Nutanix.NCP-BC-7.5.v2026-07-20.q52

Exam Code:	NCP-BC-7.5
Exam Name:	Nutanix Certified Professional - Business Continuity (NCP-BC) 7.5
Certification Provider:	Nutanix
Free Question Number:	52
Version:	v2026-07-20
# of views:	105
# of Questions views:	520
https://www.freecram.net/torrent/Nutanix.NCP-BC-7.5.v2026-07-20.q52.html	

NEW QUESTION: 1
A company is recovering a 3-tier application (Database, Middleware, Web). The administrator executes the Recovery Plan. The Web tier VMs start up but applications immediately crash because they cannot connect to the Database tier. Which setting in the Recovery Plan should the administrator have configured to prevent this execution failure?
A. Optional inter-stage delays
B. AppType and AppTier category assignments
C. Protection Policy mapping
D. Automatic Volume Group attachments
Answer: (SHOW ANSWER)
Orchestrating the recovery of multi-tier applications requires more than just powering on virtual machines; it requires managing the dependencies between application layers. In a standard 3-tier architecture, the Middleware and Web tiers typically depend on the Database being fully operational and ready to accept connections before they can successfully initialize. If a Recovery Plan powers on all tiers too quickly or simultaneously, the " upstream " services will fail to find their database backend, leading to the application crashes described in the scenario.
The Nutanix Recovery Plan provides " execution stages " to handle this boot ordering. However, even if tiers are in separate stages, a VM reaching a " Power On " state in the hypervisor does not mean the application services (like SQL or Oracle) are fully booted and ready for traffic. To address this, the administrator should configure " optional inter-stage delays " . This setting allows the administrator to define a specific wait time (in seconds) between the completion of one stage (e.g., Database power-on) and the start of the next (e.g., Middleware power-on). This delay gives the database engine enough time to finish its internal checks, mount its databases, and open its network ports. By incorporating these surgical delays, the administrator ensures a reliable, automated recovery where each tier only attempts to start once its required dependencies are fully operational at the destination site.

NEW QUESTION: 2
An administrator builds a single recovery plan that includes workloads protected with Asynchronous (async) and Nearsync replication schedules together with workloads protected by Synchronous replication (syncrep) schedules. Which statement is true regarding this configuration?
A. Mixing schedule types in the recovery plan only affects retention, not recovery success.
B. Mixing async and nearsync in the same recovery plan is supported.
C. Mixing is supported only for test failovers, but not for planned or unplanned failovers.
D. Mixing async and syncrep schedules in one recovery plan is supported.
Answer: B (LEAVE A REPLY)
Nutanix Recovery Plans are designed to orchestrate the failover of complex applications that may consist of multiple virtual machines with different protection requirements. However, there are rules regarding which types of replication schedules can be mixed within a single plan.
Nutanix supports mixing Asynchronous and NearSync replication schedules within the same Recovery Plan.
This is because both Async and NearSync are fundamentally snapshot-based, " passive " recovery models where the system pulls from recovery points stored at the destination cluster.
However, the system generally restricts or warns against mixing Synchronous replication with other types in a single orchestrated plan.
Synchronous replication (syncrep) is an " active-active " or " stretched " model that behaves differently during a failover compared to snapshot-based models. In some software versions, including entities protected by synchronous schedules with those on async/nearsync within a single plan can cause the failover task to fail validation because the orchestration logic for zero-RPO failover is distinct from snapshot restoration.
Therefore, while mixing Async and NearSync is a supported and common practice, administrators must be cautious with Synchronous workloads.

NEW QUESTION: 3
An administrator is evaluating Nearsync for critical VMs in a Nutanix Disaster Recovery environment. What is the minimum RPO, in minutes, supported by Nearsync replication?
A. 0
B. 1
C. 15
D. 60
Answer: (SHOW ANSWER)
Nutanix Disaster Recovery offers several replication tiers to meet varying business Recovery Point Objectives (RPOs). NearSync replication is a specific technology designed to fill the gap between traditional Asynchronous replication and zero-RPO Synchronous replication.
NearSync replication supports a minimum RPO of 1 minute. It achieves this by using Lightweight Snapshots (LWS), which are highly efficient and can be taken much more frequently than standard full snapshots without causing significant I/O overhead on the system. While some

versions of Nutanix software allow NearSync to reach RPOs as low as 20 seconds, the standard minimum RPO recognized in the context of the NCP-BC 7.5 evaluation for minutes-based intervals is 1 minute. Option A (0 RPO) is only achievable through Synchronous replication. Option C (15 minutes) is the typical maximum RPO for NearSync before it transitions to Async, and Option D (60 minutes) is the standard interval for Asynchronous replication.
Understanding these thresholds is essential for administrators when designing protection policies for mission- critical workloads.

NEW QUESTION: 4
An administrator is creating a recovery plan between two availability zones. The production VMs use a subnet that does not exist at the recovery AZ. To ensure the VMs can successfully connect after a failover, what must be configured?
A. A local-only schedule in the protection policy
B. A new protection policy with synchronous replication
C. Network mappings to a valid subnet in the recovery plan
D. A required matching subnet name at the recovery AZ
Answer: (SHOW ANSWER)
When failing over a virtual machine between different physical locations or availability zones, it is common for the network topology at the recovery site to differ from the primary site. If a VM is recovered into a site where its original subnet is unavailable, it will lose connectivity.
To resolve this, Nutanix Recovery Plans include a " Network Mapping " section. The administrator must configure these mappings to explicitly define which production subnet at the source site corresponds to which valid subnet at the recovery site. During the failover process, the orchestrator uses this mapping to automatically re-configure the VM ' s virtual network interface (vNIC) to connect to the correct local VLAN or overlay network at the destination. While " matching subnet names " (Option D) can simplify management, it is not a requirement if the mapping is correctly defined. Synchronous replication (Option B) does not address the networking layer, as it is a storage-level feature. Network mapping is therefore the essential task for ensuring that workloads are not only recovered but also accessible following a failover.

NEW QUESTION: 5
How does Nutanix Disaster Recovery automatically handle Container Mapping if a storage container with the same name does not exist at the recovery Nutanix cluster?
A. The data is temporarily kept in the SSD tier of the remote site pending intervention.
B. The system automatically creates a new Storage Container with the same name on the remote site.
C. The initial replication fails and generates a critical alert requiring manual mapping in Prism Central.
D. The recovery points replicate to a random storage container at the recovery Nutanix cluster.
Answer: (SHOW ANSWER)

Storage Container mapping is a fundamental requirement for successful replication between Nutanix clusters, as it defines where the replicated recovery point data will reside on the destination cluster. In the early versions of Nutanix Disaster Recovery, administrators had to manually map source containers to destination containers to ensure the replication stream had a valid landing zone. However, to simplify the deployment of BCDR solutions and reduce manual configuration errors, modern Nutanix Disaster Recovery behavior has been enhanced. If the system detects that a VM is protected but no manual container mapping has been established in Prism Central, it will proactively check the recovery site for a container with the same name as the source. If that container does not exist, the orchestrator will automatically create a new storage container with the identical name on the remote cluster to allow replication to proceed without interruption. This automation ensures that " out-of-the-box " protection policies can function immediately even if the recovery site hasn ' t been fully pre-provisioned with storage structures. This behavior prevents replication from failing due to minor administrative oversights and maintains the integrity of the data protection schedule.

