Nutanix.NCA-7.5.v2026-10-02.q42

Exam Code:	NCA-7.5
Exam Name:	Nutanix Certified Associate (NCA) 7.5 Exam
Certification Provider:	Nutanix
Free Question Number:	42
Version:	v2026-10-02
# of views:	103
# of Questions views:	420
https://www.freecram.net/torrent/Nutanix.NCA-7.5.v2026-10-02.q42.html	

NEW QUESTION: 1
In a Nutanix cluster where compute nodes run ESXi, which hypervisor is required on storage-only nodes added to the cluster?
A. Hyper-V
B. AHV
C. ESXi
D. KVM
Answer: (SHOW ANSWER)
A Nutanix storage-only node operates with AHV even when the hyperconverged compute nodes in the same cluster use VMware ESXi. The storage-only node contributes storage capacity and runs the Controller VM and required Nutanix services, but it does not host ordinary user VMs. Using AHV on the storage-only node preserves compatibility with the Nutanix storage-only architecture and avoids requiring an additional third- party hypervisor licence solely for a node that supplies storage. This creates a supported mixed-hypervisor cluster configuration: ESXi remains on the compute-capable HCI nodes, while AHV runs on the storage-only nodes. Administrators must observe the documented restrictions for cluster expansion, network segmentation, maintenance, and workload placement in this configuration.
User workloads remain on the ESXi compute nodes because the AHV storage-only nodes are not available as general VM execution targets.
Hyper-V is not used for this ESXi storage-only-node design. Although AHV is based on KVM technology, the supported Nutanix hypervisor product and configuration is specifically AHV, making generic KVM an incorrect selection. Reference: Prism 7.5 - Storage-Only Nodes . NHCF alignment: compute-only and storage- only node configurations.

NEW QUESTION: 2
An administrator needs to configure cluster fault tolerance of 1N & 1D.
How many nodes must the cluster have for this task?

A. 2
B. 3
C. 4
D. 5
Answer: (SHOW ANSWER)
The 1N & 1D cluster fault-tolerance level is supported on a three-node cluster. It is designed to maintain access to data when one complete node and one additional disk fail concurrently, provided that the failures and remaining resources satisfy the documented resiliency conditions. To achieve this protection level, the platform configures replication factor 3 for the relevant data and maintains three metadata copies.
This configuration is distinct from standard 1N/1D fault tolerance, which protects against either one node failure or one disk failure and normally uses data replication factor 2. The ampersand in 1N & 1D is technically significant because it indicates simultaneous node-and-disk protection.
A two-node cluster uses a separate witness-based architecture and cannot provide this three-node 1N & 1D configuration. Nutanix documentation identifies three as both the required and supported node count for this specific fault-tolerance mode. Larger conventional clusters use their applicable fault-tolerance configurations rather than treating 1N & 1D as an arbitrary scalable setting.
Administrators must also verify that the cluster has sufficient resilient capacity before enabling the higher protection level. References: Prism 7.5 - One Node and One Disk Cluster Fault Tolerance and Cluster Fault Tolerance .

NEW QUESTION: 3
Which component is used to present Nutanix storage as iSCSI block targets to external (non-virtualized) servers?
A. vDisk Template
B. Volume Group
C. Storage Container
D. Storage Pool
Answer: (SHOW ANSWER)
A Volume Group is the Nutanix storage component used to present block storage through iSCSI to external clients, including physical or otherwise non-virtualized servers. A Volume Group contains one or more virtual disks and can be attached directly to an AHV VM or exposed through iSCSI to qualified external initiators.
External client access is controlled by defining the relevant initiator information and establishing the iSCSI connection to the Volume Group.
A storage container is a logical storage boundary used for VM data placement and for settings such as replication factor, compression, deduplication, and erasure coding. Although Volume Group data ultimately consumes capacity from a container, the container itself is not presented as an iSCSI block target.

A storage pool aggregates the physical storage resources contributed by cluster nodes. It operates below the storage-container and workload-consumption layers and is not directly assigned to an external server. A vDisk template is also not the external block-presentation object.
The examination distinction is between the underlying storage hierarchy and the consumer-facing block- storage object. Storage pools supply capacity, containers provide logical policy boundaries, and Volume Groups organize vDisks for direct or iSCSI-based attachment.
NHCF/NCA reference: Nutanix storage operations, Volume Groups, external clients, and iSCSI connectivity.
Official references: Storage Components and Volume Group Details .

NEW QUESTION: 4
An administrator is provisioning a new Linux virtual machine.
Which timezone is recommended for this VM?
A. Local
B. Cluster
C. UTC
D. RTC
Answer: (SHOW ANSWER)
UTC is the recommended timezone selection when creating a Linux VM on AHV. Nutanix specifically states that the real-time clock of a Linux VM must operate in UTC. Selecting the UTC timezone ensures that the guest's hardware-clock interpretation is consistent with Linux conventions and avoids time shifts caused by local timezone or daylight-saving changes.
Time consistency is essential for authentication, application logging, distributed databases, monitoring, certificate validation and event correlation. The guest operating system can still display a regional timezone to users at the application or operating-system level while maintaining its underlying hardware clock in UTC.
"Local" can cause the guest to interpret the virtual hardware clock according to a regional offset, producing incorrect timestamps after migration, reboot or daylight-saving transitions. "Cluster" is not a supported timezone designation for this VM setting. RTC means real-time clock and describes the clock mechanism rather than a selectable geographic timezone.
The VM should also use a reliable NTP source after deployment so its operating-system time remains synchronized.
NHCF/NCA reference: VM Management-creating and configuring Linux virtual machines; NCA Section 2,
"Describe VM tasks." Official references: Prism 7.5-Creating an AHV VM , Prism 7.5-Time-Synchronization Recommendations .

NEW QUESTION: 5
Which Nutanix solution provides microsegmentation and application-centric security policies for VMs?
A. AHV Virtualization

B. AOS Storage
C. Flow Network Security
D. Nutanix Disaster Recovery
Answer: (SHOW ANSWER)
Flow Network Security provides application-centric microsegmentation and distributed security-policy enforcement for Nutanix workloads. It uses categories and application context to define which workloads may communicate, rather than depending exclusively on static IP addresses or physical VLAN boundaries.
Policies can follow VMs as they move or scale within supported Nutanix environments. Microsegmentation limits east-west communication between application components and reduces the potential lateral movement of malware or unauthorized users. Administrators can define policies for application tiers, environments, user groups and isolation requirements, then manage those policies centrally through Prism Central. Flow Network Security also provides traffic visualization and policy-monitoring capabilities before enforcement.
AHV supplies the virtualization and host-compute layer but is not, by itself, the security-policy product identified by the question. AOS provides distributed storage, resilience and data services. Nutanix Disaster Recovery manages workload protection, replication and recovery rather than distributed microsegmentation.
NHCF/NCA reference: Nutanix Cloud Infrastructure solutions, networking and security capabilities; NCA Section 1, "Recognize Nutanix solutions, use cases and features of NCI technologies." Official references:
Flow Network Security , Nutanix Flow .

