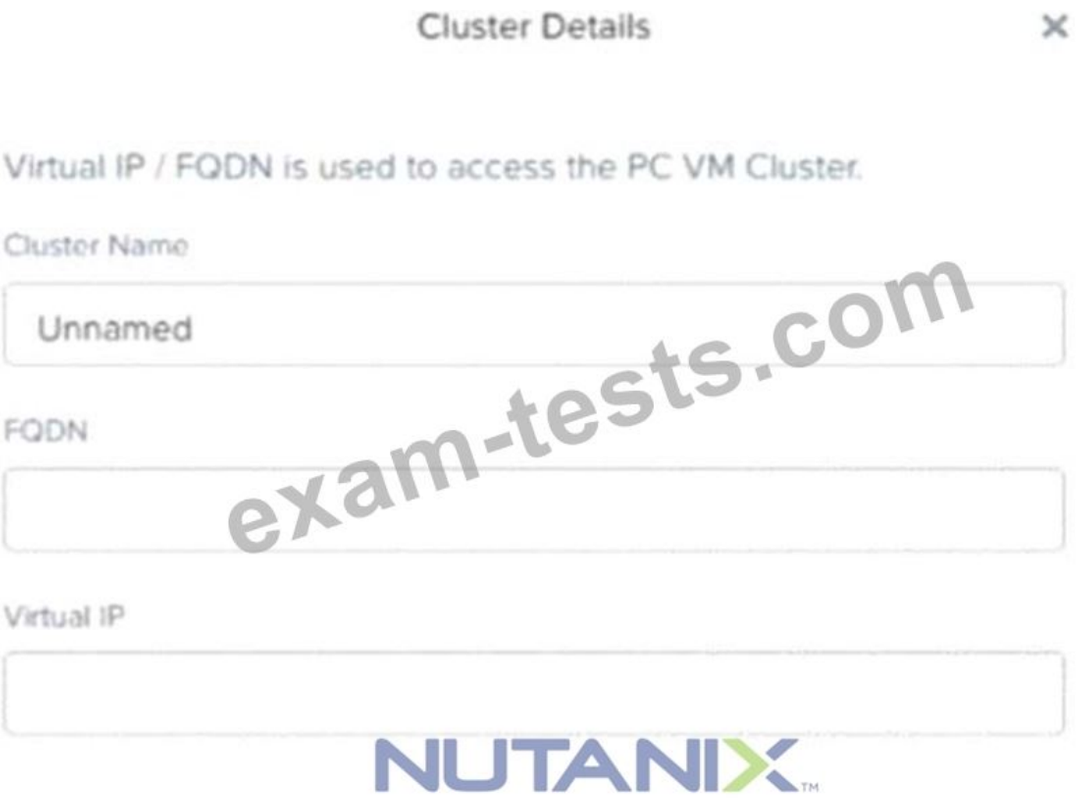

In ascale-out Prism Central deployment, what additional functionality does configuring anFQDN instead of a Virtual IPprovide?

A. Load balancing

B. Resiliency

C. Segmentation

D. SSL Certificate

Answer: A (LEAVE A REPLY)

When usingFQDN instead of a Virtual IP in a scale-out Prism Central deployment, Nutanixenables load balancing across multiple Prism Central instances.

* Option A (Load balancing) is correctbecause it ensures that requests are distributedamong multiple Prism Central nodes, improving performance and redundancy.
* Option B (Resiliency) is incorrectbecause resiliency is achieved through HA and replication, not through FQDN configuration.
* Option C (Segmentation) is incorrectbecause network segmentation is handled at the VLAN or security policy level.
* Option D (SSL Certificate) is incorrectbecause SSL certificates can be applied regardless of whether FQDN or Virtual IP is used.

