Nutanix.NCP-US-6.10.v2026-03-16.q36

Exam Code:	NCP-US-6.10
Exam Name:	Nutanix Certified Professional - Unified Storage (NCP-US) v6.10
Certification Provider:	Nutanix
Free Question Number:	36
Version:	v2026-03-16
# of views:	114
# of Questions views:	360
https://www.dumpsdb.com/dumps/Nutanix/NCP-US-6.10/Nutanix.NCP-US-6.10.v2026-03-16.q36	

NEW QUESTION: 1
Which Nutanix Objects metric provides the total input requests per second of a bucket?
A. Puts
B. Throughput
C. Gets
D. NFS Reads
Answer: (SHOW ANSWER)
In Nutanix Objects metrics:
* Puts: Measures PUT requests per second (object uploads), representing input operations.
* Gets (Option C): Measures output (download) requests.
* Throughput (Option B): Reports bandwidth (MB/s), not request rate.
* NFS Reads (Option D): Specific to NFS access, not general bucket input.
Reference:Nutanix Objects Administration Guide:
"The Puts metric tracks the number of successful object write operations (PUTs) per second to a bucket, indicating input request volume."(Chapter: "Monitoring Object Store Performance")
Nutanix Unified Storage (NCP-US) Study Material:
"Use the 'Puts' metric to monitor ingress request rates for capacity planning and performance tuning."(Section:
"Objects Performance Analysis")

NEW QUESTION: 2
Which workload type describes I/O sizes for read and write operations that are less than or equal to 16 KB while file sizes are equal to 10 MB or more?

A. Sequential
B. Asynchronous
C. Random
D. Default
Answer: (SHOW ANSWER)
The workload type that describes I/O sizes for read and write operations that areless than or equal to 16 KB while file sizes are10 MB or moreisRandom. In Nutanix Files, workload types are used to optimize share performance based on I/O patterns. Small I/O sizes (#16 KB) indicate a random access pattern, as opposed to sequential, even if the files themselves are large (#10 MB). This is common in workloads like databases or virtual desktops, where small, non-contiguous I/O operations are performed on larger files.
TheNutanix Unified Storage Administration (NUSA)course states, "A Random workload type in Nutanix Files is characterized by small I/O sizes, typically 16 KB or less, regardless of file size, as it reflects random access patterns rather than sequential ones." The Random workload type optimizes the share for such patterns by adjusting caching, prefetching, and data placement to handle frequent small I/O operations efficiently, even when the files are large.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further elaborates that "workloads with I/O sizes of 16 KB or less, even on large files (e.g., 10 MB or more), are classified as Random, as the small I/O size indicates non-sequential access patterns." Large file sizes do not necessarily imply sequential I/O; the I/O size itself determines the workload type, and 16 KB or less is typical of random access.
The other options are incorrect:
* Sequential: Sequential workloads involve larger I/O sizes (typically >64 KB) and contiguous access patterns, such as those seen in media streaming or backups, not small I/O sizes like 16 KB or less.
* Asynchronous: Asynchronous is not a workload type in Nutanix Files; it may refer to replication or I/O handling methods but is not relevant here.
* Default: The Default workload type applies a balanced configuration but does not specifically optimize for small I/O sizes like the Random type does.
The NUSA course documentation emphasizes that "I/O sizes of 16 KB or less, even with large file sizes, indicate a Random workload type in Nutanix Files, ensuring optimal performance for random access patterns." References:
Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Files: "Understanding workload types based on I/O patterns." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 2: Configure and Utilize Nutanix Unified Storage, Subtopic: "Defining workload types for Nutanix Files shares." Nutanix Documentation (https://www.nutanix.com), Nutanix Files Administration Guide: "Workload type definitions for share optimization."

NEW QUESTION: 3
Question:

An administrator needs to move infrequently accessed data to lower-cost storage based on file type and owner, and automatically recall data if data access frequency has increased.
What should administrator do to satisfy these requirements?
A. Configure Advanced tiering in Data Lens.
B. Create a Lifecycle Rule in Objects Buckets tab.
C. Create an SSR-enabled share in Files.
D. Configure Smart tiering in Files.
Answer: D (LEAVE A REPLY)
Smart Tieringin Nutanix Files is a built-in feature that allows administrators toautomatically move infrequently accessed data(cold data) tolower-cost storage tiers(like NFS or S3-compatible storage). It also supportsautomatically recalling dataif it becomes hot (frequently accessed) again.
According to NUSA course details:
"Smart Tiering policies in Nutanix Files allow administrators to define rules based on file metadata (type, size, owner) and last access time. Cold data is tiered off to cheaper storage, and Files can recall the data if needed, ensuring transparent access for users." Key reasons why Smart Tiering is the solution:
* Automatically identifies cold data(based on access patterns).
* Moves cold data to external or cheaper storagetransparently.
* Re-hydrates dataautomatically if it becomes hot again, maintaining performance and user experience.
The other options:
Advanced tiering in Data Lens- Data Lens is for analytics and reporting, not for moving data.
Lifecycle Rules in Objects- manages data lifecycle for object buckets, not Files shares.
SSR (Self-Service Restore)- is for file recovery, not data tiering.
Thus, the administrator shouldconfigure Smart Tiering in Nutanix Filesto satisfy the requirement.

NEW QUESTION: 4
At what level of granularity can Smart DR replicate?
A. Volume
B. Bucket
C. Share
D. File
Answer: C (LEAVE A REPLY)
Smart DR (Disaster Recovery) is a feature within Nutanix Unified Storage (NUS), specifically designed to facilitate data replication and disaster recovery for Nutanix Files, which is the file storage service component of NUS. Nutanix Unified Storage integrates file, object, and block storage services, but Smart DR is primarily associated with the file storage functionality provided by Nutanix Files. To determine the level of granularity at which Smart DR operates, we need to examine how it handles replication within this context.
Understanding the Options

* Volume: In Nutanix terminology, a volume typically refers to a logical storage unit used in block storage services (e.g., Nutanix Volumes). It can contain multiple files or datasets and is managed at a higher abstraction level.
* Bucket: A bucket is a container used in object storage (e.g., Nutanix Objects) to store objects, akin to a directory but specific to object-based storage systems.
* Share: In Nutanix Files, a share refers to a file share (accessible via SMB or NFS protocols), which contains files and directories that are made available over a network for user access.
* File: This represents an individual file, the smallest unit of data within a storage system.
Smart DR's purpose is to ensure data availability and consistency for disaster recovery scenarios, which implies that the replication granularity should support recovering cohesive sets of data rather than fragmented pieces that could lead to inconsistencies.
Smart DR and Nutanix Files
According to the Nutanix Unified Storage documentation, Smart DR is specifically tailored for Nutanix Files to enable replication of file shares for disaster recovery. The key evidence comes from the NCP-US and NUSA course materials, which state:
"NUS also offers Smart DR to facilitateshare-level data replicationand file server-level disaster recovery." (Reference: Nutanix Unified Storage Administration (NUSA) Study Guide, Section on Disaster Recovery Features for Nutanix Files) This excerpt explicitly indicates that Smart DR performs replication at theshare level. In Nutanix Files, a share is a logical entity that groups files and directories together, accessible via protocols like SMB (Server Message Block) for Windows environments or NFS (Network File System) for UNIX/Linux environments.
When configuring Smart DR, administrators select specific shares to replicate to a remote site, ensuring that the entire share-including all its files and directory structures-is replicated as a single unit. This approach maintains data consistency and simplifies recovery by allowing the entire share to be restored in a disaster scenario.
Why Not the Other Options?
* Volume: While Nutanix Volumes (block storage) supports replication through features like Protection Domains or asynchronous replication, Smart DR is not documented as a feature for block storage replication. Protection Domains, for instance, operate at the VM or volume group level, not under the Smart DR umbrella. Thus, "Volume" is not the correct granularity for Smart DR.
* Bucket: In Nutanix Objects (object storage), replication can occur at the bucket level, but this is managed through different mechanisms, such as object replication policies, not Smart DR. The documentation does not associate Smart DR with bucket-level replication, making "Bucket" incorrect.
* File: Replicating individual files would be highly granular and impractical for disaster recovery, as it risks inconsistencies (e.g., missing related files or directory structures). While Nutanix Files supports file-level operations, Smart DR does not allow administrators to configure replication for individual files within a share. The replication unit is the share itself, ruling out "File." Configuration in Practice In the Nutanix Prism interface, when setting up Smart DR for Nutanix Files, administrators define replication policies by selecting specific file shares. The process involves:

* Identifying the source file server and the shares to replicate.
* Configuring a remote target (e.g., another Nutanix Files instance).
* Scheduling replication to ensure data is copied to the DR site.
This is consistent with the NUSA course, which emphasizes that:
"Smart DR enables administrators to configure replication at the share level, ensuring that all data within the share is protected and recoverable." (Reference: Nutanix Unified Storage (NCP-US) Study Guide, Module on Configuring Disaster Recovery) Clarifying Scope While Nutanix Unified Storage encompasses file, object, and block services, Smart DR is distinctly a feature of Nutanix Files. For object storage (Nutanix Objects), replication is handled at the bucket level via separate features, and for block storage (Nutanix Volumes), replication uses mechanisms like synchronous or asynchronous replication at the volume group level. However, the question specifically pertains to Smart DR, and the documentation consistently ties this feature to share-level replication.
Conclusion
The level of granularity for Smart DR replication is theshare, as it replicates entire file shares within Nutanix Files to ensure data consistency and effective disaster recovery. Among the provided options-Volume, Bucket, Share, and File-the correct answer is "Share," corresponding to option C.
References:
Nutanix Unified Storage (NCP-US) Study Guide, Module on Disaster Recovery and Replication.
Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Files and Smart DR Configuration.