NEW QUESTION: 6
What Nutanix storage concept consists of groups of physical disks from one or more disk tiers?
A. Disk Groups
B. Volume Groups
C. Storage Pools
D. Storage Containers
Answer: (SHOW ANSWER)
A Nutanix storage pool is the physical aggregation of disks contributed by the nodes in a cluster. The disks may come from one or more storage tiers, such as NVMe, SSD or HDD, depending on the supported hardware configuration. Each physical storage device belongs to a storage pool, and Nutanix generally recommends a single storage pool per cluster so AOS can optimize capacity and performance across the available devices.
A storage container is a logical subdivision of storage capacity presented from the storage pool. Administrators apply settings such as compression, deduplication and other storage policies at the container level. Therefore, a container consumes capacity from a pool but is not itself the grouping of physical disks.
A volume group is a collection of virtual disks exposed to one or more VMs or external clients as block storage. "Disk group" is not the Nutanix AOS abstraction defined by the question.

The hierarchy is consequently physical disks, then the storage pool, followed by logical storage containers and workload-facing virtual disks or volume groups.
NHCF/NCA reference: Storage Concepts-AOS storage architecture, storage pools, containers and storage tiers; NCA Section 2, "Discuss storage operations." Official references: Prism 7.5-Storage Components , AOS 7.5-Nutanix Software Configuration .

NEW QUESTION: 7
On which entity is Replication Factor configured?
A. Replication Policy
B. Storage Container
C. Protection Domain
D. Distributed Storage Pool
Answer: (SHOW ANSWER)

NEW QUESTION: 8
Refer to the exhibit.



A VM is experiencing poor performance at the storage level.
After reviewing the image, which action should the administrator take?
A. Verify whether a storage policy is affecting the VM.
B. Migrate the VM to another host.
C. Increase the CVM memory size.
D. Add additional SSDs to all nodes in the cluster.
Answer: (SHOW ANSWER)
The exhibit displays the Cluster-wide Controller Latency metric, which represents the average time required for I/O operations to be processed by the cluster's storage controllers. Although the chart contains a latency spike, it does not establish that the cluster lacks SSD capacity or that a Controller VM is undersized. The first diagnostic action should be to determine whether a storage policy, particularly a storage Quality of Service configuration, is limiting the VM.
Storage QoS can impose an IOPS limit on a VM. When the workload reaches that limit, I/O requests are delayed to enforce the configured policy, producing elevated latency even when the physical storage subsystem has sufficient capacity. Nutanix documentation specifically notes that a maximum-IOPS configuration can cause higher latency to appear for the affected VM.
Migrating the VM to another host is unlikely to resolve a policy-enforced storage limit because Nutanix storage is distributed across the cluster. Increasing CVM memory or purchasing additional

SSDs would be premature without evidence of sustained resource exhaustion. The appropriate first step is therefore to inspect the VM's applicable storage policy and QoS settings. References: Prism 7.5 - Storage Overview and Prism Central 7.5 - Storage QoS .

NEW QUESTION: 9
Which bond type provides the best utilization of multiple physical NICs for VM traffic throughput?
A. Active-Backup
B. No Uplink Bond
C. Balance-slb
D. Balance-tcp
Answer: (SHOW ANSWER)
Balance-tcp provides active-active link utilization through Link Aggregation Control Protocol and is the strongest choice when the objective is to maximize aggregate VM traffic throughput across multiple physical NICs. The mode distributes traffic using flow information, allowing separate TCP flows-including flows originating from the same VM-to use different physical uplinks. This can increase host and VM bandwidth utilization beyond the capacity of a single uplink when sufficient parallel traffic flows exist.
Balance-tcp requires the corresponding physical-switch ports to be configured as a compatible LACP link aggregation group. The AHV bond and upstream switch configuration must agree; otherwise, connectivity or traffic-distribution failures can occur.
Active-Backup uses only one uplink for forwarding while the other remains available for failover, providing redundancy but not simultaneous bandwidth utilization. Balance-slb is also an active-active mode and does not require switch-side LACP, but its source-load-balancing behavior generally associates traffic sources with particular uplinks and cannot distribute an individual flow across links. "No Uplink Bond" supplies neither link aggregation nor bond-level redundancy.
NHCF/NCA reference: Networking-AHV virtual switches, uplink bonds and physical-network connectivity; NCA Section 2, "Identify virtual networking tasks." Official references: Prism Central-Creating a Virtual Switch , AHV Networking Recommendations .

NEW QUESTION: 10
A VM is experiencing delays even though its own CPU usage appears low.
Which metric should be checked to identify scheduling issues?
Freecram.net
A. CPU Ready
B. Disk capacity
C. Memory usage
D. Snapshot size
Answer: (SHOW ANSWER)
CPU Ready measures the time during which a virtual CPU is prepared to execute but is waiting for the hypervisor to schedule it on a physical processor. A VM can therefore report low apparent CPU consumption while still responding slowly: its workload is not continuously executing because the assigned vCPUs are waiting for physical CPU time.

Elevated CPU Ready commonly indicates physical CPU contention, excessive vCPU allocation, an unfavorable vCPU-to-core ratio, or competition from other VMs on the host. The administrator should correlate CPU Ready with host CPU utilization, VM sizing, the number of runnable vCPUs, and workload behavior. Assigning additional vCPUs without confirming the cause can worsen scheduling delays because the hypervisor may need to coordinate more virtual processors.
Disk capacity measures available storage space and does not identify processor scheduling latency. Memory usage can reveal pressure, swapping, or insufficient guest memory, but it does not directly represent time spent waiting for CPU scheduling. Snapshot size relates to retained point-in-time storage data and is unrelated to hypervisor CPU dispatch. Nutanix identifies CPU Ready Time as the percentage of CPU wait time, making it the appropriate metric. Reference: Nutanix Intelligent Operations - VM Details and CPU Ready Time .
NHCF alignment: workload performance interpretation.

NEW QUESTION: 11
When using the LCM Dark Site procedure in Prism Element or Prism Central, which Nutanix software component should be checked and upgraded first?
A. MSP
B. NCC
C. LCM Framework
D. Foundation VM
Answer: (SHOW ANSWER)
The LCM Framework must be current before the remaining dark-site software or firmware upgrade workflow is performed. The framework supplies the orchestration logic, compatibility metadata processing, inventory functionality, and update mechanisms that Life Cycle Manager uses to evaluate and execute component upgrades. An outdated framework may not correctly recognize newer bundles, dependencies, compatibility rules, or supported upgrade paths.
For the recommended dark-site sequence, Nutanix directs administrators to perform an LCM inventory first.
That operation checks the installed environment and updates the LCM Framework when the required framework bundle is available. Nutanix further specifies that no other software component should be upgraded during this initial framework step. NCC is upgraded and run afterward so that current health checks can be executed before broader maintenance proceeds.
Foundation VM is principally associated with imaging, discovering, and deploying Nutanix nodes; it is not the first component in an established cluster's LCM upgrade order. MSP is likewise not the controlling lifecycle framework for this procedure. Therefore, the correct operational sequence begins with the LCM Framework. Reference: LCM - Recommended Upgrade Order for the Dark Site Method . NHCF alignment:
software maintenance, LCM inventory, and lifecycle sequencing.

NEW QUESTION: 12
How can an administrator protect volume-based containerized workloads?