NEW QUESTION: 6
An administrator is testing the failover scenarios from a source availability zone (AZ1) to a destination availability zone (AZ2). The administrator successfully performs an unplanned failover from AZ1 to AZ2.
The administrator then immediately tries to use the unplanned failover option to fail the VM back to AZ1, but does not see the VM listed under Recovery Plan. What could be the cause?
A. Recovery Plan is not bi-directional.
B. Replication is not yet successful from AZ2 to AZ1.
C. Clean up of Entities is needed to perform failback.
D. Unplanned failover cannot be used to failback VMs.
Answer: (SHOW ANSWER)
Failback is a critical phase of the disaster recovery lifecycle where workloads are returned to their original primary site. When an " Unplanned Failover " is executed, the system assumes the primary site is gone or unmanaged. Once the primary site (AZ1) is restored, the administrator cannot simply " re-run " the failover in reverse without preparing the source site. The " Clean up of Entities " (Option C) is a mandatory step that must be performed on the original primary cluster after it returns to a healthy state. This process clears the stale VM registrations and metadata that existed prior to the outage, signaling to the Nutanix orchestrator that AZ1 is now ready to re-receive the workloads. Until this cleanup is completed, the Recovery Plan will not list the VMs as eligible for failback because it still detects the " ghost " of the VM at the destination. While reverse replication (Option B) is also necessary to prevent data loss during failback, the reason the VMs are entirely missing from the Recovery Plan interface during the initial failback attempt is the lack of entity cleanup. Performing this cleanup ensures a smooth transition back to primary production without metadata conflicts or registration errors.

NEW QUESTION: 7

A company ' s two clusters are part of the same availability zone. Prism Central, which holds the protection policy configuration and recovery plans, is located on Cluster A and protected by the Prism Central Backup and Restore mechanism. What is the first step to take to execute a disaster recovery plan on the datacenter hosting Cluster A?

A. Restore the Prism Central instance from Prism Element on the surviving cluster.

B. Deploy a new Prism Central instance on the surviving cluster and register the cluster to it to access the replicated VM snapshots.

C. Log into Prism Element on the surviving cluster and execute the Recovery Plan.

D. Access the Prism Central web interface via its Floating IP to trigger the " Unplanned Failover " workflow for the affected VMs.

Answer: (SHOW ANSWER)

Prism Central (PC) acts as the centralized control plane for Nutanix Disaster Recovery, storing the critical metadata for Protection Policies and Recovery Plans. If the cluster hosting the Prism Central instance suffers a total failure, the " brain " of the DR orchestration is lost, and the administrator cannot trigger the Recovery Plans through the usual interface. To address this risk, Nutanix provides the " Prism Central Backup and Restore " feature, which replicates the PC metadata to a remote cluster managed by Prism Element (PE). In a disaster scenario affecting Cluster A, the first and most critical step is to recover the management plane. The administrator must log into Prism Element on the surviving cluster and use the native recovery workflow to restore the Prism Central instance from the most recent backup. Once the PC instance is restored and booted on the surviving cluster, it will once again have access to the Recovery Plans and Protection Policies. Only after the management plane is restored can the administrator then initiate the " Unplanned Failover " tasks to bring up the protected workloads. This process ensures that the RTO is kept low by restoring the automated orchestrator rather than manually re-creating policies and plans on a new instance (Option B).

NEW QUESTION: 8

Which port needs to be opened between two Prism Centrals to ensure a successful AZ pairing?

A. 2009

B. 2020

C. 2074

D. 9440

Answer: (SHOW ANSWER)

Disaster recovery orchestration in the Nutanix environment often involves the pairing of two separate Availability Zones (AZs), each managed by its own Prism Central instance. For these two management planes to communicate, exchange metadata, and synchronize Protection Policies and Recovery Plans, specific firewall ports must be open between them.

Port 9440 is the standard port used for all Prism Central and Prism Element web interface and API traffic.

During the AZ pairing process, the primary Prism Central must reach the recovery Prism Central ' s API to establish the trust relationship and pair the sites. Other ports mentioned, such as 2009

(Option A) and 2020 (Option B), are used for data replication between Controller VMs (CVMs) but are not used for the initial management-plane pairing between Prism Central instances. Port 2074 (Option C) is typically used for Nutanix Guest Tools (NGT) communication with the cluster. Therefore, ensuring port 9440 is open between the two Prism Central instances is a mandatory prerequisite for successful disaster recovery site pairing.

NEW QUESTION: 9
An administrator configures a protection policy that replicates workloads to two different recovery AZs (multisite deployment). Replication to both recovery AZs is successful. However, when attempting to perform failover to the second recovery AZ, the administrator cannot select it as a failover target.
Which setup step was most likely missed?
A. Configuring static IP mapping
B. Increasing the snapshot retention count for both AZs
C. Enabling synchronous replication
D. Creating a separate recovery plan for the second AZ
Answer: D (LEAVE A REPLY)

NEW QUESTION: 10
After performing a failover of encrypted AHV VMs from an on-premises AZ to a recovery AZ, an administrator notices that newly written data is not encrypted. Which post-failover cleanup task must be performed to correct this issue?
Freecram.net
A. Enable replication traffic encryption.
B. Reassign storage policies and categories.
C. Reconfigure the replication schedule.
D. Recreate the recovery plan and categories.
Answer: (SHOW ANSWER)
Nutanix supports software-based data-at-rest encryption for AHV workloads. This encryption is typically managed through " Storage Policies " in Prism Central, which are applied to virtual machines based on their " Categories " . For example, a policy might state that any VM in the Compliance:Encrypted category must have its data encrypted on the disk.
When a failover occurs, the VM moves to the recovery cluster. While the data that was replicated is already encrypted, the " Active " Storage Policy at the recovery site must be verified. If the recovery cluster does not have the same Storage Policy configuration or if the VM-to-Policy association was lost during the transition, the AOS storage engine at the recovery site will treat the VM as a standard, non-encrypted workload. This results in " newly written data " (writes occurring after the failover) being stored in plain text. To fix this, the administrator must " reassign " or refresh the Storage Policies and Categories at the recovery site. This ensures that the encryption rules are actively enforced on the new cluster, maintaining the organization ' s security posture post-disaster. This cleanup step is vital for compliance-heavy environments where encryption- at-rest is a mandatory requirement.

NEW QUESTION: 11
An administrator needs to provide the QA team with a VM state from four days ago to test an older browser version. However, the current production VM contains new application code and security patches that must not be deleted or overwritten. Which restore operation should the administrator perform to meet this requirement?
A. Disk-level Attach from Snapshot
B. Self-Service Restore (SSR)
C. In-place Restore (Revert)
D. Out-of-place Restore (Clone)
Answer: (SHOW ANSWER)
Nutanix offers several ways to recover data, ranging from full VM reverts to granular file-level access. In this scenario, the key requirement is to access historical data (from four days ago) without modifying or destroying the current production VM ' s state. Self-Service Restore (SSR) is the most efficient and least disruptive method to achieve this. SSR, which requires Nutanix Guest Tools (NGT) to be installed, allows an administrator or user to mount a historical snapshot disk directly to the existing VM as a new, temporary volume. This allows the QA team to browse the file system of the " four-day-old " state and copy the specific older browser version they need while the production application continues to run undisturbed with its new code and patches. An " In-place Restore " (Option C) would be disastrous here because it would roll back the entire VM, overwriting the new patches. A " Clone " (Option D) would work but is less efficient as it consumes additional vCPU, RAM, and full storage overhead for an entirely new VM. SSR provides the surgical precision needed for file-level testing while maintaining the integrity of the live production environment.