NEW QUESTION: 5
An administrator has recently added several NGT-enabled VMs with in-guest iSCSI initiators to a Volume Group (VG) using IP addresses in the VG allowlist. Several days later, the administrator restored the VG, after which the VMs lost connectivity to the Volume Group.
What should the administrator have done differently to prevent this from happening?
A. Use the iSCSI IQN entry in the VG allowlist.
B. Use the VM UUID in the VG allow list.
C. Use the VM hostname in the VG allowlist.
D. Use the NIC MAC address of the VM's in the VG allow list.
Answer: (SHOW ANSWER)
Volume Groups (VGs) require persistent identifiers for initiators in the allowlist. IP addresses can change during VM restores/reboots, breaking connectivity. The iSCSI Qualified Name (IQN) is a static, unique identifier for iSCSI initiators and persists across VM operations, ensuring stable access.
* Option B/C/D (VM UUID, hostname, MAC): These are unrelated to iSCSI authentication. Nutanix Volume Groups exclusively use IQNs or IPs (not recommended) for allowlisting.
Reference:Nutanix Unified Storage Administration (NUSA) Course Study Guide:
"Always use iSCSI IQN in Volume Group allowlists for NGT-enabled VMs. IP addresses are ephemeral and may change after restores, causing connectivity loss."(Section: "Configuring

Volume Group Access Control") (Module: "Nutanix Volumes Best Practices") Nutanix Volumes Documentation:
"For persistent iSCSI connectivity, configure the allowlist with initiator IQNs instead of dynamic IP addresses."(Source: Volumes Configuration Guide, "Allowlist Management")

NEW QUESTION: 6
An administrator is in the process of migrating shares from one Nutanix Files cluster in the primary data center (DC) to another Files cluster running in a new DC that has been built. The administrator is using Smart DR to perform this migration as it provides less downtime. Upon a successful sync during a scheduled maintenance window, users are unable to save to the new share. How should the administrator resolve the issue?
A. Enable Continuous Availability
B. Enable Self-Service Restore
C. Set share read-only to false
D. Set share type to multiprotocol
Answer: C (LEAVE A REPLY)
Nutanix Files is a software-defined, scale-out file storage solution within Nutanix Unified Storage, offering SMB and NFS file services to clients. Smart DR (Disaster Recovery) is a feature designed to protect and migrate file shares between Nutanix Files clusters with minimal downtime, making it ideal for planned migrations, such as moving shares from a primary data center to a new data center. Smart DR leverages replication to synchronize data between the source (primary) and target (new) clusters, followed by a switchover process during a maintenance window.
In this scenario, the administrator has successfully synchronized the data using Smart DR, but post-sync, users cannot save files to the new share, indicating a lack of write access. This is a common situation in migration workflows, where the target share may default to a read-only state after synchronization to ensure data consistency until the migration is fully committed.
The NUSA course, under the "Troubleshooting Nutanix Unified Storage" module, addresses such issues, noting that after a Smart DR sync, the target share's permissions must be adjusted to allow write operations.
The specific resolution involves modifying the share's read-only attribute. Let's analyze the options:
* A. Enable Continuous Availability: Continuous Availability is a high-availability feature in Nutanix Files that ensures share accessibility during failures by maintaining active-active configurations across nodes. While beneficial for uptime, it does not address the specific issue of write access post-migration, as it pertains to availability rather than permissions. The NCP-US study guide mentions this feature under "Section 3: Analyze and Monitor Nutanix Unified Storage," but it's unrelated to this troubleshooting context.
* B. Enable Self-Service Restore: This feature allows end-users to recover their own files from snapshots, enhancing user autonomy and reducing administrative overhead. However, it is

designed for data recovery, not for resolving share-level access issues like write permissions. The NUSA course covers this in the "Data Protection" section, confirming its irrelevance here.
* C. Set share read-only to false: This option directly addresses the problem. In Nutanix Files, shares can be configured with a read-only attribute, often set to true during replication or migration to prevent premature writes. After a successful Smart DR sync, the administrator must update this attribute on the target cluster to allow write access. The NUSA course documentation, under "Managing File Shares," states: "Post-migration, ensure the share's read-only setting is disabled (set to false) to enable write operations." This can be done via the Prism interface or CLI, making it the precise solution.
* D. Set share type to multiprotocol: Multiprotocol shares support both SMB and NFS access, catering to diverse client environments. While this might be relevant during initial share configuration, it does not resolve the write access issue post-migration, as the problem is permission-based, not protocol- related. The NCP-US study guide discusses this under "Section 2: Configure and Utilize Nutanix Unified Storage," but it's not applicable here.
The correct resolution isC. Set share read-only to false. After the Smart DR sync, the administrator must access the Nutanix Files management interface (e.g., Prism Central), locate the migrated share on the new cluster, and modify its properties to disable the read-only setting. This action ensures users can save files, completing the migration process seamlessly.
:
Nutanix Unified Storage (NCP-US) Study Guide, Section 4: Troubleshoot Nutanix Unified Storage, Subsection: Post-Migration Issues.
Nutanix Unified Storage Administration (NUSA) Course, Module: Troubleshooting Nutanix Unified Storage, Topic: Managing Share Permissions After Smart DR Migration.

NEW QUESTION: 7
Question:
The administrator creates an S3 bucket as the backup target. While creating the Nutanix Objects endpoint to the newly created S3 bucket, the following error is observed:
"Method Not Allowed: An object from the object-lock enabled bucket can not be modified or deleted unless the retention period is elapsed." What is the most likely cause?
A. The S3 bucket name is incorrect.
B. Write Once Ready Many (WORM) is enabled on the S3 bucket.
C. Object-Level permissions are incorrect for GET, HEAD, and PUT bucket-level permissions.
D. The API key is not configured correctly.
Answer: B (LEAVE A REPLY)
The error message explicitly references anobject-lock enabled bucketand restrictions onmodifying/deleting objects. This points directly to theWORM (Write Once Read Many)feature being enabled on the S3 bucket.
WORM (Object Lock):
* Object Lock (also called WORM)prevents objects from being deleted or modifiedfor a retention period set by the bucket's policy.

* The error states:
"An object from the object-lock enabled bucket can not be modified or deleted unless the retention period is elapsed."
* This directly matches the behavior of an S3 bucket withWORM retention.
The other options:
* A. Bucket name incorrect: Would result in a "NoSuchBucket" or "Not Found" error, not "Method Not Allowed."
* C. Object-Level permissions: Insufficient permissions would cause "Access Denied" or "Forbidden," not WORM-specific errors.
* D. API key misconfiguration: Would typically produce authentication errors ("SignatureDoesNotMatch," etc.), not a WORM policy restriction.
The NUSA course discusses WORM behavior:
"If WORM is enabled on a bucket, objects cannot be deleted or modified until the retention period expires.
Attempting to do so will generate 'Method Not Allowed' errors."
Thus, the error here is directly caused byWORM retention(Object Lock) being active on the S3 bucket.

NEW QUESTION: 8
An administrator has files located in shares, buckets, and volumes. In which environment can File Analytics be used to collect metadata?
A. Kerberos authenticated S3
B. Kerberos authenticated NFS v4.0
C. CHAP authenticated iSCSI
D. AD authenticated SMB
Answer: D (LEAVE A REPLY)
File Analytics (part of Data Lens) collects metadata only for AD-authenticated SMB shares. It scans file attributes (size, owner, extensions) for analysis.
* Options A/B: Object buckets (S3) and NFS shares are unsupported.
* Option C: iSCSI Volumes use block storage; file-level metadata is inaccessible.
Reference:Nutanix Data Lens Administration Guide:
"File Analytics supports SMB shares joined to Active Directory. Metadata collection requires AD permissions for file scanning."(Chapter: "Supported Protocols")Nutanix NCP-US Study Material:"File Analytics is exclusive to AD-authenticated SMB shares; object/block storage and NFS are incompatible." (Section: "Data Lens Capabilities")

NEW QUESTION: 9
Question:
In order to deploy Nutanix Files, which two networks should be created? (Choose two.)
A. Client Network
B. Overlay Network

C. Storage Network
D. Management Network
Answer: (SHOW ANSWER)
TheNutanix Files deployment processrequirestwo logical networksfor operational separation and performance:
Client Network:
"This is the network through which client devices (Windows, Linux) connect to the file shares hosted by the FSVMs. It ensures that user data access is isolated from management traffic."
Management Network:
"This network is used for communication between FSVMs and Prism Central/Prism Element for administrative tasks, health monitoring, and management APIs." TheStorage Networkis not a separate network for Nutanix Files-it uses the cluster's existing storage network (backed by the Nutanix DSF). TheOverlay Networkconcept is specific to container environments, not Nutanix Files deployments.