A. Configure a replication policy
B. Enable Disaster Recovery in Prism Central
C. Use Nutanix Data Services for Kubernetes
D. Set up Objects replication
Answer: (SHOW ANSWER)
Nutanix Data Services for Kubernetes (NDK) provides application-centric protection and data-management capabilities for stateful Kubernetes workloads. It is designed to protect persistent application data together with relevant Kubernetes configuration, allowing administrators to apply policies for snapshots, replication, recovery, migration, and disaster recovery. This includes persistent block volumes backed by Nutanix Volumes and, in supported releases, file-based persistent storage.
A generic replication policy does not by itself establish Kubernetes-aware protection. Effective recovery must account for both persistent data and application objects such as namespaces, StatefulSets, services, ConfigMaps, and secrets. NDK coordinates these elements and supplies policy-driven recovery workflows at the application level. Prism Central disaster-recovery capabilities primarily protect supported VM workloads and do not automatically provide the Kubernetes-native application awareness required here. Nutanix Objects replication protects object-store data and buckets rather than persistent block volumes consumed by containerized applications.
Nutanix describes NDK as extending enterprise data services to containerized applications and providing built- in disaster recovery with policy-driven data management. This directly matches the stated requirement.
Reference: Nutanix Data Services for Kubernetes . NHCF alignment: recognizing Nutanix cloud-native products and matching NDK to stateful Kubernetes data protection.

NEW QUESTION: 13
Refer to the exhibit.



freecram.net

A cluster is showing the alert in the attached image.
What is the recommended method to quickly determine possible causes and corrective actions for this alert?
A. Copy the text of the alert and use it for a Google search.
B. Click the Source Entity of the alert to access the alert details page.
C. Click the selection box next to the alert and then the blue Email Configuration button.

D. Click the Title of the alert to access the alert details page.

Answer: (SHOW ANSWER)

The alert title is the navigation control used to open the alert details page. That page consolidates the information needed to investigate the condition, including the alert description, severity, creation time, affected entity, possible causes, and recommended corrective actions. Nutanix documentation states that an administrator displays an alert's details by clicking its title.

The Source Entity link leads to information about the affected object-for example, a disk, controller VM, host, or cluster. Entity information may be useful later in the investigation, but it is not the direct route to the alert-specific cause and resolution guidance requested in the question. The selection box merely selects the record for supported bulk or administrative operations. Email Configuration controls alert-notification delivery and does not diagnose the active alert.

Copying the title into a public search engine is also inferior to using the contextual information already maintained in Prism. Search results may apply to different software releases or environmental conditions, whereas the details page is connected to the actual alert instance and entity. Therefore, clicking the alert title is the fastest supported diagnostic path. References: Prism Central - Alerts Summary View and Prism 7.5 - Alert Details . NHCF alignment: interpreting and investigating alerts.

NEW QUESTION: 14

Which tool tests how a Nutanix environment will behave under a simulated production load?

A. Nutanix Move

B. Nutanix Collector

C. Nutanix Foundation

D. Nutanix X-Ray

Answer: (SHOW ANSWER)

Nutanix X-Ray is an automated testing and benchmarking framework that runs repeatable scenarios against hyperconverged infrastructure targets. It uses realistic workload generators and controlled infrastructure events to evaluate performance, resilience, scalability, data integrity, and availability. Its scenarios can simulate conditions such as VM cloning, snapshots, migrations, node failures, node power-offs, and sustained storage workloads.

X-Ray is intended for evaluation and test environments rather than active production clusters. The resulting metrics and reports allow architects and administrators to compare configurations and determine how an environment behaves under representative production-style demand and failure conditions.

Nutanix Move migrates workloads between supported platforms and does not benchmark infrastructure behavior. Nutanix Collector gathers environmental information for assessment and sizing activities but does not generate simulated application loads. Nutanix Foundation images nodes and deploys or configures clusters; it is not a performance-testing framework.

The requirement to test behavior under simulated load therefore maps directly to X-Ray's scenario-based testing model. References: Nutanix X-Ray FAQ and X-Ray 4.4 - Test Scenarios . NHCF alignment:

identifying Nutanix assessment, deployment, migration, and testing tools.

NEW QUESTION: 15
After deploying a VM, which tool provides additional capabilities, such as file-level restore?
A. Nutanix X-Ray
B. Nutanix Foundation
C. Nutanix Guest Tools
D. Nutanix Move
Answer: C (LEAVE A REPLY)
Nutanix Guest Tools is a collection of guest-resident services and utilities that extends the integration between a deployed virtual machine and the Nutanix platform. One of its principal capabilities is Self-Service Restore, also described as file-level restore. When the feature is enabled and NGT is installed, an authorized VM administrator can access files from a Nutanix-protected snapshot without requiring the infrastructure administrator to restore the entire VM.
NGT establishes the guest-side functionality required to expose and use supported recovery services.
Depending on the operating system and enabled components, NGT also supports functions such as application-consistent snapshot coordination and guest communication with the Nutanix platform.
The other tools serve different lifecycle stages. Nutanix X-Ray validates platform performance, resilience and infrastructure behavior. Foundation images physical nodes and creates or expands Nutanix clusters. Nutanix Move migrates workloads from supported source environments to Nutanix AHV or other supported targets.
None of these tools supplies guest-level file recovery after normal VM deployment.
NHCF/NCA reference: VM Management and Nutanix tools; VM protection, recovery points and guest integration; NCA Section 1, "Recognize the use cases and features of Nutanix tools." Official references:
Prism Central 7.5-Nutanix Guest Tools Overview , Prism 7.5-Self-Service Restore .

NEW QUESTION: 16
Which minimum licensing option includes Container Storage Interface, Load Balancing, and Role-Based Access Control?
A. NCI Ultimate
B. NKP Pro
C. NCI Pro
D. NKP Ultimate
Answer: (SHOW ANSWER)
NCI Pro is the minimum listed licensing option that includes the specified capabilities. Official Nutanix licensing documentation states that NKP Starter is included with NCI Pro and NCI Ultimate. NKP Starter provides the Nutanix CSI driver for persistent storage, role-based access control with single sign-on, and load- balancing and ingress services.

NKP Pro also contains these functions because it includes everything in Starter. However, it adds production- oriented services such as an observability stack, backup and restore, service mesh, serverless functions, broader deployment options, and additional platform capabilities. It is therefore not the minimum option among the answers.
NCI Ultimate includes the same NKP Starter entitlement but is a higher NCI tier than NCI Pro. NKP Ultimate is the highest NKP tier and adds fleet management, multitenancy, centralized observability, GitOps, cost management, application-catalog functions, and support for managing broader hybrid multicloud Kubernetes environments.
The decisive requirement is not simply which license contains the features; it is which listed option provides them at the lowest applicable tier. Since NCI Pro includes NKP Starter and Starter supplies CSI, load balancing, and RBAC, option C is technically correct.
NHCF/NCA reference: NCI and NKP editions, licensing, and Kubernetes platform capabilities. Official references: NKP Datasheet and Edition Matrix and Nutanix Cloud Platform Software Options .

Valid NCA-7.5 Dumps shared by EduDump.com for Helping Passing NCA-7.5 Exam! EduDump.com now offer the **newest NCA-7.5 exam dumps**, the EduDump.com NCA-7.5 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com NCA-7.5 dumps with Test Engine here: https://www.edudump.com/exams/Nutanix/NCA-7.5/premium/ (**91** Q&As Dumps, **35%OFF** **Special Discount Code: freecram**)

NEW QUESTION: 17
Which number of node(s) can be placed into maintenance mode at the same time in a cluster?
A. Two nodes
B. All nodes with no powered-on VMs
C. One node per block
D. One node
Answer: D (LEAVE A REPLY)
Nutanix permits only one node per cluster to be placed in maintenance mode at a time. This restriction protects cluster availability and ensures that sufficient compute, storage and metadata services remain operational while maintenance is performed. Before entering maintenance mode, Prism verifies whether the cluster has adequate resiliency and whether the workloads hosted by the node can be evacuated or otherwise handled safely.
For AHV, the maintenance workflow migrates eligible user VMs to other available hosts. The Controller VM on the affected node is then placed into an appropriate maintenance state so that storage services can continue through the remaining CVMs. Attempting to maintain multiple nodes simultaneously could exceed the cluster' s fault-tolerance capability, particularly if another unplanned node or disk failure occurs during the procedure.