References:

* Nutanix Prism Central Deployment Guide
* Nutanix Best Practices for Scale-Out Prism Central
* Nutanix Support KB: Configuring FQDN for Prism Central

NEW QUESTION: 43

A cluster has an RF3 storage container with many VMs. Before the cluster runs out of space, the administrator has created a new RF2 storage container and would like to live migrate the vDisks.
Which check should be done before performing vDisk migration?
A. Remove VMs from Protection Domains or Protection Policies.
B. Ensure storage optimization options match between storage containers
C. Validate network latency is below 5 milliseconds.
D. Verify Multichannel Support is enabled.
Answer: (SHOW ANSWER)

NEW QUESTION: 44
An administrator wants toenable application discoveryon aNutanix clusterto monitor applications. APrism Central instanceis alreadyconfigured with sufficient CPU and memory.
What other prerequisites must be met before enabling application discovery? (Choose two.)
A. Network controller is enabled
B. Sufficient Prism Central VM resources
C. API key and key ID
D. Internet connection
Answer: B,D (LEAVE A REPLY)

NEW QUESTION: 45
After addingnew workloads, theOverall Runwayis nowbelow 365 days, but thescenario still shows the cluster is in good shape.
Why?
A. Because Storage Runway is still good.
B. Because new workloads are sustainable.
C. Because there are recommended resources.
D. Because the Target is 1 month.
Answer: B (LEAVE A REPLY)
A cluster runway below 365 days does not necessarily indicate an issue if Intelligent Operations determines that workloads are sustainable.
* Option B (Because new workloads are sustainable) is correct:
* Nutanixanalyzes resource trendsand marks clusters as healthyif new workloads are within projected capacity.
* Option A (Storage Runway is still good) is incorrect:
* Storage isone component, butCPU and memoryalso affect runway calculations.
* Option C (Recommended resources) is incorrect:
* Recommendations help optimize capacity, butdo not define cluster health.
* Option D (Target is 1 month) is incorrect:
* The scenario's target window does not impact the actual runway calculation.
References:
* Nutanix Prism Central Guide #Capacity Planning & Runway Analysis

* Nutanix KB #Understanding Capacity Runway and Workload Sustainability

NEW QUESTION: 46
An administrator has deployed two Nutanix clusters and is now establishing synchronous replication between them. However, the replication is failing immediately.
Which two responses show the reason and corrective action an administrator can take to resolve the issue?
(Choose two.)
A. If the primary and the recovery clusters are in different subnets, open the ports manually for communication.
B. If the primary and the recovery clusters are on the same subnet, open the ports manually for communication.
C. Use the command modify_firewall to open the ports on eth1 interface.
D. Use the command modify_firewall to open the ports on eth0 interface
Answer: (SHOW ANSWER)
When synchronous replication fails immediately between two clusters, it is often due to blocked communication across the required ports (2009 and 2020). These must be open manually if the clusters are in different subnets or if network policies block traffic.
From theNutanix Enterprise Cloud Administration (ECA)course materials:
"Replication communication relies on specific ports, which must be allowed through the firewall. If the clusters are in different subnets or if there are external firewalls, these ports must be explicitly opened."
"The modify_firewall command is used to open or close ports on cluster nodes. For replication and remote site communication, eth1 is typically used for external replication traffic." Since eth1 is typically used for external connectivity and replication traffic, opening ports on this interface using modify_firewall resolves the communication block.

NEW QUESTION: 47
An administrator is trying to configure Metro Availability between Nutanix ESXi-based clusters. However, theCompatible Remote Sitesscreen does not list all required storage containers.
Which two reasons could be a cause for this issue? (Choose two.)
A. Source and destination hardware are from different vendors.

B. The remote site storage container has compression enabled.
C. Both storage containers must have the same name.
D. The destination storage container is not empty.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 48
After upgradingPrism Central from PC2022.1 to PC2024.1, an administrator isunable to log inwith theirIAM domain account.
What is the first troubleshooting step the administrator should take?
A. Ping the Domain Controller from the CVM.
B. Ensure port 9441 is open in the firewall.
C. Validate the trusted signing certificate of the organization.
D. Log in with a local admin account.
Answer: (SHOW ANSWER)
After a Prism Central upgrade, IAM authentication settings may require reconfiguration.
* Option D (Log in with a local admin account) is correct:
* If IAM authentication fails, thelocal admin account must be usedto check domain settings.
* Option A (Ping the Domain Controller) is incorrect:
* Network connectivity is important, but theissue is likely related to IAM settings, not network reachability.
* Option B (Check firewall port 9441) is incorrect:
* Port 9441 is used forSSO authentication, butport issues usually result in login delays, not complete failures.
* Option C (Validate signing certificate) is incorrect:
* While certificates can cause issues,local admin login should always work.
References:
Nutanix KB #Troubleshooting IAM Login Issues After a Prism Central Upgrade Nutanix Documentation #Managing User Authentication and IAM Integration