NEW QUESTION: 12
An organization uses a Recovery Plan to protect a SQL Cluster that relies on Volume Groups (VGs). The VGs are configured with hypervisor attachments. The administrator executes a Planned Failover to migrate the SQL Cluster to the Recovery Site. The Failover task completes successfully, but the database administrators report that the database is offline. What is the possible cause of this issue?
A. An administrator failed to configure Protection Domains replication.
B. VGs configured with hypervisor attachments do not reattach automatically after recovery.
C. VGs cannot be protected by Recovery Plans.
D. An administrator failed to configure mapping for the Storage Network in the Recovery Plan.
Answer: (SHOW ANSWER)
Volume Groups (VGs) are used in Nutanix to provide block storage to virtual machines, often for high- performance applications like SQL Server. VGs can be attached to VMs in two ways: " In-guest " (using iSCSI initiators) or " Hypervisor-attached " (where the hypervisor handles the connection).

In a disaster recovery scenario managed by Recovery Plans, there is a technical distinction in how these attachments are handled. While Nutanix Disaster Recovery can replicate the data in the Volume Groups and recover the VM itself, VGs that are " hypervisor-attached " do not currently support automatic reattachment as part of the orchestration workflow in all AOS versions. When the failover completes, the VM powers on at the recovery site, but the VGs are not connected to the VM at the hypervisor level. This causes the application (SQL) to find its data disks missing, resulting in the database staying offline. To resolve this, the administrator must manually reattach the Volume Groups to the VM at the recovery site using the Prism interface or a script. This limitation highlights the need to verify specific application storage configurations during DR testing. For automated recovery of VGs, using " In-guest " iSCSI attachments is often preferred because the iSCSI initiator within the guest OS can re-establish the connection to the new cluster ' s Data Services IP once the network is active.

NEW QUESTION: 13
After a failover and stabilization of workloads at the recovery AZ, an administrator removes several VMs from the protection policy because they are no longer required. Which action should be performed as part of post-failover cleanup to prevent unnecessary storage consumption?
A. Recreate the recovery plan.
B. Reconfigure the replication schedule.
C. Delete the associated recovery points.
D. Disable Nearsync replication.
Answer: (SHOW ANSWER)
Post-failover cleanup is a vital operational step to ensure that the environment remains optimized and that resources are not being wasted on workloads that are no longer part of the production environment. When virtual machines are removed from a Protection Policy after a failover, the orchestration engine stops creating new recovery points for them; however, the existing schedule metadata and historical snapshots may still consume space on both the primary and recovery clusters. Reconfiguring the replication schedule (Option B) is the specific task that refreshes the policy ' s state and forces the system to re-evaluate the storage requirements for the remaining entities. This action ensures that the background cleanup processes (managed by Cerebro and Curator) can properly identify and purge the snapshots associated with the removed VMs once their retention period has passed. Failing to reconfigure the schedule or the associated protection policy can leave " orphaned " snapshots in the storage tier, leading to a steady increase in storage utilization over time. By reconfiguring the schedule, the administrator confirms the new steady-state of the DR environment, ensuring that storage is only consumed by the active, required workloads.

NEW QUESTION: 14
When designing a Disaster Recovery strategy to an NC2 cluster using MST, which technical limitation impacts the Recovery Plan and workload compatibility?

A. The " Zero Compute " configuration requires manual restoration of workloads (automated failover is not supported), and MST currently only protects AHV VMs.
B. To ensure resilience on NC2, both MST DR and " Cluster Protect " features must be configured simultaneously on the same cluster.
C. Automated failover is natively supported in the " Zero Compute " configuration but strictly requires configuring an external IPAM to manage and reassign IP addresses.
D. The " Pilot Light " configuration allows to select multiple on-premise clusters as sources within a single protection policy to automate the failover of AHV and ESXi VMs.
Answer: (SHOW ANSWER)
Nutanix Cloud Clusters (NC2) on AWS or Azure allows for flexible disaster recovery models, including " Pilot Light " and " Zero Compute " . " Zero Compute " is a cost-effective model where data is replicated to cloud storage (like S3 or Azure Blob) without needing a running cluster at the recovery site. Multi-Site Tooling (MST) or Nutanix Disaster Recovery orchestration can be used to manage this process.
In a " Zero Compute " setup, the orchestration of an automated failover is technically complex because there is no active cluster management plane at the recovery site until the failover is triggered. A key technical requirement for successful automated failover in this model is the management of networking and IP addresses. Because the VMs are being brought up in a cloud environment that likely has a different network subnet than the on-premises environment, the system must be able to dynamically assign and manage IP addresses. Therefore, configuring an external IPAM (IP Address Management) system is a strict requirement to ensure that when VMs are restored, they receive valid networking configurations that allow them to communicate with external users and other services. Without a properly configured IPAM and network mapping in the Recovery Plan, the restored workloads would be isolated and non-functional. This highlights the need for careful network planning when designing DR strategies that utilize cloud-native storage and on- demand compute resources.

NEW QUESTION: 15
An administrator wants to protect the snapshots created on the cluster. Only authorized users should be allowed to modify or delete the snapshots on the cluster. How can the administrator harden the security of the snapshots?
A. Configure CVM and AHV hardening.
B. Configure a backup user with admin privileges.
C. Configure cluster lockdown and use SSH keys.
D. Configure Approval policy for snapshots.
Answer: (SHOW ANSWER)
Hardening a Nutanix cluster for Business Continuity involves implementing security controls that protect recovery data from unauthorized destruction. In many modern cyberattacks, specifically ransomware, the goal is not just to encrypt production data but also to delete the snapshots and backups that would allow for recovery. To satisfy the requirement that only authorized

modifications occur, the administrator must move beyond simple Role-Based Access Control (RBAC) and implement an " Approval Policy " .
The Approval Policy feature in Prism Central forces a secondary layer of validation for sensitive operations like snapshot deletion. When a user tries to delete a recovery point, the Nutanix orchestration engine captures the request and holds it in a " Pending " state. The request must be reviewed and approved by a designated person with approval authority. This ensures that even if an administrator ' s account is compromised, the attacker cannot immediately wipe out the snapshots needed to restore the environment. This " multi-person control " is a recognized security best practice for protecting mission-critical data. While cluster lockdown (Option C) prevents unauthorized SSH access, it does not manage the logical management of data through the Prism UI. An Approval Policy provides the specific " modification and deletion " control requested, ensuring that the cluster ' s recovery points remain a reliable " source of truth " following a security incident or administrative error.

NEW QUESTION: 16
A deployment uses native encryption for replication traffic between two clusters. An administrator subsequently enables network segmentation to isolate this traffic. Which specific maintenance step must be performed to ensure the encrypted replication functions correctly on the new segmented network?
A. Manually add the new segmented IP addresses to the external firewall whitelist.
B. Update the IPSec VPN gateway configuration to include the new subnet range.
C. Regenerate the Key Management Server (KMS) tokens for all self-encrypting drives (SEDs).
D. Delete the existing certificates for Cerebro and Stargate services and restart the services.
Answer: (SHOW ANSWER)
Nutanix allows for the encryption of replication traffic " on the wire " using IPSec tunnels between clusters.
When an administrator enables network segmentation for Disaster Recovery, the replication traffic is moved from the default management network to a new, dedicated subnet and set of virtual interfaces.
Because the native replication encryption relies on IPSec gateways to secure the communication, the underlying VPN configuration must be aware of the network paths it is protecting. If the traffic is now originating from a new segmented subnet, the IPSec VPN gateway configuration at both the primary and recovery sites must be updated to include these new subnet ranges. If this step is missed, the encrypted tunnel will not recognize the segmented traffic as valid for the VPN, causing replication to fail even if the physical network is reachable. Regenerating KMS tokens (Option C) relates to data-at-rest encryption on physical drives, not traffic-in-transit. Restarting services (Option D) might be part of the general segmentation workflow, but the specific requirement for " encrypted replication " stability in this context is the alignment of the IPSec VPN policy with the new segmented network topology.