The limitation is cluster-wide, not one node per hardware block. The absence of powered-on VMs also does not authorize every node to enter maintenance mode because AOS storage and cluster-control services still depend on quorum and service availability.
NHCF/NCA reference: Cluster maintenance, node maintenance, resiliency and hardware-maintenance procedures; NCA Section 3, "Identify hardware maintenance tasks." Official references: Prism 7.5-Node Maintenance Mode , AOS 7.5-ESXi Node Maintenance .

NEW QUESTION: 18
What is the purpose of NCC?
A. To migrate VMs between different hypervisor types
B. To configure network microsegmentation for workloads
C. To deploy new Nutanix clusters across multiple sites
D. To diagnose cluster health and identify configuration issues
Answer: (SHOW ANSWER)
Nutanix Cluster Check is cluster-resident diagnostic software that evaluates the health, configuration, and operational state of a Nutanix environment. NCC includes plugins and checks covering cluster services, storage, networking, hardware, hypervisor integration, data protection, security-related settings, and other platform dependencies. Its results identify passed checks, warnings, failures, and conditions that require administrative attention.
NCC is used during routine health validation, before and after upgrades, and when troubleshooting an active problem. Keeping NCC current is important because newer releases contain updated checks for current platform versions, supported configurations, and known failure conditions. Administrators can run broad check groups or individual plugins from Prism and the Controller VM command line.
VM migration between hypervisors is performed through Nutanix Move. Network microsegmentation is provided through Flow Network Security and related policy controls. Cluster imaging and deployment are handled through Foundation or Foundation Central. None of those functions represents NCC's role.
Nutanix defines NCC as software that helps diagnose cluster health and identify configurations qualified and recommended by Nutanix. Reference: Nutanix Cluster Check Guide . NHCF alignment: performing health checks, interpreting results, and preparing a cluster for maintenance or support investigation.

NEW QUESTION: 19
On which entity is Replication Factor configured?
A. Replication Policy
B. Protection Domain
C. Distributed Storage Pool
D. Storage Container
Answer: (SHOW ANSWER)

Replication Factor is configured on a storage container. It determines how many copies of the data associated with that container are maintained across the Nutanix distributed storage fabric. The official Prism 7.5 definition states that replication factor is the number of copies of data maintained on a storage container within a cluster. For example, RF2 maintains two copies, whereas RF3 maintains three copies and consequently requires a cluster capable of providing the corresponding fault-tolerance level.
A storage pool aggregates physical storage resources, but it does not independently establish the replication policy for guest data. A protection domain defines which entities are protected and controls snapshot or replication schedules for disaster recovery; it is not the administrative location for configuring local data-copy count. "Replication Policy" is therefore a distractor representing data-protection behavior rather than an AOS storage entity.
Replication factor also affects usable capacity because additional copies consume physical storage. Administrators must therefore align the container's replication factor with the cluster's available fault domains and required resilience. Prism requires RF3 to be configured individually for the applicable containers after the cluster supports the necessary fault-tolerance level.
NHCF/NCA reference: Cluster storage configuration, storage containers, distributed storage, and data resilience. Official references: Replication Factor and Replication Factor Configuration .

NEW QUESTION: 20

An administrator needs to receive an email notification when a specific virtual machine consistently exceeds
90% CPU utilization.
Which feature should be used?

Freecram.net

A. Storage container quota
B. Custom Pulse Policy
C. Recovery Plan
D. Custom Alert Policy
Answer: (SHOW ANSWER)
A Custom Alert Policy should be created for the target VM's CPU-utilization metric. The policy can define the VM entity, the CPU threshold of 90%, the duration or evaluation condition required to establish sustained utilization, and the alert severity. Prism Central evaluates enabled user-defined alert policies periodically and generates an alert when the configured conditions are satisfied.
Email notification additionally requires alert-email configuration. The administrator must define valid SMTP settings or the applicable notification channel and ensure the intended recipients receive alerts matching the policy. This separates metric evaluation from notification delivery: the Custom Alert Policy detects the condition, while alert-email settings determine how the resulting alert is distributed.
A storage-container quota monitors storage capacity rather than VM CPU utilization. Pulse transfers diagnostic telemetry to Nutanix Support and does not provide customer-defined VM CPU threshold policies.

A Recovery Plan orchestrates disaster-recovery operations and is unrelated to performance-alert evaluation.
Using a sustained-duration condition is important because a brief CPU spike does not necessarily indicate a capacity or performance problem. The policy should reflect how long utilization must exceed 90% before notification is justified.
NHCF/NCA reference: Custom alert policies, metric thresholds, alert evaluation, and email notifications.
Official references: User-Defined Alert Policies , Alert Metrics , and Configuring Alert Emails .

NEW QUESTION: 21
What resource should be used to ensure a successful upgrade of the hypervisors and AOS within a vSphere- based Nutanix cluster?
A. Nutanix Validated Designs
B. Hardware Replacement Documentation
C. Compatibility and Interoperability Matrix
D. Partner Software Validation
Answer: (SHOW ANSWER)
The Nutanix Compatibility and Interoperability Matrix must be consulted before upgrading AOS or the ESXi hypervisor in a vSphere-based Nutanix cluster. The matrix identifies supported combinations of AOS releases, ESXi versions, Prism components, hardware platforms and related software. This prevents the administrator from introducing a hypervisor release that the installed AOS version does not support.
Upgrade planning must consider the current and target versions together. If the target AOS version requires a different supported hypervisor level, the administrator must follow the prescribed upgrade order. Nutanix Lifecycle Manager and one-click upgrade workflows perform compatibility checks, but the matrix remains the authoritative planning resource for confirming the intended version combination.
Nutanix Validated Designs provide architectural and deployment guidance rather than an exhaustive version- compatibility decision. Hardware replacement documentation governs component service procedures. Partner Software Validation may establish support for particular third-party products, but it does not replace Nutanix' s platform compatibility matrix for AOS and ESXi.
NHCF/NCA reference: Life Cycle Management-inventory, software updates, upgrade planning and compatibility validation; NCA Section 3, "Identify software maintenance tasks." Official references: AOS
7.5-ESXi Host Upgrade , LCM-Recommended Upgrade Order .

NEW QUESTION: 22
What happens to pinned and Replication Factor 1-enabled VMs on a cluster node that is entering maintenance mode?
A. VMs remain powered on.
B. VMs are live migrated to another node.