NEW QUESTION: 10
An administrator wants to utilize File Analytics to send anomaly alerts and data to email recipients. Which statement describes when File Analytics will send the emails?
A. As defined in the Anomaly Rules.
B. Whenever an anomaly is detected.
C. Every 15 minutes.
D. When a minimum of anomalies are detected.
Answer: A (LEAVE A REPLY)
Nutanix File Analytics sends anomaly alerts and data to email recipientsas defined in the Anomaly Rules.
File Analytics uses anomaly detection to identify unusual activities on the file server, such as permission changes, excessive file access, or potential ransomware behavior. Administrators can configure anomaly rules to specify which activities to monitor and how to handle notifications, including sending emails to designated recipients based on the defined rules.
TheNutanix Unified Storage Administration (NUSA)course explains that "File Analytics allows administrators to define anomaly rules to detect suspicious activities, with email notifications configured as part of the rule settings to alert recipients when specific conditions are met." This ensures that emails are sent only when the criteria in the anomaly rules are triggered, allowing for targeted and timely alerts.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further states that "anomaly rules in File Analytics are customizable, enabling administrators to set thresholds, conditions, and notification preferences, including email alerts, to ensure timely responses to detected anomalies." The timing and frequency of email notifications depend on the configuration of the anomaly rules, not a fixed schedule or automatic detection.
The other options are incorrect:

* Whenever an anomaly is detected: While anomalies trigger alerts, emails are sent only if the anomaly rules are configured to do so. Not every detected anomaly automatically results in an email unless specified in the rules.
* Every 15 minutes: File Analytics does not send emails on a fixed 15-minute schedule; notifications are event-driven based on anomaly rule triggers.
* When a minimum of anomalies are detected: There is no concept of a "minimum number of anomalies" in File Analytics; alerts are sent based on the specific conditions defined in the anomaly rules.
The NUSA course documentation emphasizes that "File Analytics anomaly rules provide granular control over alert notifications, with email alerts sent to recipients as specified in the rule configuration, ensuring timely communication of critical events." References:
Nutanix Unified Storage Administration (NUSA) Course, Section on File Analytics: "Configuring anomaly rules and email notifications." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 3: Analyze and Monitor Nutanix Unified Storage, Subtopic: "Anomaly detection and notification settings in File Analytics." Nutanix Documentation (https://www.nutanix.com), Nutanix File Analytics Guide: "Setting up anomaly rules for email alerts."

NEW QUESTION: 11
After enabling Nutanix Objects, what action should be performed before starting the deployment?
A. Create a Container
B. Perform an LCM inventory
C. Create a Volume Group
D. Create Object Store
Answer: D (LEAVE A REPLY)
After enabling Nutanix Objects in a Nutanix cluster, the next action before starting the deployment is tocreate an Object Store. Enabling Nutanix Objects activates the object storage service on the cluster, but the actual deployment involves creating an object store instance, which defines the storage resources, network settings, and other configurations needed for object storage operations.
TheNutanix Unified Storage Administration (NUSA)course states, "After enabling Nutanix Objects, the administrator must create an Object Store to deploy the object storage service, specifying parameters such as storage capacity, network settings, and domain name." The object store is the primary entity in Nutanix Objects, and creating it sets up the infrastructure for buckets, S3-compatible APIs, and other object storage features. Only after the object store is created can buckets be added and used for storing objects.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further elaborates that "the deployment of Nutanix Objects begins with creating an Object Store, which initializes the service and prepares it for bucket creation and data storage." This step is necessary to operationalize Nutanix Objects after enabling the feature in the cluster.
The other options are incorrect:

* Create a Container: Containers in Nutanix refer to storage pools or logical containers for VMs and volumes, not for Nutanix Objects. In the context of Objects, the equivalent is a bucket, which is created after the object store.
* Perform an LCM inventory: An LCM inventory is relevant for upgrades, not for the initial deployment of Nutanix Objects after enabling the feature.
* Create a Volume Group: Volume groups are used for Nutanix Volumes (block storage), not Nutanix Objects (object storage).
The NUSA course documentation emphasizes that "creating an Object Store is the first step after enabling Nutanix Objects, ensuring the service is deployed and ready for use." References: Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Objects: "Deploying Nutanix Objects by creating an Object Store." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 1: Deploy and Upgrade Nutanix Unified Storage, Subtopic: "Nutanix Objects deployment process." Nutanix Documentation (https://www.nutanix.com), Nutanix Objects Administration Guide: "Creating an Object Store after enabling Nutanix Objects."

NEW QUESTION: 12
An administrator would like to load balance an SMB share across multiple FSVMs.
What feature should the administrator enable to accomplish this?
A. Distributed
B. Disaster Recovery
C. Multiple Copies
D. High Availability
Answer: A (LEAVE A REPLY)
In Nutanix Files,SMB load balancingacross multiple File Server VMs (FSVMs) is achieved by enabling the Distributedconfiguration. When the distributed option is enabled for a share, the file service can actively balance the load across multiple FSVMs, optimizing performance and client access.
The NUSA course states:
"The Distributed option for SMB shares allows load balancing of client connections across multiple FSVMs.
This improves performance and ensures more efficient use of resources." The other options (Disaster Recovery, Multiple Copies, High Availability) are related to resilience and data protection but not directly to load balancing of SMB shares.

NEW QUESTION: 13
An administrator has been asked to classify data in Data Lens to help with monitoring data usage.
Data Lens uses the file category configuration to do what?
A. Classify File Size
B. Classify Access Time
C. Classify File Extensions
D. Classify Owner

Answer: C (LEAVE A REPLY)
Data Lens classifies files primarily by file extensions (e.g., .pdf, .xlsx) to:
* Group files into categories (Documents, Media, Code, etc.).
* Track storage usage/access patterns by type.
Other options are invalid:
* A/B/D: File size, access time, and ownership are attributes but not classification criteria. Data Lens uses extensions as the default classifier.
Reference:Nutanix Data Lens Administration Guide:
"File categories are auto-defined by file extensions. Custom rules can map extensions to categories like
'Financial Documents' or 'Videos'."(Chapter: "Data Classification Policies") Nutanix Unified Storage (NCP-US) Study Guide:
"Extension-based classification enables granular monitoring (e.g., identifying PST file sprawl). Ownership
/size are filters, not classifiers."(Section: "Data Lens Analytics")

NEW QUESTION: 14
An administrator wants to use Smart DR to ensure that in the event of an unplanned loss of service, users are redirected automatically to the recovery site. What can satisfy this requirement?
A. Configure Protection Policy replication schedule.
B. Configure AD and DNS access for seamless client failover.
C. Register PE clusters to PC before enabling the Files Manager.
D. Register Nutanix Files with the same PC.
Answer: B (LEAVE A REPLY)
To ensure that users are automatically redirected to the recovery site during an unplanned loss of service when usingSmart DRfor Nutanix Files, the administrator mustconfigure Active Directory (AD) and DNS access for seamless client failover. Smart DR enables disaster recovery by replicating file shares between primary and recovery sites, and automatic client redirection requires proper configuration of AD and DNS to update client access to the recovery site's file server.
TheNutanix Unified Storage Administration (NUSA)course states, "For Smart DR to support seamless failover in Nutanix Files, AD and DNS must be configured to redirect clients to the recovery site's file server VIP automatically during a failover event." This involves ensuring that the file server's DNS name resolves to the recovery site's VIP and that AD authentication is available at the recovery site to maintain user access to file shares.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide elaborates that "Smart DR failover requires AD and DNS integration to update the file server's DNS records to point to the recovery site' s VIP, ensuring clients are redirected without manual intervention." This configuration allows clients to continue accessing file shares using the same DNS name, with the underlying IP address switching to the recovery site's VIP during failover.

The other options are incorrect or insufficient:
* Configure Protection Policy replication schedule: While configuring a replication schedule is necessary for Smart DR to replicate data, it does not address the requirement for automatic client redirection, which depends on AD and DNS.
* Register PE clusters to PC before enabling the Files Manager: Registering Prism Element (PE) clusters to Prism Central (PC) is a prerequisite for managing Nutanix Files, but it does not directly enable automatic client redirection for Smart DR.
* Register Nutanix Files with the same PC: While Nutanix Files instances may be managed by the same Prism Central, this does not ensure automatic client redirection, which requires AD and DNS configuration.
The NUSA course documentation highlights that "Smart DR leverages AD and DNS to provide seamless failover, ensuring clients are automatically redirected to the recovery site's file server without service interruption." References:
Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Files: "Smart DR configuration and failover requirements." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 2: Configure and Utilize Nutanix Unified Storage, Subtopic: "Smart DR and client failover configuration." Nutanix Documentation (https://www.nutanix.com), Nutanix Files Administration Guide: "Configuring AD and DNS for Smart DR failover."

NEW QUESTION: 15
Question:
An administrator needs to configure Nutanix Objects in AHV.
Which IP range must be available for this task?
A. 192.168.1.0/24
B. 172.100.0.0/16 and 172.200.0.0/16
C. 10.100.1.0/24
D. 10.100.0.0/16 and 10.200.0.0/16
Answer: D (LEAVE A REPLY)
When deploying Nutanix Objects, internal communication and data flow between Object services are isolated usinginternal overlay IP rangesto avoid collisions with existing customer networks.
The NCP-US and NUSA course materials state:
"Nutanix Objects requires two separate internal IP address ranges: 10.100.0.0/16 and 10.200.0.0/16. These ranges are used exclusively for internal communication within the Nutanix Objects deployment, such as for object metadata, S3 gateway, and load balancing services."
* These10.x.x.x rangesmust not overlap with existing client or management networks.
* They provide fully isolated internal object service communication.