Valid NCP-BC-7.5 Dumps shared by EduDump.com for Helping Passing NCP-BC-7.5 Exam! EduDump.com now offer the **newest NCP-BC-7.5 exam dumps**, the EduDump.com NCP-BC-7.5 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com NCP-BC-7.5 dumps with Test Engine here: https://www.edudump.com/exams/Nutanix/NCP-BC-7.5/premium/ (**112** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 17
An administrator protects workloads between two ESXi clusters using asynchronous (async) replication. After failover, VMs power on successfully but are placed in the default folder and resource pool instead of their original vSphere folders. Why did this occur?
A. Protection Domains must be recreated on the cluster.
B. vCenter server permissions were not configured correctly.
C. Async replication doesn ' t preserve inventory mapping.
D. Snapshots were not consolidated before failover.
Answer: (SHOW ANSWER)
Nutanix Asynchronous (Async) replication is a foundational data protection technology that allows for the periodic transfer of virtual machine snapshots from a primary cluster to a remote site. When using Async replication between two ESXi-based Nutanix clusters, the primary objective is data persistence rather than deep integration with the vSphere management hierarchy. Because Async replication operates at the storage level, it captures the virtual disk (vDisk) data and the VM ' s configuration files (.vmx) but does not natively synchronize the logical inventory metadata maintained by the vCenter Server, such as VM folders, resource pools, or specific vApp configurations.
When a failover is executed, the Nutanix cluster registers the VM to the remote ESXi host and vCenter. Since the inventory mapping details (original folder and resource pool) are not part of the standard Async replication payload, the vCenter Server treats the failing-over VM as a new registration and places it into the default root folder and the default cluster resource pool. To maintain consistent organizational structures during a disaster recovery event, administrators typically rely on Prism Central-based " Recovery Plans. " Recovery Plans offer advanced orchestration that allows for the manual definition of inventory mappings, ensuring that VMs are placed in the correct folders and resource pools upon recovery. Understanding this limitation of legacy Async replication is crucial for administrators designing BCDR solutions that require high levels of organizational consistency post-failover.

NEW QUESTION: 18
An administrator is troubleshooting a hung replication. When checking the replication status, the complete percent and bytes completed do not increase. The administrator checks the Stargate logs on the destination cluster and confirms the writes are failing with " nfs status 28 " . Checking the container info on the destination cluster, free space on the container is 0. The administrator

confirms that there is sufficient space available in the storage pool. What could be the reason for free space on the container to be 0?

An administrator is troubleshooting a hung replication. When checking the replication status, the complete percent and bytes completed do not increase:

```
nutanix@CVM:~$ ncli pd list-replication-status

  Id                        : 77523370
  Protection Domain         : PD
  Replication Operation     : Sending
  Start Time                : 07/09/2020 23:33:25 EST
  Remote Site               : Destination_Cluster
  Snapshot Id               : 77365268
  Aborted                   : false
  Paused                    : false
  Bytes Completed           : 19.67 GiB (21,117,143,286 bytes)
  Complete Percent          : 84.86114
```

The administrator checks the Stargate logs on the destination cluster and confirms the writes are failing:

```
Write file failed with nfs status 28
Backend write failed for data log oplog_lws_usage at offset 16777216 length: 312
```

Checking the container info on the destination cluster, free space on the container is 0:

```
nutanix@CVM:~$ ncli ctr ls
  Id                        : 0005339d-5cee-[illegible]-0000-00000000bc53::3019595
  Uuid                      : 23d1a770-[illegible]-b227-959ee8f7e116
  Name                      : Destination_container
  Storage Pool Id           : 0005339d-5cee-[illegible]-0000-00000000bc53::2457
  Storage Pool Uuid         : d0[illegible]d8b-d02a-4b54-a924-df052e732b4a
  Free Space (Physical)     : 0 bytes
  Used Space (Physical)     : 1.95 TiB (2,147,985,125,376 bytes)
  Max Capacity (Physical)   : 1.95 TiB (2,147,483,648,000 bytes)
  Used by others (Physical) : 0 bytes
  Explicit Res. (Logical)   : 0 bytes
  Thick Prov. (Logical)     : 500 GiB bytes
  Advertised Capacity       : 1.95 TiB (2,147,483,648,000 bytes)
  Replication Factor        : 2
  Oplog Replication Factor  : 2
  NFS Whitelist Inherited   : true
  Compression               : off
  Compression Delay         : off
  Fingerprint On Write      : off
  On-Disk Dedup             : off

  Erasure Code              : off
  Placement Policy Name     : rf2_default
  Software Encryption       : off
```

freecram.net

A. Advertised capacity is set for the container.

B. Replication factor needs to be set to 1.

C. Thick Provisioning is consuming all available space.

D. Compression is turned off.

Answer: (SHOW ANSWER)

In a Nutanix environment, " Advertised Capacity " is a logical limit (a quota) that an administrator can set on a specific storage container. Even if the physical storage pool has terabytes of free

space, if a container ' s advertised capacity is reached, the hypervisor and the Nutanix services will perceive the container as being
100% full.
In this troubleshooting scenario, the " nfs status 28 " error code is a standard indicator for " No space left on device " . Since the administrator confirmed the pool is healthy, the logical limit on the container is the only remaining bottleneck. Replication jobs will " hang " because the destination Stargate service cannot write the incoming data blocks to a container that reports zero free space. To resolve this, the administrator must either increase the advertised capacity value in Prism or remove the limit entirely to allow the container to consume the available physical space in the pool. This highlights a common configuration " trap " where logical quotas-intended for multi-tenant management-can inadvertently block critical BCDR processes like replication if not monitored or sized correctly according to the protected VM growth.

NEW QUESTION: 19
An organization is implementing a Disaster Recovery solution to protect critical workloads by failing them over to a NC2-on-Azure cluster. The administrator needs to ensure that when VMs failover to Azure, their IP addresses are predictable and consistent with the application requirements. Which configuration strategy should the administrator use in the Recovery Plan to achieve predictable IP assignment?
A. Rely on the IPAM networks of the NC2 cluster to maintain the IP assignments.
B. Configure custom IP mapping to assign specific IP addresses to the VMs.
C. Rely on the default behavior where the last octet of the source IP is preserved.
D. Configure In-Guest scripts for NGT to maintain the static IP assignments.
Answer: (SHOW ANSWER)
Maintaining " predictable and consistent " IP addresses during a failover to a public cloud environment like Nutanix Cloud Clusters (NC2) on Azure is a common requirement for multi-tier applications that rely on static IP references.
To achieve this level of precision, the administrator should configure " Custom VM IP Mapping " within the Nutanix Recovery Plan. This feature allows the administrator to explicitly define the exact IP address that each VM should receive upon recovery at the destination site. While some Nutanix network configurations attempt to preserve the " last octet " of an IP address (Option C), this is not always possible or predictable in a cloud environment with existing subnets. Relying on default IPAM behavior (Option A) often results in dynamic, unpredictable assignments. Using in-guest NGT scripts (Option D) is an overly complex and manual workaround that is natively superseded by the Recovery Plan ' s IP mapping feature. By using custom IP mapping, the administrator ensures that the database, application, and web tiers all come up with the exact addresses required for the application to function immediately without manual intervention.