C. VMs are restarted on another node.
D. VMs are powered off.
Answer: (SHOW ANSWER)
Pinned VMs and Replication Factor 1-enabled VMs with node affinity cannot be safely live migrated during host maintenance. Prism therefore identifies these VMs as non-migratable and requires them to be powered off before the node can complete entry into maintenance mode.
A pinned VM has a placement constraint binding it to a particular host. Moving it automatically would violate that constraint. An RF1-enabled VM has only one copy of its applicable vDisk data. When that data is associated with the node entering maintenance, another host cannot provide the same resilient storage access expected for an ordinary RF2 or RF3 VM. Powering off prevents the workload from continuing in an unsupported or inaccessible state.
Normal live-migratable VMs are moved to other nodes, subject to available CPU and memory capacity. HA- protected workloads can also be relocated according to the maintenance workflow. Pinned and RF1-affined VMs are handled separately, and Prism prompts the administrator to confirm their shutdown.
Before proceeding, the administrator should identify business impact, obtain an approved outage window, confirm cluster resiliency, and verify that affected applications can be restarted after the node exits maintenance mode.
NHCF/NCA reference: Node maintenance, VM migration restrictions, affinity, and RF1 behavior.
Official references: Node Maintenance Mode and Putting a Host into Maintenance Mode .

NEW QUESTION: 23
Which task must be performed first when preparing to upgrade an Internet-connected Nutanix cluster using Life Cycle Manager?
A. Run an inventory to detect available updates and current versions.
B. Execute upgrade plan to apply selected software components immediately.
C. Upload software bundles directly into the cluster repository storage.
D. Restart cluster services to prepare nodes for maintenance operations.
Answer: (SHOW ANSWER)
The administrator must first run an LCM inventory. Inventory establishes the cluster's current software and firmware state, refreshes component metadata, and obtains information about compatible updates available through the connected Nutanix source. LCM requires this baseline before it can accurately identify upgrade candidates and construct a valid update plan.
Executing an upgrade immediately is premature because LCM must first know which component versions are installed and which update combinations are supported. The upgrade workflow subsequently performs dependency evaluation and prechecks before applying selected components.
Directly uploading software bundles is principally associated with dark-site or disconnected workflows. An Internet-connected cluster can retrieve the required metadata and available bundles through its configured online source. Restarting cluster services is neither the initial preparation step nor a normal manual prerequisite for LCM discovery.

Inventory also refreshes the LCM framework when required and makes the resulting component information available in the management interface. After inventory completes, the administrator reviews available updates, selects applicable components, executes prechecks, resolves reported issues, and then initiates the upgrade plan. This sequence minimizes compatibility and dependency risks.
NHCF/NCA reference: Life Cycle Manager, connected-site operation, inventory, prechecks, and upgrade planning. Official references: Performing LCM Inventory and Pre-Upgrade Configurations and Recommendations .

NEW QUESTION: 24
Which option best describes Nutanix Konnector?
A. Connect an AHV hypervisor to external storage.
B. Connect an AHV hypervisor to cloud storage in the public cloud.
C. Automate onboarding of Kubernetes clusters to Prism Central.
D. Connect non-NKP Kubernetes clusters to NKP.
Answer: (SHOW ANSWER)
Nutanix Konnector is designed to automate the onboarding and registration of Kubernetes clusters with Prism Central. Nutanix describes it as a lightweight component that establishes the integration required for Prism Central to discover Kubernetes clusters and correlate them with the underlying Nutanix infrastructure.
After onboarding, Prism Central can provide infrastructure-oriented visibility into Kubernetes environments, including the relationship between clusters, virtual machines, persistent storage, networking context and relevant infrastructure alerts. Starting with supported NKP and Prism Central releases, Konnector can automatically onboard NKP-managed clusters, reducing the need for separate manual registration steps.
Option D is incorrect because Konnector's primary function is not to convert or attach a non-NKP cluster to the NKP management plane. Its management destination is Prism Central. Options A and B describe storage- connectivity functions that Konnector does not perform. External-storage and public-cloud storage integration are separate AHV, NCI or cloud-service capabilities.
The supplied answer D therefore requires correction to C.
NHCF/NCA reference: Nutanix tools, Kubernetes platform integration and centralized management; NCA Section 1, "Recognize the use cases and features of Nutanix tools." Official references: Nutanix Konnector
1.1 , Prism Central-Kubernetes Cluster Management .

NEW QUESTION: 25
An administrator needs to ensure that business-critical VMs maintain normal CPU performance. Which action can the administrator take to generate alerts when VMs are oversubscribed during peak hours?
A. Enable Pulse monitoring across the Nutanix cluster.
B. Create a custom alert based on VM CPU Ready time.

C. Increase the number of vCPU cores assigned to the critical VMs.
D. Check the average CPU usage for the VM in Prism.
Answer: (SHOW ANSWER)
CPU Ready Time measures the percentage of time that a VM's virtual CPUs are ready to execute but cannot be scheduled on physical CPU resources. It is therefore a direct indicator of processor contention or CPU oversubscription. A VM can show moderate CPU utilization while still suffering poor response times because its vCPUs are waiting for execution time.
Prism Central supports user-defined alert policies based on infrastructure metrics, including VM CPU Ready Time. The administrator can define a threshold and evaluation period so that an alert is generated when scheduling delays exceed an acceptable level during peak demand. This produces proactive notification instead of requiring periodic manual review.
Pulse sends diagnostic information to Nutanix Support but does not create the workload-specific threshold requested. Increasing vCPU allocation without first validating contention may worsen CPU Ready Time because the hypervisor must schedule additional virtual processors. Reviewing average CPU usage is retrospective and does not automatically generate an alert. The most precise control is therefore a custom CPU Ready alert. References: Prism Central 7.5 - Alert Metrics and Intelligent Operations - Chart Metrics .

NEW QUESTION: 26
An administrator is troubleshooting resource availability within the cluster. As part of the troubleshooting, the administrator wants to build a list of powered-off VMs on the Nutanix cluster.
In order for these VMs to be listed in the Prism Central VM Efficiency widget, how many days must the VMs be powered off to be considered a Dead VM?
A. 5
B. 15
C. 30
D. 45
Answer: (SHOW ANSWER)
Prism Central classifies a virtual machine as a Dead VM after it has remained powered off for at least 30 days. The VM Efficiency capability uses this duration to distinguish temporarily stopped workloads from VMs that are sufficiently inactive to justify administrative review. Consequently, a VM powered off for only
5 or 15 days does not meet the classification threshold, while 45 days exceeds-but does not define-the required minimum.
A Dead VM is not automatically deleted. The classification is an efficiency finding intended to identify resources that may no longer be required. Even while powered off, the VM can continue consuming storage through its virtual disks, snapshots, recovery points, and associated configuration objects. Administrators should therefore validate application ownership, retention requirements, protection policies, and compliance obligations before deleting or archiving it.
The VM Efficiency widget is part of Prism's broader behavioral and capacity-optimization functionality.

Other inefficiency classifications can identify overprovisioned or constrained VMs by evaluating historical resource behavior. The 30-day rule is specifically associated with the powered-off Dead VM condition and should not be confused with utilization-based efficiency findings.
NHCF/NCA reference: Prism Central monitoring, VM efficiency, resource optimization, and behavioral analysis. Official reference: Behavioral Learning Tools and VM Efficiency , specifically the Dead VM classification.