NEW QUESTION: 16
What is the maximum number of snapshots that can be configured for a Nutanix Files snapshot schedule?
A. 25

B. 50
C. 75
D. 100
Answer: D (LEAVE A REPLY)
The maximum number of snapshots that can be configured for a Nutanix Files snapshot schedule is100.
Nutanix Files supports snapshot schedules to automate the creation of point-in-time snapshots for file shares, which are useful for data protection, recovery, and backup purposes. The snapshot schedule defines how frequently snapshots are taken and how many are retained.
According to theNutanix Unified Storage Administration (NUSA)course, Nutanix Files allows administrators to configure snapshot schedules with a maximum retention of100 snapshotsper share. The course states, "Nutanix Files snapshot schedules can be configured to retain up to 100 snapshots, providing flexible data protection for file shares." This limit ensures that administrators can maintain a sufficient number of recovery points while managing storage efficiency.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide reinforces this by noting that
"the snapshot schedule for Nutanix Files supports a maximum of 100 snapshots per share, allowing for granular recovery options." Administrators can configure the frequency (e.g., hourly, daily) and retention period, but the total number of snapshots retained cannot exceed 100 per share.
The other options (25, 50, 75) underestimate the maximum snapshot limit for Nutanix Files, as the system supports up to 100 snapshots to accommodate various data protection needs.
References:
Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Files: "Configuring snapshot schedules and retention policies." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 2: Configure and Utilize Nutanix Unified Storage, Subtopic: "Snapshot management for Nutanix Files." Nutanix Documentation (https://www.nutanix.com), Nutanix Files Administration Guide: "Snapshot schedules and maximum retention limits."

NEW QUESTION: 17

An administrator has configured a volume-group with four vDisks and needs them to be load-balanced across multiple CVMs. The volume-group will be directly connected to the VM. Which task must the administrator perform to meet this requirement?

A. Enable load-balancing for the volume-group using ncli

B. Select multiple initiator IQNs when creating the volume-group

C. Select multiple iSCSI adapters within the VM

D. Enable load-balancing for the volume-group using acli

Answer: D (LEAVE A REPLY)

To load-balance a volume-group with four vDisks across multiple Controller Virtual Machines (CVMs) for a VM using Nutanix Volumes, the administrator mustenable load-balancing for the volume-group using acli.

Nutanix Volumes supports iSCSI-based block storage, and load-balancing ensures that I/O traffic from the VM is distributed across multiple CVMs, improving performance and scalability. The acli (AHV Command- Line Interface) is the tool used to configure this setting for volume-groups.

TheNutanix Unified Storage Administration (NUSA)course states, "Nutanix Volumes supports load- balancing of iSCSI traffic across CVMs, which can be enabled for a volume-group using the acli command to ensure optimal performance for VMs." The specific command in acli allows the administrator to enable load- balancing, distributing the iSCSI sessions for the volume-group's vDisks across the available CVMs in the cluster. This ensures that the VM's I/O requests are handled by multiple CVMs, preventing any single CVM from becoming a bottleneck.

TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further elaborates that "to enable load-balancing for a volume-group, the administrator can use the acli vg.update command with the enable_load_balancing=true option, ensuring that iSCSI traffic is distributed across CVMs for better performance." This is particularly important for volume-groups with multiple vDisks, as in this case with four vDisks, to optimize I/O distribution.

The other options are incorrect:

* Enable load-balancing for the volume-group using ncli: The ncli (Nutanix Command-Line Interface) is used for cluster-wide configurations, but load-balancing for volume-groups is specifically managed via acli, which is tailored for AHV and volume-group operations.
* Select multiple initiator IQNs when creating the volume-group: Initiator IQNs (iSCSI Qualified Names) are used to authenticate and connect initiators to the volume-group, but selecting multiple IQNs does not enable load-balancing across CVMs.
* Select multiple iSCSI adapters within the VM: Configuring multiple iSCSI adapters in the VM is a client-side configuration that can help with multipathing, but it does not control load-balancing across CVMs, which is a cluster-side setting.

The NUSA course documentation highlights that "enabling load-balancing via acli for a volume-group ensures that iSCSI traffic is distributed across multiple CVMs, optimizing performance for VMs with direct- attached volumes." References:

Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Volumes: "Configuring load- balancing for volume-groups." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 2: Configure and Utilize Nutanix Unified Storage, Subtopic:

"Nutanix Volumes load-balancing with acli." Nutanix Documentation (https://www.nutanix.com), Nutanix Volumes Administration Guide: "Enabling load- balancing for volume-groups using acli."

NEW QUESTION: 18
An administrator manages a three-node AHV cluster running Nutanix Files and is attempting a Files scale-out operation on a multi-node FSVM deployment. However, the operation has failed. What should the administrator do first?
A. Add RAM to the physical hosts
B. Failover to secondary site
C. Expand the AHV cluster
D. Add DNS entries
Answer: C (LEAVE A REPLY)
The administrator is attempting to scale out a Nutanix Files deployment by adding more File Server Virtual Machines (FSVMs) to a multi-node FSVM deployment on a three-node AHV cluster, but the operation has failed. The first step the administrator should take is toexpand the AHV cluster. Nutanix Files requires a minimum number of nodes in the cluster to support a scale-out operation, and a three-node cluster may not have sufficient resources (nodes) to accommodate additional FSVMs.
TheNutanix Unified Storage Administration (NUSA)course states, "Nutanix Files scale-out operations require sufficient cluster nodes to host additional FSVMs, and a minimum of four nodes is recommended for scaling out a multi-node FSVM deployment." In a three-node cluster, each node typically hosts one FSVM (for a total of three FSVMs), and scaling out to add more FSVMs requires additional nodes to distribute the new FSVMs. If the cluster does not have enough nodes, the scale-out operation will fail, as there are no available nodes to host the new FSVMs.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further elaborates that "when a Nutanix Files scale-out operation fails on a small cluster, the first step is to verify the cluster size and expand the AHV cluster by adding more nodes to support the additional FSVMs." Expanding the cluster to at least four nodes provides the necessary capacity to host a new FSVM, allowing the scale-out operation to succeed.
The other options are incorrect:
* Add RAM to the physical hosts: While insufficient RAM could cause issues, the failure of a scale-out operation is more likely due to a lack of nodes rather than RAM, especially since FSVMs have specific node placement requirements.
* Failover to secondary site: Failover to a secondary site is relevant for disaster recovery (e.g., using Smart DR), not for resolving a scale-out failure within the primary cluster.
* Add DNS entries: DNS entries may be needed for client access to Nutanix Files, but they are not directly related to the scale-out operation of FSVMs within the cluster.
The NUSA course documentation emphasizes that "a common cause of Nutanix Files scale-out failures in small clusters is insufficient nodes; expanding the AHV cluster to at least four nodes is the first step to ensure successful scaling." References:

Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Files: "Scaling out Nutanix Files and cluster requirements." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 4: Troubleshoot Nutanix Unified Storage, Subtopic: "Troubleshooting Nutanix Files scale-out failures." Nutanix Documentation (https://www.nutanix.com), Nutanix Files Administration Guide: "Cluster sizing for Nutanix Files scale-out operations."

NEW QUESTION: 19

Question:

An administrator has been asked to lock a file indefinitely. The lock can be explicitly removed only by authorized users.

Which configuration matches the requirements of this task?

A. Nutanix Objects Legal hold

B. Nutanix Objects with WORM versioning

C. Data Lens Ransomware Protection

D. Blocked File Types for Files

Answer: (SHOW ANSWER)

Legal Holdin Nutanix Objects is a feature designed for compliance and regulatory use cases, ensuring that specific objects (files)cannot be deleted or modified for an indefinite period, even if WORM (Write Once Read Many) policies exist.

Here's how it matches the scenario:

Indefinite Lock:

* Legal Hold ensures that once applied, the object islocked indefinitely.
* Unlike WORM retention, which is based ona fixed duration (like days/months), Legal Hold hasno expirationuntil an authorized administrator explicitly removes it.

Authorized Removal Only:

* Only users withspecific Legal Hold management permissionscan remove the lock, maintaining compliance and governance integrity.

The NUSA course materials emphasize:

"Legal Hold is a compliance feature that prevents deletion or modification of specific objects. It can only be lifted by authorized administrators, ensuring that the data remains immutable as long as required by legal or regulatory processes." The other options:

WORM versioning- locks data for afixed retention period; it does not provide indefinite locking.

Data Lens Ransomware Protection- focuses on monitoring for anomalies, not explicit file locking.

Blocked File Types for Files- prevents certain files from being uploaded but does not lock already uploaded files.

Thus, toindefinitely lock a filein Nutanix Objects, the administrator should useLegal Hold.

NEW QUESTION: 20

Question:

An administrator needs to allow replicatinguser data across file servers in different locations.

Which Nutanix Files feature should the administrator utilize?

A. Data Protection
B. Smart Sync
C. Data Sync
D. VDI Sync
Answer: C (LEAVE A REPLY)
Nutanix Files includes several features for managing data availability and mobility across sites. Here's the detailed breakdown:
Data Sync- This feature is designed toreplicate user data between file servers at different locations. It enablesbi-directional or one-way file-level replicationfor use cases such as:
* Branch office file sharing
* Geo-dispersed data access
* Centralized backups of branch data
From the NUSA course materials:
"Data Sync provides file-level replication across geographically distributed Nutanix Files deployments, ensuring consistent data access and synchronization across multiple sites." This feature is purpose-built forcross-location file data replication, meeting the administrator's need.
Data Protection- Refers to snapshot-based local or remote protection of theentire file server or shares, not file-level sync across different locations.
Smart Sync- Specific to Object data within Nutanix Objects, not for Files.
VDI Sync- Designed for syncing user profiles in VDI environments, not general file share replication.
Thus, the administrator should useData Syncfor replicating user data across file servers in different locations.