NEW QUESTION: 20

A Protection Policy is configured with a Nearsync replication schedule (15-minute RPO). An administrator observes that the system has temporarily transitioned to an hourly replication schedule. Which scenario would cause this automatic transition from Nearsync to Asynchronous (hourly) replication?

Freecram.net

A. The Nutanix Cloud Infrastructure license expired.

B. The bandwidth is insufficient for the change rate.

C. A new VM was added to the protection category, triggering a full initial sync.

D. A Test Failover operation is currently running on the recovery site.

Answer: (SHOW ANSWER)

NearSync replication is a Nutanix technology that provides RPOs as low as 1 minute by using " Lightweight Snapshots " (LWS). However, NearSync has specific operational requirements to maintain its aggressive schedule . If the system detects that it cannot maintain the 15-minute RPO using lightweight snapshots, it will automatically " downshift " to a standard Asynchronous (hourly) replication mode to ensure that protection is not lost entirely.

One of the most common triggers for this temporary transition is the addition of a new virtual machine to the category associated with the Protection Policy. When a new VM is added, the Nutanix cluster must perform a

" full initial sync " (also known as a seed replication) to transfer all the VM ' s data to the remote site. Because a full seed replication involves transferring much more data than a standard delta update, it often exceeds the time window and processing capabilities of the NearSync LWS engine. To handle this large data transfer safely, the system switches to the Async (hourly) schedule, which uses more robust but slower snapshot mechanisms. Once the initial sync of the new VM is complete and the clusters are in a " consistent " state, the system will automatically attempt to " upshift " back to the NearSync schedule to restore the 15-minute RPO.

This behavior ensures that the system handles growth and changes in the environment while maintaining the highest possible level of protection without administrative intervention.

NEW QUESTION: 21

An administrator configures disaster recovery between an on-premises AZ and a Nutanix Cloud AZ. After enabling DR and assigning protection policies to several VMs with volume groups, the administrator initiates failover and notices the VMs do not have their volume groups attached. Investigation reveals:

* The clusters are successfully paired in Prism Central.
* Protection policies and replication schedules are configured correctly.
* Recovery plan validation completes successfully.
* Required firewall ports between sites are confirmed open.
* Management IP addresses are reachable in both directions.

Despite this, the administrator observes the volume groups are still not attached. Which network-related prerequisite is most likely misconfigured?

A. Nearsync replication does not support volume groups.

B. The clusters do not have a Data Services IP addresses configured.

C. The recovery plan has not been validated.
D. The hypervisor management interfaces are not on the same subnet.
Answer: (SHOW ANSWER)
Nutanix Volume Groups (VGs) provide block-level storage to virtual machines, often used for high- performance databases or shared storage clusters. Unlike standard vDisks, which are managed directly as part of the VM ' s hardware definition, VGs are presented as iSCSI targets. In a Disaster Recovery scenario, when a VM is recovered at a new site, it must be able to " find " and re-attach these Volume Groups to resume operation.
For this attachment to succeed, especially across different clusters or into a Nutanix Cloud AZ, both the source and recovery clusters must have a " Data Services IP Address " configured. The Data Services IP acts as the cluster-wide iSCSI target portal. When a VM fails over, the Nutanix orchestration engine uses this IP to coordinate the iSCSI login and mapping of the Volume Group to the guest OS. If this IP is missing on either cluster, the internal iSCSI redirection and handoff mechanisms will fail, even if the general management networking and Protection Policies are perfectly healthy. This is a common " Day 1 " configuration error because management connectivity (Prism) works without a Data Services IP, but block-level services like VGs and certain NGT features strictly require it. Verifying and configuring the Data Services IP at both availability zones is the primary fix for Volume Group attachment failures during DR failover.

NEW QUESTION: 22
An administrator is validating a newly created Recovery Plan and receives the following warning:
IP addresses xxx.xxx.xxx.xxx cannot be preserved/mapped for VM REPORTVM01.
IP addresses cannot be preserved/mapped for entities
IP might be already in use or will be used by some other VM for recovery. IP cannot be mapped. Only one IP address can be preserved for a vNIC in an IP address management enabled network.
What can the administrator do to resolve this warning without changing the existing IP address assigned to the VM?
A. Disable IPAM on the recovery site subnet to allow Nutanix DR to preserve static IP addresses during failover.
B. Add an additional vNIC to the VM and assign the IP address to the new vNIC to bypass the IPAM conflict.
C. Configure Custom VM IP Mapping in the Recovery Plan to define explicit IP address mapping for the VM.
D. Remove and re-add the VM to the Recovery Plan to trigger a fresh IP address mapping validation.
Answer: C (LEAVE A REPLY)
IP address management (IPAM) in Nutanix networks ensures that IP assignments are tracked to prevent conflicts. When a Recovery Plan is validated, the orchestrator checks if it can successfully assign the required IPs at the recovery site. The warning occurs when the system

detects a potential conflict or cannot automatically " guarantee " the preservation of a static IP, often due to overlapping pools or the way IPAM is configured on the destination subnet.
To resolve this without changing the VM ' s assigned IP, the administrator should use " Custom VM IP Mapping " . This feature allows the administrator to override the default automatic assignment and explicitly state which IP the VM must receive at the recovery site. By manually defining this mapping, the administrator
" claims " the IP within the Recovery Plan ' s logic, satisfying the validation check and ensuring that the VM boots with its correct identity during a failover. Disabling IPAM (Option A) is often not feasible in production environments, and adding a vNIC (Option B) complicates the VM ' s OS configuration. Custom IP mapping is the purpose-built administrative tool for resolving IP assignment ambiguity and ensuring predictable recovery in complex, IPAM-enabled networking environments.

NEW QUESTION: 23
An administrator logs into Prism Central (PC) and notices that a replication of a recovery point is still ongoing from Cluster A to Cluster B on a new setup.
* The bytes being replicated aren ' t changing.
* Both site ' s virtual IP addresses are reachable.
* Connectivity between PC/PE is good.
* There is a dedicated LAN between the clusters.
What could be causing this replication to be stuck?
A. Acropolis is down at the Remote Site.
B. Protection Policy is incorrectly configured.
C. Port 2020 is blocked between CVMs.
D. Port 2030 is blocked between CVMs.
Answer: (SHOW ANSWER)
When a replication job is marked as " ongoing " but the progress (bytes completed) is stagnant, it typically indicates a mismatch in the configuration or a failure in the orchestration logic rather than a total network blackout, especially when virtual IPs are confirmed reachable. In a " new setup, " a very common root cause is an incorrectly configured Protection Policy.
This can include issues such as mismatched category assignments, incorrect target site selection, or an invalid replication schedule that the system is unable to initialize properly. If the policy itself is misconfigured-for instance, if it is attempting to replicate to a storage container that hasn ' t been mapped or lacks permissions- the Cerebro service may initiate the task but then fail to progress the data transfer. While blocked ports (Options C and D) are common culprits in replication failures, they usually result in immediate " unreachable
" alerts or job failures rather than a " stuck " status where connectivity is otherwise reported as healthy. If the Acropolis service were down at the remote site (Option A), the cluster would likely report the site as disconnected or show a critical service alert in Prism. Therefore, the most logical first area to investigate in a new, non-progressing setup is the logic and settings defined within the Protection Policy itself.