NEW QUESTION: 27
Which protocols does the Network Visualizer leverage to discover and display network connections and statistics?
A. BGP and OSPF
B. DHCP and PXE
C. CDP/LLDP and SNMP
D. STP and VTP
Answer: (SHOW ANSWER)
Nutanix Network Visualizer uses Link Layer Discovery Protocol or Cisco Discovery Protocol to determine the relationship between an AHV host interface and the corresponding first-hop physical-switch port. It can then use configured SNMP access to retrieve additional switch, port and operational statistics. The combined protocol set in option C therefore enables both topology discovery and the display of useful physical-network information.
LLDP is vendor-neutral, while CDP is primarily associated with Cisco networking equipment. Either protocol can advertise the local port identity needed to construct the topology mapping. SNMP supplements that discovery by providing management-plane information obtained from supported switches. Nutanix documentation notes that when LLDP or CDP information is unavailable, SNMP data may be used to assist mapping, although complete and accurate visualization depends on supported devices and appropriate configuration.
BGP and OSPF exchange routing information; they are not the protocol set used by Network Visualizer for host-to-switch discovery. DHCP and PXE provide address allocation and network boot services. STP prevents Layer 2 loops, while VTP distributes VLAN information in certain Cisco environments.
NHCF/NCA reference: Networking-monitoring networks in Prism Element and Prism Central; NCA Section 2, "Identify virtual networking tasks." Official reference: Prism 7.5-Network Visualization Prerequisites .

NEW QUESTION: 28
When migrating ESXi VMs to an AHV-based cluster using Nutanix Move, what should be in place to ensure a successful migration?
A. The hosts must be in maintenance mode.
B. VMware Tools must be installed and up to date.
C. The VMs' hardware version must be 9 or later.
D. The ESXi version must be 7.x or later.

Answer: (SHOW ANSWER)
Nutanix Move requires VMware Tools to be installed and current on guest VMs for supported ESXi-to-AHV migration workflows. VMware Tools provides guest-level information and operations needed during migration preparation and cutover. Its services support tasks such as obtaining guest details, coordinating shutdown or restart actions, and completing supported guest-side preparation.
Placing ESXi hosts in maintenance mode would normally evacuate or stop workloads and is not a standard migration prerequisite. Move is designed to migrate selected running workloads while the source infrastructure remains operational until cutover. Option C is incorrect because official Move requirements specify virtual hardware version 7 or later for the relevant support condition, not version 9 or later. Option D is also too restrictive: supported Move releases accommodate documented ESXi versions beyond only the 7.x family, subject to the current compatibility matrix.
Administrators must also verify network reachability, credentials, guest-operating-system support, source and target compatibility, and adequate target resources. However, among the choices, the explicit guest requirement is current VMware Tools. Reference: Nutanix Move 6.2 - Requirements .
NHCF alignment:
Nutanix Move prerequisites, cross-hypervisor migration, and workload preparation.

NEW QUESTION: 29
How should a license be removed from a cluster for reuse?
A. Disable license features from the cluster.
B. Unregister the cluster from Prism Central.
C. Delete the cluster license summary file from Prism Central.
D. Unlicense the cluster using License Manager interface.
Answer: (SHOW ANSWER)
A cluster must be unlicensed through the License Manager interface when its allocated entitlement needs to be released for reuse. The unlicensing workflow updates the authoritative licensing state and returns the applicable entitlement or capacity to the organization's available license pool. This is the controlled administrative process recognized by Nutanix licensing services.
Disabling licensed features does not remove the license assignment. It changes feature usage or configuration while leaving the entitlement associated with the cluster. Unregistering a cluster from Prism Central only terminates its centralized management relationship; registration and licensing are independent administrative states. A cluster can therefore be unregistered while still consuming licenses. Deleting a locally generated summary file is equally ineffective because the file reports licensing information but does not constitute the authoritative license allocation.
The administrator should verify connectivity and licensing status, open License Manager, select the applicable cluster, and use the unlicense operation. Any required synchronization with the Nutanix licensing service must be allowed to complete before the entitlement is treated as available elsewhere. This is especially important before decommissioning hardware, rebuilding a cluster, or transferring capacity to another Nutanix environment.
NHCF/NCA reference: License types, cluster licensing, entitlement management, and Prism administrative functions. Official references: Nutanix License Manager Guide and Prism Licensing .

NEW QUESTION: 30

What crucial step must be taken before successfully stopping a Nutanix cluster?

A. Shut down all Nutanix Controller VMs

B. Disable the Remote Support Tunnel

C. Disable the protection domains

D. Shut down all guest virtual machines

Answer: (SHOW ANSWER)

All guest virtual machines must be shut down before stopping a Nutanix cluster. A cluster shutdown terminates the AOS services provided by the Controller VMs and ultimately shuts down the hypervisor hosts.

Leaving user workloads running would expose those VMs to abrupt power loss, interrupted write operations and possible application or filesystem corruption.

The supported sequence begins with a controlled shutdown of the guest operating systems. Administrators should verify that no user VMs remain powered on, check cluster health and then execute the supported cluster-stop procedure. The cluster process coordinates the shutdown of AOS services and Controller VMs; therefore, administrators must not manually power off every CVM as the first action.

Disabling the Remote Support Tunnel has no bearing on the safe shutdown of storage and compute services.

Protection domains generally do not need to be disabled for an ordinary cluster shutdown, although synchronous-replication configurations can require additional site-specific preparation.

NHCF/NCA reference: Cluster maintenance, controlled shutdown and service-management procedures; NCA Section 3, "Identify hardware maintenance tasks." Official references: AHV-Shutting Down a Cluster , Shutting Down DR Clusters .

NEW QUESTION: 31

Which Nutanix product provides unstructured file-storage services through SMB and NFS for home directories and user profiles on a Nutanix cluster?

A. Files

B. Move

C. Volumes

D. Objects

Answer: (SHOW ANSWER)

Nutanix Files is the scale-out file-storage service designed to provide SMB shares and NFS exports for unstructured data. Typical use cases include user home directories, roaming or virtual-desktop profiles, departmental shares, application logs, archives, and general enterprise file services. Files deploys File Server VMs that present the client-facing file protocols while using the Nutanix distributed storage layer underneath.

Nutanix Move is a workload-migration utility used to transfer VMs between supported source and target platforms; it does not provide file shares. Nutanix Volumes supplies block-storage services,

commonly exposed through iSCSI, for applications that require block devices rather than SMB or NFS file systems.
Nutanix Objects provides S3-compatible object storage and is intended for applications using object APIs rather than conventional file protocols.
The presence of SMB, NFS, home directories, and user profiles makes the use case specifically file-oriented.
Nutanix documentation identifies Files as a repository for unstructured data and lists home directories and user profiles among its standard workloads. Reference: Nutanix Files Storage Technical Note . NHCF alignment: differentiating Files, Volumes, Objects, and migration services by protocol and workload type.

Valid NCA-7.5 Dumps shared by EduDump.com for Helping Passing NCA-7.5 Exam!
EduDump.com now offer the **newest NCA-7.5 exam dumps**, the EduDump.com NCA-7.5 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com NCA-7.5 dumps with Test Engine here:
https://www.edudump.com/exams/Nutanix/NCA-7.5/premium/ (**91** Q&As Dumps, **35%OFF**
Special Discount Code: freecram)

NEW QUESTION: 32
In a 2-node cluster, which configuration option is supported?
A. Compression
B. Network Segmentation
C. Deduplication
D. Cluster Expansion
Answer: (SHOW ANSWER)
Compression is supported on a Nutanix two-node cluster. Compression reduces physical capacity consumption by encoding eligible data more efficiently before it is retained in the distributed storage fabric. It remains available because its operation does not require the additional node count or data-placement structure demanded by certain other storage-efficiency features.
Nutanix explicitly identifies deduplication as unsupported for the two-node configuration while stating that compression is supported. Cluster expansion is also not a supported scale-out path for a cluster originally created as a two-node cluster. If a larger cluster is required, the environment must be designed and deployed as the appropriate multi-node configuration rather than treating the two-node deployment as an incrementally expandable starting point.
Two-node clusters also have specialized architecture. They use RF2 across the nodes and require a Witness VM in a separate failure domain to resolve leadership when communication between the two nodes is lost.
The Witness does not contain workload data; it provides quorum-related arbitration during applicable failures.