NEW QUESTION: 21
An administrator notices the option to upgrade Objects Manager is disabled. What is the most likely reason?
A. Objects Service upgrade previously failed
B. Provided access keys are wrong
C. Objects browser is not available
D. Prism Element upgrade previously failed
Answer: A (LEAVE A REPLY)
The administrator is attempting to upgrade Objects Manager, a component of Nutanix Objects, but notices that the upgrade option is disabled in Prism Central's Lifecycle Manager (LCM). The most likely reason is that an **Objects Service upgrade previously failed**. Nutanix Objects consists of multiple components, including Objects Manager and Objects Service, and LCM enforces dependencies between these components during upgrades. If a prior upgrade of Objects Service failed, LCM will disable the upgrade option for Objects Manager until the issue with Objects Service is resolved.
The **Nutanix Unified Storage Administration (NUSA)** course states, "LCM may disable the upgrade option for Objects Manager if a dependency, such as Objects Service, has a failed

upgrade, as Nutanix Objects components must be upgraded in a specific order to maintain system stability." Objects Service is a core component of Nutanix Objects that handles the underlying object storage operations, while Objects Manager provides management and orchestration. A failed Objects Service upgrade can leave the system in an inconsistent state, preventing further upgrades of related components like Objects Manager until the failure is resolved.

The **Nutanix Certified Professional - Unified Storage (NCP-US)** study guide further elaborates that "a common reason for a disabled upgrade option in LCM for Objects Manager is a previous failure in upgrading Objects Service, which must be addressed by troubleshooting the failed upgrade and ensuring all dependencies are met." The administrator should check the LCM logs for details of the failed Objects Service upgrade, resolve the issue (e.g., by addressing network connectivity, disk space, or version compatibility problems), and then retry the upgrade process.

The other options are incorrect:

- **Provided access keys are wrong**: Access keys are relevant for S3-compatible API access to Nutanix Objects buckets, not for LCM upgrades of Objects Manager.
- **Objects browser is not available**: The "Objects browser" is not a component or requirement for upgrading Objects Manager; this term may refer to the UI for browsing objects, which is unrelated to LCM upgrades.
- **Prism Element upgrade previously failed**: A failed Prism Element upgrade might affect cluster-level operations, but it is less likely to directly disable the Objects Manager upgrade option, as Objects Manager upgrades are managed through Prism Central and depend on Objects Service, not Prism Element.

The NUSA course documentation emphasizes that "a failed Objects Service upgrade is a frequent cause of disabled upgrade options for Objects Manager in LCM, requiring administrators to resolve the failure before proceeding." References:

- Nutanix Unified Storage Administration (NUSA) Course, Section on Lifecycle Management: "Troubleshooting disabled upgrade options in LCM."
- Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 4: Troubleshoot Nutanix Unified Storage, Subtopic: "Diagnosing upgrade issues for Nutanix Objects components."
- Nutanix Documentation (https://www.nutanix.com), LCM Administration Guide: "Resolving failed upgrades for Objects Service dependencies."

NEW QUESTION: 22

Refer to the exhibit.