NEW QUESTION: 49
An administrator is planning for an upcoming maintenance window. The administrator would like to minimize the chance of an upgrade failure during the maintenance window to ensure the updates will complete without issue.
What action should the administrator take to reduce the risk of any potential failures during an upgrade?
A. Reboot each CVM one-at-a-time to ensure the reboots are successful.
B. Upgrade AOS to the latest version from LCM.
C. Run an Upgrade Precheck from LCM.
D. Reboot each host one-at-a-time to ensure the reboots are successful.
Answer: (SHOW ANSWER)

NEW QUESTION: 50
In anRF2 cluster, what is theminimum number of nodes requiredto allow ahost removal?
A. 2
B. 3
C. 4
D. 5
Answer: B (LEAVE A REPLY)
Replication Factor (RF2) means that each piece of data is stored twice across different nodes to ensure availability.
* Option B (3 nodes) is correct:
* In anRF2 cluster, data redundancy requires at leastthree nodestoensure data protectionwhen one node is removed.
* If a node is removed from a3-node cluster, Nutanix automatically redistributes data across the remaining nodes.
* Option A (2 nodes) is incorrect:
* RF2 requiresat least three nodesto maintainfault tolerance.
* A2-node cluster cannot provide full redundancywithout a Witness node.
* Option C (4 nodes) and Option D (5 nodes) are incorrect:
* While larger clusters provide more redundancy, theminimum requirement is 3 nodes.
References:
* Nutanix Bible #Replication Factor (RF) and Fault Tolerance
* Nutanix Prism Element Guide #Managing Node Failures and Removals

NEW QUESTION: 51
An administrator is configuringProtection Policiesto replicate VMs to aNutanix Cloud Cluster (NC2)over the internet.
To comply with security policies, how should data be protected during transmission?
A. Configure Data on a self-encrypting drive.
B. Configure VMs to use UEFI Secure Boot.
C. Enable Data-at-Rest Encryption.
D. Enable Data-in-Transit Encryption.
Answer: D (LEAVE A REPLY)
Data-in-Transit Encryption ensures that replication traffic is encrypted while being sent over the internet.
* Option D (Enable Data-in-Transit Encryption) is correct:
* This encryptsreplicated data between clusters, ensuring security againstman-in-the-middle attacks.
* Option A (Self-encrypting drive) is incorrect:
* This protects data at rest, not during transmission.
* Option B (UEFI Secure Boot) is incorrect:
* Secure Boot prevents unauthorized OS modifications, but doesnot encrypt network traffic.

* Option C (Data-at-Rest Encryption) is incorrect:
* This encryptsstored data but does not secure replication traffic.
References:
* Nutanix Security Guide #Configuring Data-in-Transit Encryption
* Nutanix KB #Protecting Replication Traffic Over Public Networks

NEW QUESTION: 52
An administrator has configuredAHV Metro Availability with Witnessand is testing failover scenarios.
During testing, the administratordisconnects the primary and recovery clustersbutPrism Central remains connected to the recovery site.
What are two expected system behaviors? (Choose two.)
A. Guest VM I/O operations pause (freeze) until connectivity is restored.
B. Guest VM I/O operations pause (freeze) until connectivity between Prism Central and the primary site is restored.
C. Guest VMs failover automatically to the recovery cluster.
D. Guest VMs continue to run on the primary cluster.
Answer: A,C (LEAVE A REPLY)
When connectivity between Metro clusters is lost, Nutanix Metro Availability ensures data integrity using Witness for automatic failover.
* Option A (Guest VM I/O operations pause until connectivity is restored) is correct:
* Metro Availability enforces data consistency, so I/O operationspauseuntil failover is confirmed.
* Option C (Guest VMs failover automatically to the recovery cluster) is correct:
* TheWitness VM detects the failureand initiates anautomatic failoverto the secondary cluster.
* Option B is incorrect:
* Prism Centraldoes not control VM failover in Metro Availability.
* Option D is incorrect:
* The primary cluster isunreachable, so VMs cannot continue running there.
References:
Nutanix Metro Availability Guide#How Witness Handles Failover Scenarios Nutanix KB#I/O Freezing and Failover Behavior in Metro Clusters