Network configuration is similarly subject to two-node guidelines and limitations, making Network Segmentation an incorrect selection in this answer set.
The question tests recognition that a two-node cluster retains selected AOS capabilities but is not functionally identical to a conventional three-node-or-larger cluster.
NHCF/NCA reference: Two-node cluster capabilities, limitations, Witness operation, and storage efficiency.
Official references: Two-Node Clusters and Two-Node Cluster Guidelines .

NEW QUESTION: 33
When managing licenses manually, what is the first step to start the licensing workflow?
A. Use the Prism Element interface to download a license summary file.
B. Contact Support to request a license summary file.
C. Use Prism Central Admin Center to download a cluster summary file.
D. Use Prism Central Admin Center to download a license summary file.
Answer: (SHOW ANSWER)
The manual licensing workflow begins by downloading a cluster summary file from the Licensing area in Prism Central Admin Center. The cluster summary file contains the deployment and consumption information that Nutanix License Manager requires to identify the relevant clusters and determine which entitlements can be allocated.
After downloading the cluster summary file, the administrator uploads it to Nutanix License Manager in the Support Portal, selects or allocates the appropriate licenses and generates a license summary file. That license summary file is subsequently downloaded from the portal and uploaded back into Prism Central. Therefore, the cluster summary file is the workflow's initial outbound file, while the license summary file is the resulting inbound file.
Option D reverses this sequence: Prism Central does not generate the license summary file at the start. Option A confuses the file type and management interface. Prism Element may appear in certain legacy-license procedures, but the question specifically addresses Prism Central Admin Center licensing. Contacting Support is not required for the standard manual workflow.
NHCF/NCA reference: Platform administration and licensing tasks; NCA Section 2, "Describe licensing tasks." Official references: Manage Cluster-Based Licenses , Nutanix License Manager Guide .

NEW QUESTION: 34
What would cause storage compression to become automatically disabled?
A. Storage writes are high
B. Storage reads are high
C. Disk capacity is high
D. Metadata usage is high
Answer: (SHOW ANSWER)
Nutanix AOS automatically disables compression when metadata usage becomes high. The Prism 7.5 storage documentation explicitly identifies this condition and states that an alert is generated

when automatic disablement occurs. Consequently, option D is the verified answer, replacing the supplied answer A.
Compression requires AOS metadata resources to track compressed extents and the associated data layout. If metadata consumption reaches a protective threshold, continuing compression could place additional pressure on the metadata services required for normal storage operation. AOS therefore prioritizes platform stability and metadata availability by disabling compression and notifying the administrator. The corresponding health alert should be investigated to determine why metadata usage has increased and whether workload characteristics or another storage condition requires remediation.
High write activity may increase storage workload and influence when post-process compression can execute, but it is not the documented automatic-disable trigger in this question. High reads similarly do not directly cause compression to be disabled. High physical-capacity consumption generates its own capacity warnings and corrective requirements but is distinct from the metadata threshold governing compression.
NHCF/NCA reference: Storage Concepts-capacity optimization, compression and storage-container administration; NCA Section 2, "Discuss storage operations." Official references: Prism 7.5-Compression , Prism 7.5-Storage Alert Reference .

NEW QUESTION: 35
An administrator needs to increase the available storage of an existing AHV virtual machine by adding a new vDisk.
Which method is supported to perform this task?
A. Create a new VM and migrate the workloads to it.
B. Increase the size of the AHV host's local disk.
C. Modify and increase the storage-pool size from within Prism Element.
D. Edit the VM and create a new disk, specifying the storage container.
Answer: (SHOW ANSWER)
An administrator can add storage directly to an existing AHV VM by editing its configuration and selecting Add New Disk. In the disk configuration, the administrator specifies the device type, disk size, storage container, and supported controller or bus settings. The new vDisk is then attached to the VM and presented to the guest operating system, where it can be initialized, partitioned, and formatted.
Creating another VM is unnecessary because AHV supports modifying the virtual hardware of an existing VM. Increasing a host's local disk does not directly create guest storage; physical disks contribute capacity to the Nutanix distributed storage fabric rather than being individually assigned to a VM. Similarly, changing the size of a storage pool is an infrastructure-capacity operation and does not attach a new virtual disk to the selected workload.
A storage container holds the vDisk and provides access to the cluster's distributed storage resources.
Selecting the appropriate container during disk creation is therefore part of the supported workflow.

Reference: Prism 7.5 - Creating and Attaching a VM Disk and Prism 7.5 - Managing an AHV VM . NHCF alignment: VM storage administration and vDisk management.

NEW QUESTION: 36
A virtual machine needs to be placed on a different storage location within the cluster.
Which task should be performed?
A. Change vDisk Placement
B. Assign a Category
C. Set Image Placement
D. Apply a Storage Policy
Answer: (SHOW ANSWER)
Moving a VM's existing storage to a different storage container requires a vDisk placement change, implemented through live vDisk migration across storage containers. Nutanix supports migrating the attached vDisks of a running VM to a target container without deleting or shutting down the VM, subject to the documented requirements and limitations. The administrator can move an individual vDisk or all supported vDisks belonging to the VM.
A category assigns metadata used for organization, governance, automation, and policy association; it does not directly relocate existing disk data. Image placement policies determine which clusters receive images from the image service, not where an existing VM's attached vDisks are stored. A storage policy controls storage characteristics such as QoS or other supported data-service properties. It is not the general storage- container migration task.
The keyed answer D is particularly problematic because Nutanix documents that a vDisk cannot be migrated to another container while a storage policy is configured for the associated VM. Applying such a policy can therefore prevent the requested migration rather than accomplish it. References: Migrating a vDisk to Another Container and Storage Policy Management . NHCF alignment: storage operations and vDisk management.

NEW QUESTION: 37
Which component is required to prevent split-brain scenarios on a two-node ROBO cluster?
A. Prism Central
B. Prism Element
C. Foundation Central
D. Witness VM
Answer: (SHOW ANSWER)
A two-node cluster lacks the natural majority available in a conventional cluster containing three or more voting members. If communication between its two nodes is interrupted, each node could otherwise interpret the other node as failed and attempt to continue independently. That condition creates a split-brain risk because both sides could consider themselves authoritative.
The external Witness VM supplies the independent arbitration required for the two-node configuration. It is deployed in a separate failure domain and maintains communication with the cluster. During a node or communication failure, witness participation allows the cluster to distinguish

the surviving authoritative side and avoid both nodes independently continuing cluster operations. Separating the witness from the two data- bearing nodes is critical because placing all voting components in the same failure domain would undermine the tie-breaking design.
Prism Element is the local cluster-management interface, while Prism Central provides centralized management across clusters. Neither substitutes for an independent quorum witness. Foundation Central supports centralized discovery and deployment activities and does not arbitrate two-node cluster ownership.
Nutanix documentation identifies the external Witness VM as part of the supported two-node ROBO configuration. References: Prism 7.5 - Two-Node Clusters and Witness for Two-Node Clusters .
NHCF alignment: valid cluster configurations, availability, and failure handling.