NEW QUESTION: 24
An administrator at a retail company is preparing to apply a critical OS security patch to a production SQL Server VM. Before starting, the administrator ensures the VM is part of a Protection Domain and manually creates a local snapshot (Recovery Point). Ten minutes after the patch is applied, the SQL service fails to start, and the database logs indicate corruption in the system registry. The administrator needs to return the VM to its functional state from ten minutes ago as quickly as possible. The administrator decides to perform a Restore from the local snapshot. Under which condition would an In-place Restore (Revert) fail, forcing the administrator to use the Clone (Out-of-place) option instead?
A. NGT (Nutanix Guest Tools) version was updated immediately after the snapshot.
B. The VM ' s memory was increased from 16GB to 32GB to accommodate the patch installation.
C. The VM is currently powered on and cannot be shut down.
D. A new 50GB data disk was added to the VM configuration after the snapshot was taken.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 25
What must be installed on a guest VM prior to enabling Self-Service Restore?
A. Post_Thaw Scripts
B. NGT
C. Pre_Freeze Scripts
D. VirtIO
Answer: (SHOW ANSWER)
Nutanix Guest Tools (NGT) is a software bundle that must be installed on the guest virtual machine (VM) to enable advanced management and data protection features within the Nutanix ecosystem. For the specific functionality of Self-Service Restore (SSR), NGT provides the necessary communication drivers and the Nutanix SSR agent. This agent allows end-users or administrators to mount recovery point snapshots as local drives directly within the guest operating system. Without NGT, the cluster lacks the specialized hooks into the guest OS required to perform granular file-level recovery without restoring the entire VM. Furthermore, NGT facilitates the creation of application-consistent snapshots by interacting with the Windows Volume Shadow Copy Service (VSS), which is often a prerequisite for high-integrity backups of Tier-1 applications like SQL Server or Exchange. By installing NGT, the VM is granted the intelligence to browse and attach historical vDisks, effectively enabling the SSR workflow and reducing the operational burden on the infrastructure team for simple file recovery requests.

NEW QUESTION: 26
An administrator wants to run a test failover and maps the primary production virtual network to the recovery site production virtual network for the test. What are the two most likely outcomes? (Choose two.)
A. Test failover can overlap with production networking and can lead to IP conflicts.

B. Using production networks on recovery AZ for test failover provides the most accurate validation.
C. It defeats the intended isolation between production and test workflows.
D. Test failover does not require network mapping, so the selection of production or test networks will have no impact.
Answer: (SHOW ANSWER)
The fundamental purpose of a " Test Failover " in Nutanix Disaster Recovery is to provide a non-disruptive way to validate the recovery process. During this test, virtual machines are cloned from their most recent snapshots and powered on at the recovery site. If an administrator incorrectly maps the " Test Network " to the actual " Production Network " (VLAN/Subnet) at the recovery site, they create two significant risks.
First, this configuration can lead to immediate IP address conflicts. Because the test VMs have the exact same IP addresses as the live production VMs, having them active on the same logical production network will disrupt routing and potentially cause service outages for real users.
Second, it defeats the " isolation " between production and test workflows. A valid DR test should occur in a " sandbox " where test applications can crash or exhibit errors without any possibility of affecting the live environment. By using the production network, the administrator is effectively running a " live " test on the production infrastructure, which is a dangerous violation of standard BCDR best practices. To avoid these outcomes, the administrator must map the test network to a dedicated, isolated VLAN or overlay network at the recovery AZ that has no physical or logical path back to the production segment.

NEW QUESTION: 27
A security team has detected a ransomware attack that began encrypting files on the primary site at 2:00 AM.
By 2:45 AM, the primary site is still accessible but workloads are compromised. The administrator needs to recover to a known-good state from before the attack.
Which statement shows the action and reasoning that the administrator should consider?
A. Test Failover, to validate the recovery point in isolation before committing to a full failover.
B. Planned Failover, as the primary site is still accessible and a clean final snapshot can be created before failing over.
C. Planned Failover, as it guarantees zero data loss by capturing the most current state of all VMs.
D. Unplanned Failover, as it allows the administrator to select a specific recovery point in time.
Answer: (SHOW ANSWER)
In a Nutanix Disaster Recovery context, choosing the right failover type is critical during a ransomware event.
A " Planned Failover " is designed for maintenance scenarios. Its first step is to synchronize the current data state to the recovery site to ensure zero data loss. In a ransomware attack, the " current state " is encrypted or corrupted data. If an administrator performs a Planned Failover at

2:45 AM, the system will capture and replicate the encrypted files, effectively " backing up " the ransomware to the recovery site.
To successfully recover from ransomware, the administrator must perform an " Unplanned Failover. " Despite the name, this failover type is the correct choice for this scenario because it allows the administrator to manually select a historical recovery point from the destination cluster ' s storage. By choosing a recovery point from 1:00 AM (before the attack began), the administrator can restore the virtual machines to a known- good, unencrypted state. This is a primary use case for maintaining a multi-day or multi-week snapshot retention policy. While a " Test Failover " (Option A) could be used to verify the 1:00 AM snapshot in an isolated network first, the ultimate production recovery is achieved via an Unplanned Failover that specifically targets a point-in-time snapshot that predates the infection.

NEW QUESTION: 28
A third-party backup solution is configured to use Nutanix application-consistent snapshots. Backups are failing for several Windows VMs, although crash-consistent snapshots succeed. What is the most likely cause of the failure?
A. Bandwidth throttling is configured too low.
B. The protection policy uses asynchronous replication.
C. Replication traffic encryption is disabled.
D. NGT is not installed or not communicating.
Answer: (SHOW ANSWER)
Nutanix supports two main types of snapshots: crash-consistent and application-consistent. A crash-consistent snapshot is taken at the storage layer without any awareness of the guest OS state; it is equivalent to pulling the power plug on a server. An application-consistent snapshot is more complex because it requires the Nutanix cluster to " quiesce " the applications (such as SQL Server or Exchange) inside the guest OS, flushing pending writes to disk to ensure data integrity. To perform this quiescence on Windows VMs, the Nutanix cluster must communicate with the Windows Volume Shadow Copy Service (VSS). This communication is handled exclusively by Nutanix Guest Tools (NGT). If NGT is not installed on the VM, or if the NGT agent is not running/communicating with the cluster, the system cannot trigger the VSS hardware provider. In this situation, any request for an " application- consistent " snapshot from a third-party backup solution will fail because the cluster lacks the necessary hooks into the guest OS. However, a " crash-consistent " snapshot will still succeed because it does not require any guest-level interaction. Therefore, when app-consistent backups fail but crash-consistent ones succeed, the most likely root cause is a missing or broken NGT installation on the affected virtual machines.

NEW QUESTION: 29
During routine monitoring, an administrator notices that replication from the primary Nutanix cluster to the recovery cluster has stopped. New recovery points are not being created, and the DR dashboard shows replication errors. Investigation reveals:
* The primary and recovery clusters are both online and healthy.

* Network segmentation was recently enabled to isolate replication traffic from other workloads.
* No recent changes were made to firewall rules, protection policies or RPO settings.
Which action should the administrator take first to investigate the replication failure?
A. Delete and recreate the protection policy to reset replication.
B. Increase the RPO interval setting in order to reduce the replication load.
C. Verify required ports are open and reachable between clusters.
D. Perform a test failover at the recovery site.
Answer: (SHOW ANSWER)
In a Nutanix environment, replication health is heavily dependent on the available bandwidth and the consistency of the network path between the source and target sites. When replication fails following an infrastructure change like network segmentation, it often indicates that the system is unable to sustain the current data change rate over the new path. While verifying ports (Option C) is a logical infrastructure step, the provided answer emphasizes managing the replication load by increasing the Recovery Point Objective (RPO) interval. Increasing the RPO (e.g., from 1 hour to 4 hours) reduces the frequency of snapshot transfers, which can help determine if the replication engine (Cerebro) can successfully complete a transfer when given a longer window or reduced pressure. If increasing the RPO resolves the failure, the administrator can conclude that the network change introduced a bottleneck or limited throughput that is insufficient for the original " heavy " schedule. This diagnostic step helps differentiate between a total connectivity blackout and a performance-related failure caused by the newly implemented network segmentation. Once the baseline replication is restored at a higher RPO, the administrator can then proceed to optimize the network or firewall settings to support the more aggressive RPO requirements originally set by the business.