NEW QUESTION: 53
ADisaster Recovery administratorhas set up aProtection Policy for 50 workloads, all configured similarly.
TheRPO is 60 minuteswith aspecified retention of 10 local copies, 5 remote copies, and crash consistency.
After activation,recovery points are not appearing at the DR site, even though they arevisible on the production side.
What is the most likely issue?
A. Nutanix Guest Tools (NGT) is not installed on the source VMs.

B. Windows updates need to be applied to all affected VMs.
C. The storage container name on the DR cluster does not match the production cluster.
D. The storage container RF factor does not match in both clusters.
Answer: C (LEAVE A REPLY)
For Disaster Recovery to function correctly, the source and destination storage containersmust have identical names.
* Option C (Storage container name mismatch) is correct:
* If thestorage container name at the DR site does not match, Nutanix cannot mapsnapshots and replication data.
* This causesfailover operations to fail, even though data exists.
* Option A (NGT not installed) is incorrect:
* NGT is needed for application-consistent snapshots, butnot required for crash-consistent snapshots.
* Option B (Windows updates) is incorrect:
* OS updates do not affect replication availability.
* Option D (Storage RF factor mismatch) is incorrect:
* Replication works across different RF factors, but performance may vary.
References:
* Nutanix Disaster Recovery Guide#Requirements for Remote Replication
* Nutanix KB#Storage Container Mapping for Protection Domains

NEW QUESTION: 54
An administrator wants toensure that user VMs on AHV hostscan take advantage ofbandwidth beyond a single adapter in a bond.
Which uplink Bond Type should the administrator configure to accomplish this?
A. No Uplink Bond
B. Active-Active
C. Active-Active with MAC pinning
D. Active-Backup
Answer: (SHOW ANSWER)
Active-Active bonding allows multiple network interfaces to be used simultaneously, improving bandwidth and redundancy.
* Option B (Active-Active) is correct:
* This mode enablesload balancing across all available adapters, providinghigher throughput and fault tolerance.
* Option A (No Uplink Bond) is incorrect:
* Without a bond, VMscannot benefit from multiple adapters.
* Option C (Active-Active with MAC pinning) is incorrect:
* MAC pinningbinds traffic to a single NIC, limiting bandwidth distribution.
* Option D (Active-Backup) is incorrect:
* This modeonly provides failover, not increased bandwidth.

References:
* Nutanix AHV Networking Guide #Bonding Modes and Load Balancing
* Nutanix KB #Optimizing Network Throughput in AHV

NEW QUESTION: 55
An administrator noticeshigh CPU usage on a VMand wants to determine whether addingmore vCPUs would improve performance.
Which two metrics should be analyzed to make this decision? (Choose two.)
A. VM CPU Ready Time
B. VM CPU Usage
C. Host CPU Usage
D. Host Memory Swap Out Rate
Answer: (SHOW ANSWER)
When diagnosing CPU performance issues, CPU Ready Time and CPU Usage are the key indicators of whether more vCPUs are needed.
* Option A (VM CPU Ready Time) is correct:
* High CPU Ready Time means the VM is waiting for CPU resources, indicatingCPU contention.
* Option B (VM CPU Usage) is correct:
* If CPU usage isconsistently high, adding more vCPUsmay improve performance.
* Option C (Host CPU Usage) is incorrect:
* Host-wide CPU usage does not indicate whether a specific VM needs more vCPUs.
* Option D (Host Memory Swap Out Rate) is incorrect:
* Memory swappingaffects RAM performance, not CPU allocation.
References:
* Nutanix Prism Central Guide#Analyzing VM CPU Performance
* Nutanix KB#Understanding CPU Ready Time and VM Performance