NEW QUESTION: 38
Which Nutanix product provides capabilities for advanced insights and automation for troubleshooting applications?
A. Nutanix Cloud Infrastructure
B. Nutanix Enterprise AI
C. Nutanix Cloud Clusters
D. NCM Intelligent Operations
Answer: (SHOW ANSWER)
NCM Intelligent Operations provides centralized monitoring, advanced operational analytics, capacity planning, anomaly detection, resource optimization, reporting, and automation. It evaluates infrastructure and workload behavior across managed environments and helps administrators correlate performance conditions with affected applications and resources. Its automation capabilities can initiate predefined actions or playbooks when specified events, alerts, or operational conditions occur.
Nutanix Cloud Infrastructure supplies the core compute, storage, virtualization, networking, resilience, and management foundation. Although it exposes operational metrics through Prism, it is not the dedicated advanced analytics and automation product described in the question. Nutanix Enterprise AI is used to deploy, manage, monitor, and test large language models and inference endpoints. Nutanix Cloud Clusters extends Nutanix infrastructure into supported public-cloud environments but does not itself provide the specialized operational-intelligence feature set.
Intelligent Operations is specifically positioned as an end-to-end operations, monitoring, optimization, and planning solution. It can identify anomalous behavior, inefficient VMs, resource constraints, and capacity risks while reducing manual troubleshooting through automated workflows. References: NCM Intelligent Operations Overview and Nutanix Cloud Monitoring and Intelligent Operations .
NHCF alignment:
recognizing Nutanix management and operations products.

NEW QUESTION: 39
Which bond type is the default for newly created AHV clusters?
A. Active-Backup

B. No Uplink Bond
C. Balance-slb
D. Balance-tcp
Answer: (SHOW ANSWER)
Active-Backup is the default bond configuration for newly created AHV clusters. In this mode, one physical uplink actively forwards traffic while another eligible uplink remains available as a standby. If the active link fails, AHV moves traffic to the backup link, providing straightforward network redundancy without requiring LACP or a switch-side link-aggregation configuration.
The principal advantage of Active-Backup is operational compatibility. Because only one uplink forwards traffic at a given time, the upstream switch ports do not need to be combined into a link aggregation group.
However, the configuration does not aggregate the bandwidth of multiple physical NICs during normal operation.
Balance-slb provides active-active source load balancing using MAC pinning and can improve aggregate host bandwidth without requiring LACP. Balance-tcp uses active-active LACP and requires matching configuration on the physical switches. Both are deliberately configured alternatives rather than the default.
"No Uplink Bond" would eliminate bond-level redundancy and is not the standard newly created cluster configuration.
NHCF/NCA reference: AHV networking basics, virtual switches, uplinks and bond types; NCA Section 2,
"Identify virtual networking tasks." Official references: Prism 7.5-Network Configuration , Prism Central- Creating a Virtual Switch .

NEW QUESTION: 40
How many Controller VMs (CVMs) will a five-node HCI cluster have?
A. 3
B. 4
C. 5
D. 6
Answer: (SHOW ANSWER)
A conventional five-node Nutanix HCI cluster has five Controller VMs-one CVM running on each node.
AOS operates as a distributed storage and data-services system rather than relying on a single centralized storage controller. The per-node CVM architecture provides local access to the node's storage devices while cooperating with the CVMs on the other nodes to form the distributed AOS storage fabric.
The CVMs collectively provide services associated with data placement, replication, metadata, caching, compression, deduplication, snapshots, and cluster storage management. If a CVM becomes unavailable, the host can access storage services through another CVM, subject to cluster

health and connectivity. The remaining CVMs also maintain the distributed cluster services required for continued operation.
The number of CVMs therefore normally matches the number of HCI nodes: three nodes produce three CVMs, four nodes produce four, and five nodes produce five. The count does not represent the replication factor or the number of simultaneously tolerated failures.
Compute-only nodes are a specific exception because they do not run a CVM. However, the question explicitly describes a five-node HCI cluster, so each node contributes storage and runs a Controller VM.
NHCF/NCA reference: Nutanix HCI architecture, AOS, nodes, and Controller VMs. Official references:
Infrastructure Resilience and Compute-Only Nodes .

NEW QUESTION: 41
What are two requirements to upgrade Nutanix Guest Tools on a VM? (Choose two.)
A. VM powered off
B. DNS entry for the VM
C. VM powered on
D. Empty CD-ROM drive
Answer: (SHOW ANSWER)
The VM must be powered on so that the Nutanix Guest Agent and the guest operating system can participate in the NGT installation or upgrade workflow. The NGT software runs inside the guest, so a powered-off VM cannot execute the installer, communicate status, or complete the upgrade operation.
An available IDE or SATA CD-ROM slot is also required because Prism mounts the Nutanix Guest Tools ISO to the VM during configuration, installation, or upgrade. If the VM's virtual CD-ROM is already occupied by another ISO, the NGT media cannot be attached through the expected workflow. Nutanix documentation therefore requires at least one empty supported CD-ROM device. After installation or reconfiguration, supported versions can automatically eject the mounted NGT media.
A DNS entry is not a universal prerequisite for upgrading NGT. Although functional name resolution may support broader management operations, the upgrade requirement is based on VM power state and the ability to mount the NGT ISO. Powering off the VM directly conflicts with the installation requirement. References:
Prism Central 7.5 - Setting Up NGT and Nutanix Guest Tools Requirements . NHCF alignment: VM guest- tools administration.

NEW QUESTION: 42
In an AHV Metro environment, which feature must be installed on a Windows VM to ensure successful cross- cluster migration?
A. VirtIO
B. VMware Tools

C. Witness service
D. NGT
Answer: (SHOW ANSWER)
Nutanix Guest Tools must be installed and configured on the protected Windows VM when required by the AHV Metro and recovery-plan workflow. Nutanix DR documentation specifies supported NGT versions as a guest-level prerequisite for protected entities participating in relevant cross-cluster operations.
NGT supplies guest integration that allows the Nutanix platform to coordinate VM-specific operations during protection and recovery workflows. It can support application-consistent processing, guest communication and the handling of attached resources where applicable. For Windows workloads, installing NGT ensures that the guest participates correctly in the supported Nutanix DR workflow rather than being treated solely as an unmanaged virtual-machine object.
VirtIO supplies the storage and network drivers required for Windows to operate efficiently on AHV, but it is not the Metro-specific guest integration requested by the question. VMware Tools applies to VMware guest integration and is not the corresponding AHV component. The Witness service is an external availability component used to arbitrate failures between Metro sites; it is not installed inside each Windows VM.
NHCF/NCA reference: Nutanix Guest Tools, VM protection and Nutanix Disaster Recovery capabilities; NCA Section 1, "Recognize Nutanix solutions, use cases and features." Official references: Prism Central
7.5-Creating a Recovery Plan , Prism Central 7.5-Cross-Cluster Live Migration .

Valid NCA-7.5 Dumps shared by EduDump.com for Helping Passing NCA-7.5 Exam!
EduDump.com now offer the **newest NCA-7.5 exam dumps**, the EduDump.com NCA-7.5 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com NCA-7.5 dumps with Test Engine here:
https://www.edudump.com/exams/Nutanix/NCA-7.5/premium/ (**91** Q&As Dumps, **35%OFF**
Special Discount Code: freecram)