NEW QUESTION: 30
Which two VM disk types does Self-Service restore support? (Choose two.)
A. SCSI
B. IDE
C. PCI
D. VGs
Answer: (SHOW ANSWER)
Self-Service Restore (SSR) relies on the Nutanix Guest Tools (NGT) agent to mount a historical snapshot disk to the guest OS via the hypervisor ' s virtual controller. For this to work seamlessly, the guest OS and the NGT agent must support the virtual bus types used by the virtual machine. Nutanix SSR natively supports the two most common virtual disk controller types: SCSI (VirtIO) and IDE.
These are the standard formats for most virtual disks in Nutanix AHV and ESXi environments. When a user requests a file-level restore, the system takes the vDisk from the recovery point and attaches it as a new SCSI or IDE device to the running VM, which the guest OS then scans and mounts. Conversely, specialized storage entities like Volume Groups (VGs) (Option D) and certain high-performance direct-attached PCI-based storage (Option C) are not supported by the

standard SSR workflow. Volume Groups, which are often used for shared storage between VMs (like SQL Clusters), require a different recovery method because they are not part of the standard VM-centric snapshot metadata that the SSR agent can browse. Therefore, SSR is primarily designed for the standard virtual disks that form the OS and application drives of typical virtualized workloads.

NEW QUESTION: 31
Which combination of Replication Schedules within the same Recovery Plan will cause the recovery to fail?
A. Including entities protected by synchronous schedules with entities protected by asynchronous or nearsync schedules.
B. Placing multiple categories of synchronous VMs from different clusters into the same Recovery Plan.
C. Combining entities protected by asynchronous replication schedules with entities protected by nearsync replication schedules.
D. Integrating Volume Groups (VG) and Guest VMs that share the exact same asynchronous schedule.
Answer: (SHOW ANSWER)
A Recovery Plan in Nutanix is the orchestration engine that defines the power-on sequence, network mappings, and script execution for disaster recovery. While a single cluster can support various types of replication-Asynchronous, NearSync, and Synchronous-there are strict architectural limits on how these can be mixed within a single automated plan.
The primary conflict arises when mixing Synchronous (0 RPO) schedules with snapshot-based Asynchronous or NearSync schedules. Synchronous replication is an " Active-Standby " or " Active-Active " storage state where the destination site already has the data in a " mounted " but inactive container. In contrast, Asynchronous and NearSync recovery involve " restoring " a VM from a specific recovery point (snapshot), which includes registering the VM and attaching the disks during the failover process. Because the orchestration logic for a Zero-RPO synchronous failover is fundamentally different from a snapshot restoration, Nutanix prevents placing them in the same Recovery Plan. Attempting to do so will cause validation errors and lead to execution failure, as the orchestrator cannot simultaneously manage real-time failover and snapshot-based restoration for the same logical group. To successfully protect an environment with both types of RPO requirements, administrators must create separate Recovery Plans: one dedicated to Synchronous workloads and another for snapshot-based Asynchronous/NearSync workloads, ensuring that each set of VMs follows its appropriate recovery workflow.

Valid NCP-BC-7.5 Dumps shared by EduDump.com for Helping Passing NCP-BC-7.5 Exam! EduDump.com now offer the **newest NCP-BC-7.5 exam dumps**, the EduDump.com NCP-BC-7.5 exam **questions have been updated** and **answers have been corrected** get the

newest EduDump.com NCP-BC-7.5 dumps with Test Engine here: https://www.edudump.com/exams/Nutanix/NCP-BC-7.5/premium/ (**112** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 32
An administrator migrates a guest VM from a legacy Protection Domain-based DR configuration to a Prism Central (PC)-based Protection Policy. Immediately after migration, the administrator considers deleting the legacy Protection Domain snapshots to reclaim storage. According to Nutanix guidance, when is it safe to delete the legacy Protection Domain snapshots?
A. Immediately after assigning the VM to a Protection Policy
B. After performing a planned failover
C. After validating the recovery plan
D. After the first recovery point for the VM is available in PC
Answer: (SHOW ANSWER)
Transitioning from legacy Protection Domains (which are cluster-centric) to modern Prism Central-based Protection Policies (which are management-plane centric) is a common task during Nutanix environment upgrades. This migration involves a " handoff " where the responsibility for snapshots and replication shifts from the local Cerebro service on the cluster to the global orchestration provided by Prism Central.
The primary risk during this migration is the creation of a " protection gap. " If an administrator deletes the legacy snapshots immediately after assigning a new policy, and that new policy has not yet successfully completed its first replication cycle, the virtual machine is effectively unprotected. If a disaster occurs at that exact moment, there would be no valid recovery points at the destination site to restore from. Nutanix best practice dictates that legacy artifacts must be preserved until the new system is verified as operational. Once the first recovery point for the VM appears in the Prism Central " Recovery Points " tab, it confirms that the new Protection Policy has successfully captured the VM ' s state and replicated it to the recovery AZ. At this stage, the legacy snapshots in the Protection Domain become redundant and can be safely deleted to reclaim storage space without compromising the organization ' s ability to meet its Recovery Point Objective (RPO) and Recovery Time Objective (RTO).

NEW QUESTION: 33
An organization requires strict control over data destruction. An administrator configures an Approval Policy for DR snapshots to prevent accidental deletion. A user attempts to delete a recovery point protected by this policy. What is the immediate outcome of this deletion attempt within the Prism Central interface?
A. The operation is blocked, and an email notification is sent to the super-admin for review.
B. The operation triggers an approval request, which changes the action status to " Pending Approval " .
C. The snapshot is marked for deletion but remains recoverable for 24 hours in the recycle bin.
D. The operation fails immediately with an " Access Denied " error message.

Answer: (SHOW ANSWER)
Security and compliance are critical components of a Business Continuity strategy, particularly regarding the integrity of backups and snapshots. To prevent accidental or malicious deletion of critical recovery points, Nutanix Prism Central allows administrators to configure " Approval Policies " . These policies enforce a " four-eyes " principle where certain high-risk actions require validation by a second authorized user.
When an Approval Policy is active for DR snapshots, and a user attempts a deletion, the system does not execute the command immediately. Instead, the Nutanix orchestration engine captures the request and creates a " Pending Approval " task. The user who initiated the request will see the status update in the Task console, and an notification is sent to the designated " Approvers " . The snapshot remains fully protected and available for recovery until an authorized approver reviews and explicitly grants permission for the deletion.
This mechanism provides a vital layer of defense against ransomware or administrative errors that could otherwise lead to the loss of a known-good recovery state. By integrating approval workflows into the storage management layer, Nutanix ensures that data destruction is a deliberate and verified process, meeting the strict control requirements of modern enterprises.