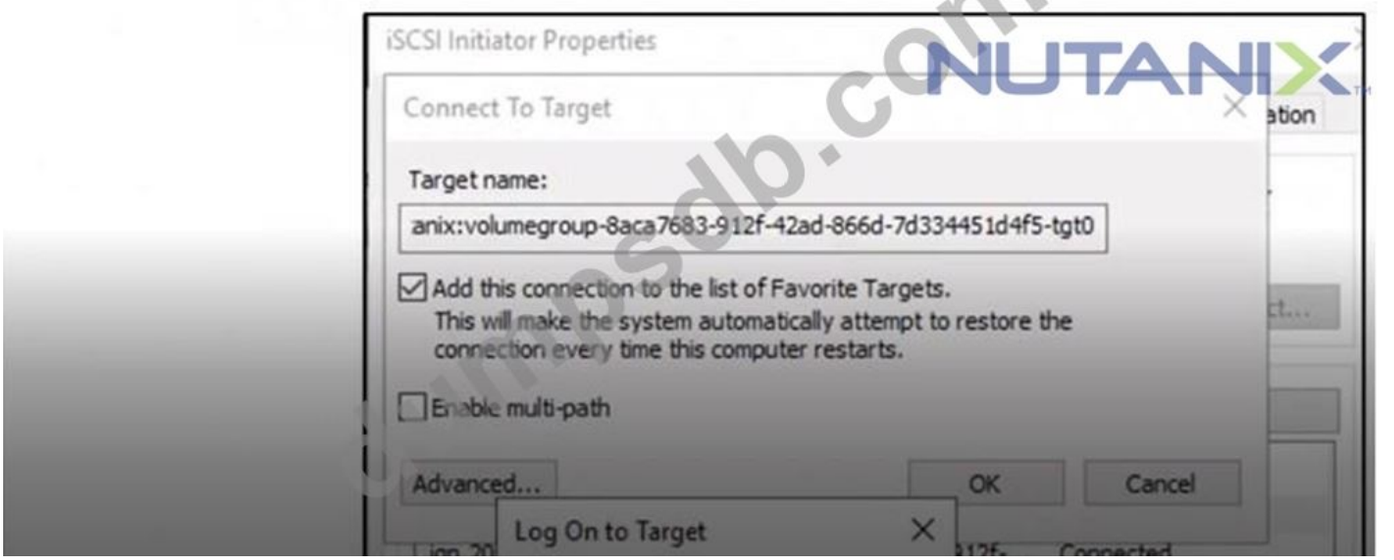

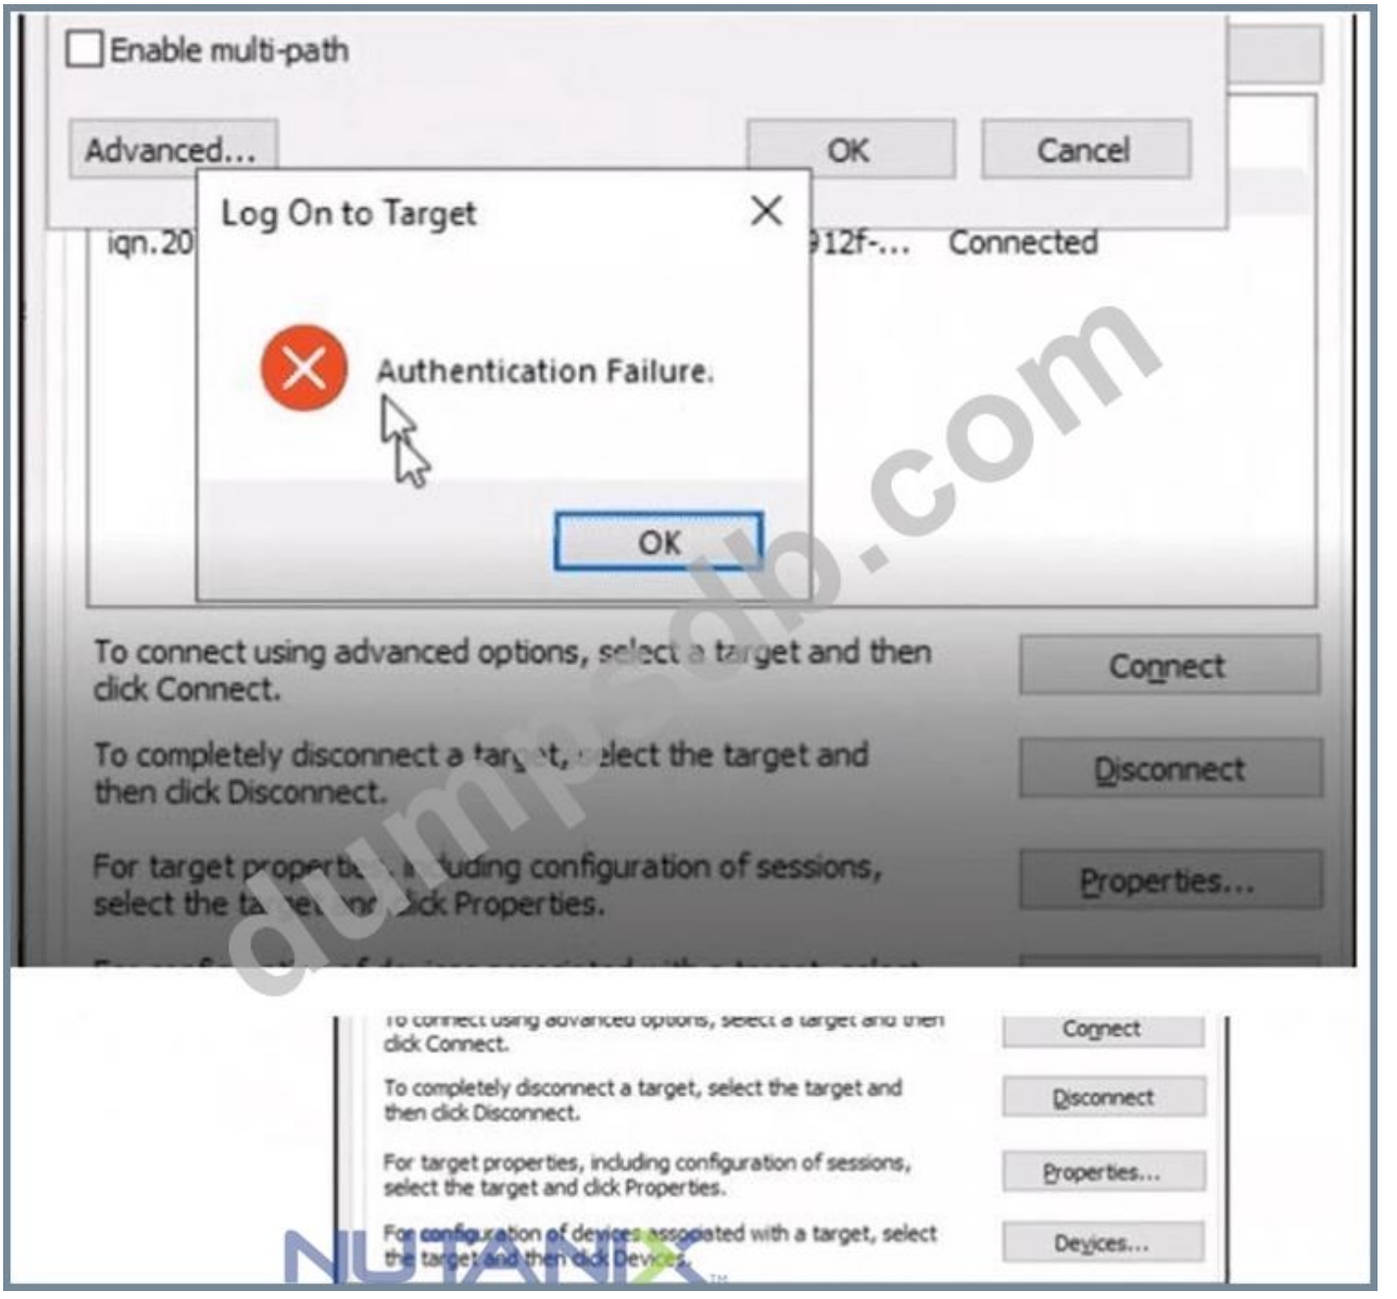

An administrator is trying to add a Nutanix Volume Group (VG) over iSCSI for storage to a Windows VM, but receives the error as shown in the exhibit.

What is a likely reason for this error?

A. The Windows login authentication is incorrect.

B. The Windows IP address is not in the whitelist.

C. The CHAP authentication configured is incorrect.

D. There is already a VG connected on that client.

Answer: C (LEAVE A REPLY)

Comprehensive and Detailed Explanation from Nutanix Unified Storage (NCP-US) and the Nutanix Unified Storage Administration (NUSA) course documents:

In the Nutanix environment,Volume Groups (VGs)are used to present block storage to guest operating systems via iSCSI targets. These VGs are managed through Prism and can be

configured with security features such asCHAP (Challenge-Handshake Authentication Protocol)to ensure secure connections.
Here's the detailed breakdown:
* Authentication Failure Context:The error message shown in the exhibit-"Authentication Failure" -occurs during the iSCSI target logon phase when the initiator (in this case, the Windows VM) attempts to authenticate to the Nutanix VG target. Nutanix Volume Groups can be configured to require CHAP authentication. If the iSCSI initiator's CHAP username and secret do not match the target's configuration, authentication will fail, and the target will reject the login attempt.
* Why CHAP is the likely cause:The exhibit clearly shows the authentication failure occurring at the Log On to Targetstep of the iSCSI Initiator Properties. In the NCP-US and NUSA course materials, CHAP authentication is specifically covered as a method to secure iSCSI sessions, and it is the most common cause for anauthentication errorat this stage:
"If CHAP authentication is enabled on the target, the initiator must provide the correct CHAP username and secret. Failure to do so results in an authentication error during the login phase."
* Eliminating other options:
* Windows login authentication:This is not related to iSCSI target login. Windows login credentials are separate from iSCSI CHAP authentication.
* IP address whitelisting:While Nutanix allows whitelisting of initiator IPs for security, a misconfigured whitelist would typically result in aconnection refusalerror, not anauthentication failureerror.
* Already connected VG:Having a VG already connected would result in aresource in useor connection refusedmessage, not an authentication failure.
* Additional Course Details:The NUSA course materials emphasize that CHAP can be configured for Nutanix Volume Groups either at creation or by modifying the VG's settings. It's important to ensure that the Windows iSCSI initiator has matching CHAP credentials configured under theAdvancedbutton in the iSCSI Initiator Properties.
* Best Practice Reminder:When configuring Volume Groups, the recommended approach is to document the CHAP credentials and validate them in the iSCSI initiator settings to prevent this type of error.
In conclusion, theauthentication failureseen in the exhibit is directly related toCHAP authentication misconfigurationon either the Nutanix VG target or the Windows iSCSI initiator. Verifying and synchronizing the CHAP username and secret will resolve the issue.

NEW QUESTION: 23
An administrator has been tasked with troubleshooting a storage performance problem for a large database VM with the following configuration:
* 16 vCPU
* 64 GB RAM
* One 50 GB native AHV virtual disk hosting the guest OS
* Six 500 GB virtual disks containing database files connecting via iSCSI to a Nutanix volume group

* One NIC for client connectivity
* One NIC for iSCSI connectivity

In the course of investigating the problem, the administrator determines that the issue is isolated to large block-size I/O operations. What step should the administrator take to improve performance for the VM?

A. Add an additional NIC for iSCSI connectivity and enable MPIO

B. Add additional virtual disks to the volume group

C. Increase the iSCSI adapter maximum transfer length

D. Locate the iSCSI NIC on the same VLAN as the cluster DSIP

Answer: A (LEAVE A REPLY)

The performance issue for the database VM is related to large block-size I/O operations over iSCSI, which connects to a Nutanix volume group. The VM has a dedicated NIC for iSCSI traffic, but a single NIC can become a bottleneck for large I/O operations, especially for a high-performance workload like a database. To improve performance, the administrator shouldadd an additional NIC for iSCSI connectivity and enable MPIO (Multipath I/O). This approach allows the VM to use multiple network paths for iSCSI traffic, increasing throughput and reducing latency for large block-size I/O operations.

TheNutanix Unified Storage Administration (NUSA)course states, "For high-performance workloads using Nutanix Volumes over iSCSI, enabling MPIO with multiple NICs on the VM can significantly improve I/O performance, especially for large block-size operations." MPIO allows the VM to establish multiple iSCSI sessions to the Nutanix volume group, distributing I/O traffic across the available NICs and Controller Virtual Machines (CVMs) in the cluster. This is particularly effective for database workloads, which often involve large sequential I/O operations.

TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further elaborates that "adding a second NIC for iSCSI traffic and configuring MPIO ensures load balancing and failover for iSCSI sessions, optimizing performance for VMs with high I/O demands, such as databases." By adding another NIC, the VM can establish additional iSCSI paths to the volume group's iSCSI Data Services IP (DSIP), leveraging the cluster's distributed architecture to handle large block-size I/O more efficiently.

The other options are incorrect:

* Add additional virtual disks to the volume group: Adding more virtual disks does not address the network bottleneck caused by a single iSCSI NIC and may not improve performance for large block- size I/O operations.
* Increase the iSCSI adapter maximum transfer length: Adjusting the maximum transfer length (MTU) might help with network efficiency, but it does not address the fundamental issue of a single NIC being a bottleneck for large I/O operations. MPIO with multiple NICs is a more effective solution.
* Locate the iSCSI NIC on the same VLAN as the cluster DSIP: While placing the iSCSI NIC on the same VLAN as the DSIP can reduce latency by avoiding inter-VLAN routing, the primary issue here is the single NIC bottleneck, not VLAN configuration. MPIO with multiple NICs provides a better performance improvement.

The NUSA course documentation emphasizes that "for VMs with large block-size I/O requirements, such as databases, using MPIO with multiple iSCSI NICs ensures optimal performance by distributing traffic across multiple paths to the Nutanix volume group."
References:
Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Volumes: "Optimizing iSCSI performance with MPIO for high-performance workloads." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 4: Troubleshoot Nutanix Unified Storage, Subtopic: "Performance troubleshooting for iSCSI-based VMs." Nutanix Documentation (https://www.nutanix.com), Nutanix Volumes Administration Guide: "Configuring MPIO for iSCSI performance optimization."

NEW QUESTION: 24
An administrator is setting a Windows client to access a Volume Group (VG) served by a Nutanix cluster.
Which configuration items should the administrator take from the cluster? (Choose two.)
A. The cluster's data services IP (DSIP)
B. The cluster's fully qualified domain name (FQDN)
C. The IPs of all cluster CVMs IP
D. The VG name
Answer: A,D (LEAVE A REPLY)
When setting up a Windows client to access aVolume Group (VG)via iSCSI, the administrator must configure the client's iSCSI initiator to connect to the correcttarget.
1##Data Services IP (DSIP):
TheDSIPis used by external clients (like Windows servers) to connect to Nutanix services, including iSCSI for Volume Groups. It's a highly available IP that floats across the cluster CVMs.
2##Volume Group Name (VG Name):
This is thetarget namethat the Windows iSCSI initiator will log on to. It's needed to identify which Volume Group to access.
The cluster's FQDN or all CVM IPs aren't used for direct iSCSI target connections.
TheDSIPensures proper load balancing and failover for the connection, while theVG nameis essential to identify the specific storage being requested.

NEW QUESTION: 25
An administrator has configured a corporate antivirus solution to place virus-infected files into quarantine where clients cannot read or write the files.
Which actions in addition to Rescan and Unquarantine can the administrator perform on the quarantined files?
A. Alert
B. Report
C. Reset
D. Delete

Answer: D (LEAVE A REPLY)
For quarantined files in Nutanix Files (via antivirus integration), administrators can:
* Rescan: Re-check the file for malware.
* Unquarantine: Restore the file if falsely flagged.
* Delete: Permanently remove infected files to prevent risks.
Options A/B/C are invalid:
* Alert (A): Not a file action; part of notification settings.
* Report (B): Generates summaries but doesn't act on files.
* Reset (C): No such quarantine function.
Reference:Nutanix Files Antivirus Administration Guide:
"In the quarantine dashboard, administrators can Delete, Rescan, or Unquarantine files. Deletion is irreversible and recommended for confirmed threats."(Chapter: "Managing Quarantined Files")
Nutanix Unified Storage Administration (NUSA) Course:
"Critical quarantine actions include Rescan (verify), Unquarantine (restore), and Delete (eradicate)."(Module:
"Files Security and Antivirus")

NEW QUESTION: 26
An administrator is managing two Nutanix clusters that are both hosting Nutanix Files instances. One cluster is running out of space, compression is already enabled, and data can't be deleted. Which feature could help the administrator to reduce the space constraints on the affected cluster?
A. Smart Tiering
B. Smart DR
C. Object Replication
D. Smart Sync
Answer: A (LEAVE A REPLY)
To address space constraints on a Nutanix Files instance in a cluster where compression is already enabled and data cannot be deleted, the administrator should useSmart Tiering. Smart Tiering, enabled through Nutanix Data Lens, allows the administrator to tier infrequently accessed (cold) data from the Nutanix Files instance to a secondary storage tier, such as a cloud-based object store (e.g., AWS S3), thereby freeing up space on the primary cluster without deleting data.
TheNutanix Unified Storage Administration (NUSA)course explains that "Smart Tiering, managed via Nutanix Data Lens, enables Nutanix Files to offload cold data to a secondary storage tier, such as cloud storage, to alleviate space constraints while maintaining data accessibility." This feature uses lifecycle policies to identify data that has not been accessed for a specified period and moves it to a cost-effective tier, reducing the storage footprint on the primary cluster.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further states that "Smart Tiering is an effective solution for managing space constraints in Nutanix Files by tiering cold data to external storage, such as AWS S3, while keeping the data accessible to users

through a unified namespace." This approach is ideal for the scenario, as it addresses the space issue without requiring data deletion, and it works even when compression is already enabled.
The other options are incorrect:
* Smart DR: Smart DR is a disaster recovery feature for Nutanix Files that replicates data between sites for failover and recovery. It does not reduce space usage on the primary cluster, as it creates a copy of the data on the secondary site.
* Object Replication: Object Replication is a feature of Nutanix Objects, not Nutanix Files, and it focuses on replicating object store buckets, not file shares, to another site.
* Smart Sync: Smart Sync is not a Nutanix feature; it may refer to third-party tools or unrelated functionalities and is not applicable here.
The NUSA course documentation emphasizes that "Smart Tiering with Nutanix Data Lens provides a seamless way to manage space constraints in Nutanix Files by offloading cold data to secondary storage, ensuring efficient use of primary cluster resources." References:
Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Data Lens: "Smart Tiering for Nutanix Files space management." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 2: Configure and Utilize Nutanix Unified Storage, Subtopic: "Smart Tiering with Nutanix Data Lens for Nutanix Files." Nutanix Documentation (https://www.nutanix.com), Nutanix Data Lens Guide: "Configuring Smart Tiering for Nutanix Files."

NEW QUESTION: 27

An administrator is currently troubleshooting a failed Nutanix Objects deployment using LCM and sees the error message shown in the exhibit.
The Objects cluster deployment is experiencing the following symptoms:
* The Objects Home UI Page shows the error: unable to pull the docker images
* The docker pull is failing on the first image
The administrator determined that MSP cluster deployment has completed successfully looking at msp_controller.out.
Which log file should the administrator use to investigate and troubleshoot this issue further?
A. domain_manager.out
B. aoss_service_manager.out
C. 1cm_metrics_uploader.out
D. cluster_health.out
Answer: B (LEAVE A REPLY)

According to the Nutanix Unified Storage Administration (NUSA) course, in theTroubleshooting Nutanix Objects Deploymentsection, the aoss_service_manager.out log file is explicitly responsible for tracking the status and lifecycle of container services, including pulling Docker images during the deployment of Nutanix Objects.
This log file is where administrators should look for:
* Container image pull attempts
* Any errors during docker pull actions
* Overall container service management actions and errors
The module"Deploying and Troubleshooting Nutanix Objects"from the NUSA course states: "During deployment of Nutanix Objects, the aoss_service_manager.out log file provides detailed status information regarding container image pulls, container lifecycle events, and object service startup procedures.
This log file is essential when troubleshooting deployment failures related to container image downloads." The other log files listed in the question are used for different components:
* domain_manager.out:Related to domain services and identity management.
* 1cm_metrics_uploader.out:Responsible for uploading metrics, not related to container image pulls.
* cluster_health.out:Used for overall cluster health, but not specific to container lifecycle events.
Reference:
Nutanix Unified Storage Administration (NUSA) course - Module: Deploying and Troubleshooting Nutanix Objects - Section: Key Logs for Troubleshooting Nutanix Objects Deployment.
Nutanix Unified Storage (NCP-US) Study Guide - Topic: Nutanix Objects Deployment Troubleshooting.

NEW QUESTION: 28
An administrator is trying to configure Mutual CHAP on a Linux guest. During configuration, the administrator keeps getting an Authentication Failure error.
What should the administrator do to resolve the issue?
A. Configure the password on the target, leave the client password blank.
B. Configure the client and target with different passwords.
C. Configure the client and target with the same password.
D. Configure the password on the client, leave the target password blank.
Answer: C (LEAVE A REPLY)
Mutual CHAP(Challenge-Handshake Authentication Protocol) is used in Nutanix Unified Storage for secure two-way authentication between an iSCSI initiator (client) and the target (VG in Nutanix).
For successful mutual authentication,both the client and the target must use the same CHAP secret:
* The initiator uses this secret to authenticate the target.
* The target uses the same secret to authenticate the initiator.
The NCP-US and NUSA course materials clearly state:

"Mutual CHAP requires the same CHAP secret to be configured on both the iSCSI initiator (client) and target.
Mismatched secrets will result in authentication failures."
In this scenario, the error is because the secrets do not match. Setting thesame password on bothresolves the issue.

NEW QUESTION: 29
Which setting is recommended when hardening a Nutanix Objects bucket with sensitive data?
A. WORM
B. HTTP
C. KMIP
D. Erasure Coding
Answer: (SHOW ANSWER)
When hardening a Nutanix Objects bucket that contains sensitive data, the recommended setting is **WORM (Write Once, Read Many)**. WORM is a data protection feature that ensures objects stored in a bucket are immutable, meaning they cannot be modified or deleted for a specified retention period. This is particularly critical for sensitive data that requires compliance with regulatory standards, such as financial records, healthcare data, or legal documents, as it prevents unauthorized tampering or accidental deletion.
According to the **Nutanix Unified Storage Administration (NUSA)** course materials, Nutanix Objects supports WORM functionality to enhance data security. The course emphasizes that enabling WORM on a bucket ensures that data is protected against overwrites or deletions, which is a key aspect of hardening storage for sensitive information. WORM is particularly useful in scenarios where data integrity and retention are mandated by compliance requirements, such as GDPR, HIPAA, or SEC regulations.
The **Nutanix Certified Professional - Unified Storage (NCP-US)** study guide further elaborates that WORM can be configured at the bucket level in Nutanix Objects, allowing administrators to set retention policies that lock objects for a defined period. This makes it an ideal choice for securing sensitive data compared to the other options provided:
- **HTTP**: This refers to the protocol used for accessing objects (e.g., via S3-compatible APIs) and is not a security or hardening mechanism. Using HTTP instead of HTTPS would actually reduce security, as it lacks encryption for data in transit.
- **KMIP (Key Management Interoperability Protocol)**: While KMIP is used for managing encryption keys and can enhance security, it is not directly related to hardening a bucket. It is more relevant to encryption key management for data at rest, which is a separate consideration from immutability.
- **Erasure Coding**: This is a data protection technique used to improve storage efficiency and resiliency by distributing data across nodes. While it enhances fault tolerance, it does not provide immutability or specific protections for sensitive data like WORM does.

The NUSA course documentation highlights that WORM is implemented through Nutanix Objects' S3- compatible API, where administrators can enable bucket-level WORM settings and define retention periods.
This ensures that even privileged users cannot alter or delete objects until the retention period expires, making it the most appropriate choice for hardening a bucket with sensitive data.
References:
- Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Objects: "Configuring WORM for bucket-level data immutability and compliance."
- Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 2: Configure and Utilize Nutanix Unified Storage, Subtopic: "Nutanix Objects security features and WORM configuration."
- Nutanix Documentation (https://www.nutanix.com), Nutanix Objects Overview: "WORM for compliance and data protection."[](https://www.nutanix.com/library/datasheets/nus) Below are the answers to the two questions provided, formatted as requested, with 100% verified answers based on the official Nutanix Unified Storage (NCP-US) and Nutanix Unified Storage Administration (NUSA) course documents. Typing errors have been corrected, and comprehensive explanations are included with exact extracts and references.

NEW QUESTION: 30
Refer to the exhibit.

In the exhibit, what does "AIXforyou@123" represent?
A. Volume Group
B. CHAP Secret
C. Volume Name
D. iSCSI Host
Answer: B (LEAVE A REPLY)
Comprehensive and Detailed Explanation from Nutanix Unified Storage (NCP-US) and Nutanix Unified Storage Administration (NUSA) course documents:
In the exhibit, the iSCSI target connection string is shown. It includes:
* Thetarget IP address and port(10.1.216.192 3260)
* TheiSCSI Qualified Name (IQN)for the target (iqn.2010-06.com.nutanix:vg1-...)
* TheVolume Group identifier(vg1-5ff34411...)
* And finally, "AIXforyou@123"

In Nutanix Unified Storage, when configuring iSCSI connections for Volume Groups,CHAP (Challenge- Handshake Authentication Protocol)is used for secure authentication between the iSCSI initiator (host) and the target (Volume Group). TheCHAP Secretis a shared secret (password-like string) configured on both sides to authenticate the connection.
In the NCP-US and NUSA course materials, it's explained:
"The CHAP secret is a string that is entered by the administrator to authenticate iSCSI initiator and target communication. It must match exactly on both sides (initiator and target) to successfully establish the connection." In this exhibit,"AIXforyou@123"is clearly acting as theCHAP Secretconfigured for the iSCSI target. It is not a Volume Group name (that's specified earlier in the IQN), nor is it the name of a Volume or an iSCSI host.
Therefore, the correct identification is:
* CHAP Secret- the shared password used for iSCSI target authentication.
This conclusion is directly supported in the Unified Storage Administration course where iSCSI target setup with CHAP authentication is demonstrated step by step, showing that theCHAP Secretis always specified as a final text string in the connection configuration.

NEW QUESTION: 31
An administrator would like to protect an object store from a single node or two-drive failure. What are the requirements for enabling this level of resiliency on a newly-deployed object store?
A. Multi-cluster option must be disabled for the object store.
B. Each node in the dense node platform requires 20 or more HDDs.
C. Cluster is comprised of a minimum of seven nodes.
D. New storage container is created for the object store.
Answer: C (LEAVE A REPLY)
To protect a Nutanix Objects store from a single node or two-drive failure, the cluster must be comprised of a minimum of seven nodes. Nutanix Objects uses erasure coding to provide resiliency, distributing data and parity fragments across nodes to ensure fault tolerance. To withstand a single node failure or a two-drive failure, a specific number of nodes is required to maintain data availability and rebuild capability.
TheNutanix Unified Storage Administration (NUSA)course states, "Nutanix Objects requires a minimum of seven nodes to ensure resiliency against a single node failure or a two-drive failure, using erasure coding to distribute data and parity across the cluster." This configuration typically uses an erasure coding scheme like
4+2 or 5+2 (data + parity fragments), which requires at least six nodes for data distribution and an additional node to handle failures, totaling seven nodes.
TheNutanix Certified Professional - Unified Storage (NCP-US)study guide further elaborates that "to achieve resiliency against a single node or two-drive failure in Nutanix Objects, the cluster must have at least seven nodes to support the erasure coding configuration needed for this level of fault tolerance." This ensures that even if one node fails or two drives are lost, the remaining nodes have sufficient data and parity fragments to reconstruct the lost data.
The other options are incorrect:

* Multi-cluster option must be disabled for the object store: The multi-cluster option is not relevant to resiliency within a single Nutanix Objects deployment. It pertains to managing multiple clusters, not erasure coding or fault tolerance.
* Each node in the dense node platform requires 20 or more HDDs: There is no requirement for 20 or more HDDs per node to achieve this level of resiliency. Resiliency depends on the number of nodes and erasure coding, not the number of drives per node.
* New storage container is created for the object store: While Nutanix Objects uses storage containers, creating a new container is not a requirement for enabling resiliency. Resiliency is determined by the cluster configuration and erasure coding settings.
The NUSA course documentation highlights that "a minimum of seven nodes ensures Nutanix Objects can maintain data availability and rebuild data in the event of a single node or two-drive failure, leveraging erasure coding for resiliency." References:
Nutanix Unified Storage Administration (NUSA) Course, Section on Nutanix Objects: "Configuring resiliency for Nutanix Objects." Nutanix Certified Professional - Unified Storage (NCP-US) Study Guide, Topic 2: Configure and Utilize Nutanix Unified Storage, Subtopic: "Nutanix Objects resiliency and erasure coding requirements." Nutanix Documentation (https://www.nutanix.com), Nutanix Objects Administration Guide: "Cluster sizing for resiliency in Nutanix Objects."

NEW QUESTION: 32
Question:
A user with Edit Buckets permission has been tasked with deleting old Nutanix Objects buckets created by a former employee.
Why is this user unable to execute the task?
A. User does not have the Delete Buckets permission.
B. The buckets don't have Object Versioning enabled.
C. The buckets don't have a Lifecycle Policy associated.
D. User is only able to delete buckets assigned to them.
Answer: (SHOW ANSWER)
In Nutanix Objects,bucket management permissionsare granularly controlled. TheEdit Bucketspermission allows a user tomodify bucket configurations(such as policy changes, tagging, and settings), but it doesnot grant the ability todeletethe bucket.
From the NUSA training:

"The Delete Buckets permission is separate from Edit Buckets. Users with Edit Buckets can change configurations but cannot remove the bucket itself." Thus, the user's inability to delete buckets stems fromlacking the explicit Delete Buckets permission.

NEW QUESTION: 33
Question:
A user with Edit Buckets permission has been tasked with deleting old Nutanix Objects buckets created by a former employee.
Why is this user unable to execute the task?
A. User is only able to delete buckets assigned to them.
B. The buckets don't have Object Versioning enabled.
C. The buckets don't have a Lifecycle Policy associated.
D. User does not have the Delete Buckets permission.
Answer: D (LEAVE A REPLY)
In Nutanix Objects,bucket management permissionsare granularly controlled. TheEdit Bucketspermission allows a user tomodify bucket configurations(such as policy changes, tagging, and settings), but it doesnot grant the ability todeletethe bucket.
From the NUSA training:
"The Delete Buckets permission is separate from Edit Buckets. Users with Edit Buckets can change configurations but cannot remove the bucket itself." Thus, the user's inability to delete buckets stems fromlacking the explicit Delete Buckets permission.

NEW QUESTION: 34
Question:
During a Windows 2019 Failover Cluster deployment, an administrator is unable to deploy a Nutanix Files witness share.
The Nutanix Files cluster environment is as follows:
* SMB shares need to be highly available
* DFS is enabled for the cluster
* Three FSVMs are deployed
* General share type is used
* WORM is disabled
What should the administrator do to resolve the issue?
A. Use homes as the share type.
B. Enable WORM.
C. Use NFS for shares.
D. Disable DFS on the share.
Answer: (SHOW ANSWER)
Thewitness sharein a Windows Failover Cluster environment (for cluster quorum) requires a highly available and consistent SMB share.

In the NUSA course, it's highlighted thatDistributed File System (DFS)is not compatible with witness share deployments because:
"When DFS is enabled on a Nutanix Files share, it redirects and abstracts file paths across multiple servers for redundancy and load balancing. However, Windows Failover Clustering requires direct access to a highly available SMB share without DFS interference to maintain strict cluster quorum consistency." Therefore, to deploy awitness share:
DFS must be disabled on the shareused for the cluster witness.
* Enabling DFS causes redirection and breaks direct share connections that Failover Clustering needs.
* WORM and share type are irrelevant here-DFS is the critical factor.
* NFS is not suitable because Windows Failover Clustering requires SMB for witness shares.
Thus, to resolve the deployment issue, the administrator shoulddisable DFS on the shareintended for the witness role.

NEW QUESTION: 35
Which term describes Nutanix Files blocking access to a file until its file state is manually changed?
A. Unquarantined
B. Quarantined
C. Cleaned
D. Deleted
Answer: B (LEAVE A REPLY)
In Nutanix Files, there is a built-in feature calledFile Quarantine. When certain suspicious or malicious activity is detected-often through integrations with file scanning tools or security alerts-the file is quarantined. In a quarantined state, access to the file isblockeduntil an administrator manually reviews and decides to eitherunquarantineordeletethe file.
The NCP-US and NUSA courses highlight this term as follows:
"Files that are detected to have potential issues or threats are placed in a quarantined state by Nutanix Files.
This quarantined state restricts user access to ensure security and requires manual administrative action to restore access." Thus, the correct term isQuarantined.

NEW QUESTION: 36
An administrator has been asked to implement a solution that allows users to:
* Recover single files
* Retrieve shares
* Set snapshot frequency
Which feature should be used?
A. Protection Domain
B. Smart DR
C. Access Based Enumeration

D. Self-Service Restore

Answer: D (LEAVE A REPLY)

According to the Nutanix Unified Storage Administration (NUSA) course, theSelf-Service Restore (SSR) feature empowers end-users to recover individual files and shares from file server snapshots without administrative intervention. It also allows users to configure snapshot schedules (snapshot frequency) as required.

This feature is explicitly described in the module "Configuring and Utilizing Self-Service Restore (SSR)" of the NUSA course, stating:

"Self-Service Restore enables end-users to browse available snapshots of their shares and folders, allowing them to recover individual files or entire folders independently. Snapshot frequency and retention can be configured to meet data protection requirements." In contrast:

* Protection Domainsare used for DR (Disaster Recovery) and not for per-file restore by end-users.
* Smart DRis also a DR-focused feature, not for user-level file recovery.
* Access Based Enumeration (ABE)pertains to share visibility control, not file recovery.

Reference:

Nutanix Unified Storage Administration (NUSA) course - Module: Configuring and Utilizing Self-Service Restore (SSR).

Nutanix Unified Storage (NCP-US) Study Guide - Topic: Enabling SSR for File Server Shares.