NEW QUESTION: 34
An administrator configures a VM-VM Anti-Affinity policy for a web application cluster to ensure high availability. The application is protected by a Protection Policy in Prism Central. A failover occurs, moving the VMs to the recovery site.
What determines whether the VM-VM Anti-Affinity rules are active on the recovery site?
A. The rules are automatically active because Anti-Affinity policies are intrinsic to the VM metadata replicated by Cerebro.
B. The rules are inactive because Anti-Affinity policies are not supported in Prism Central-based DR.
C. The rules are active only if the corresponding Anti-Affinity Policy exists on the recovery site.
D. The rules are active only if the " Enforce Anti-Affinity " checkbox is selected in the Recovery Plan.
Answer: (SHOW ANSWER)
VM-VM Anti-Affinity rules are used in Nutanix to ensure that specific virtual machines (such as the nodes of a database cluster or redundant web servers) never run on the same physical host, thereby protecting the application from a single host failure. These rules are typically managed through Prism Central (PC) in a modern deployment.
However, in the context of Disaster Recovery (DR) between two different clusters or availability zones, there is a technical limitation regarding the replication of these affinity policies. While the virtual machine ' s data and primary metadata (CPU, RAM, Disks) are replicated by the Cerebro service, the affinity rules are often considered local cluster constraints. In many versions of Nutanix Disaster Recovery (specifically the initial iterations of Leap/Prism Central-based DR), these affinity policies are not natively " synchronized " or enforced at the recovery site during a failover. This means that when the VMs are recovered at the destination cluster, the hypervisor ' s

dynamic scheduler (ADS) may place them on the same host if that host has the most available resources. To maintain anti-affinity at the recovery site, an administrator would typically need to manually recreate the affinity policies at the destination AZ or use specialized orchestration scripts.
Understanding this limitation is vital for ensuring that high-availability requirements are still met after a disaster recovery event has occurred.

NEW QUESTION: 35
An administrator initiates an unplanned failover due to a primary AZ outage. What is the worst-case scenario for data loss and which recovery point will be used by default?
A. Maximum data loss equals the RPO in the protection policy, VMs recover from the most recent recovery point.
B. Maximum data loss equals the longest RPO, and the administrator must always manually select a recovery point.
C. Data loss depends only on recovery plan, and the system uses recovery points at the recovery AZ by default.
D. Maximum data loss is always zero, and the system uses the most recent recovery point by default.
Answer: (SHOW ANSWER)
An " Unplanned Failover " is triggered when the primary site is completely offline or unreachable due to a disaster. Because the primary site cannot be contacted to perform one last synchronization of data, the recovery must rely on the recovery points that were successfully replicated to the destination cluster before the outage occurred.
In this scenario, the worst-case data loss is equivalent to the Recovery Point Objective (RPO) defined in the Protection Policy. For example, if the RPO is set to 1 hour, and the disaster occurs 59 minutes after the last successful replication, the business will lose 59 minutes of data. By default, the Nutanix orchestration engine will automatically select and use the most recent recovery point available at the recovery site to restore the virtual machines. While an administrator can manually choose an older recovery point if needed (e.g., in the case of a ransomware attack), the default behavior is to minimize data loss by using the newest available point. This highlights why setting an appropriate RPO is the most important design decision for managing potential data loss in an unplanned failure event.

NEW QUESTION: 36
An administrator enables network segmentation for Disaster Recovery on a primary cluster to isolate replication traffic. The primary and recovery clusters are in a " brownfield " configuration, meaning the recovery cluster does not yet have network segmentation enabled.
How is the configuration of the recovery cluster handled in this scenario?
A. The system automatically reconfigures the recovery cluster using the new virtual IP address communicated by the primary cluster.

B. The recovery cluster rejects the connection until the firewall rules on the recovery site are manually updated to allow traffic on port 3260.
C. The system creates a temporary VPN tunnel between the interface eth0 of the recovery cluster and ntnx0 of the primary cluster.
D. The administrator must manually enable network segmentation on the recovery cluster.
Answer: D (LEAVE A REPLY)
Network segmentation in a Nutanix cluster allows for the physical or logical isolation of different types of traffic (management, storage, and backplane). When applied to Disaster Recovery, it allows an administrator to dedicate a specific network interface (typically ntnx0) and a dedicated subnet for replication traffic, ensuring it does not contend with VM or management traffic for bandwidth.
However, Disaster Recovery is a site-to-site operation that requires a symmetric network configuration. If an administrator enables segmentation on the primary cluster, the Cerebro service at that site will now attempt to originate and receive all replication traffic on the newly defined segmented IP addresses. If the recovery cluster is still in its original " brownfield " state without segmentation, it will be looking for traffic on its management IP addresses (on eth0). This creates a mismatch where the two sites are effectively " speaking different languages " or looking for each other on the wrong networks. Nutanix does not automatically propagate segmentation settings between clusters because each site may have different physical network topologies or VLAN requirements. Therefore, the administrator must manually enable and configure network segmentation on the recovery cluster to match the primary. Only when both sites have consistent segmented interfaces and the routing/firewall rules are updated to allow communication between these new subnets will the replication link be restored and functional.

NEW QUESTION: 37
An administrator is managing a mission-critical Inventory-VM that is part of a Protection Domain (PD) named PD_Production.
* Monday: A scheduled snapshot of PD_Production is successfully taken.
* Tuesday: Due to a configuration error during a cleanup task, the administrator accidentally removes Inventory-VM from the PD_Production Protection Domain. The VM continues to run, but it is no longer being snapped or replicated.
* Wednesday: A database corruption occurs on Inventory-VM. The administrator finds the local snapshot from Monday and performs an In-place Restore (Revert) to recover the data.
Following the successful completion of the In-place Restore, what is the status of Inventory-VM regarding its data protection?
A. The VM is restored, but it is placed into a new, system-generated PD named Restored_Entities.
B. The VM is restored to Monday ' s data state and is automatically added back into the PD.
C. The VM is restored to Monday ' s data state but remains outside of any PD.
D. The restore fails because the VM is no longer a member of the PD that owns the snapshot.
Answer: C (LEAVE A REPLY)

In Nutanix Protection Domains, the relationship between a virtual machine and a snapshot is tied to the metadata at the time the snapshot was created . When an administrator removes a VM from a PD (Option B), the cluster stops taking new snapshots, but the historical snapshots (recovery points) remain in the storage system until their retention period expires .
If a disaster or corruption occurs (Option Wednesday), the administrator can still use those historical snapshots to perform a " Revert " operation. The Nutanix AOS engine can successfully roll back the VM ' s vDisks to the Monday state because the data exists. However, the " Revert " operation only impacts the data on the disks; it does not re-establish the logical configuration of the Protection Domain. Because the VM was explicitly removed from the PD on Tuesday, the system treats it as an " unprotected " VM after the restore is finished. It will not be included in the next scheduled snapshot cycle of PD_Production. To resume ongoing protection, the administrator must manually add the Inventory-VM back into the PD_Production Protection Domain. This scenario illustrates a critical operational point: restoring data does not automatically restore the protection schedule, and administrators must verify the " Protection Status " of VMs following any manual recovery task to ensure they are not left vulnerable to future failures.

NEW QUESTION: 38
The Disaster Recovery dashboard in Prism Central shows that the protection domain is out of sync and new recovery points are no longer being created at the recovery site. No recent changes were made to the protection policy configuration or RPO settings. During initial troubleshooting, an administrator has observed the following:
* Production VMs are running normally at the primary site.
* The recovery cluster is reachable from Prism Central.
* No storage capacity alerts are seen on either cluster.
* A recent network change was implemented to enable network segmentation between clusters.
* New firewall rules were introduced recently as part of a security hardening effort.
Which action should the administrator take first when evaluating this issue?
A. Delete and recreate the protection policy immediately.
B. Verify if network segmentation was enabled while replication was running.
C. Manually trigger a failover to re-establish replication.
D. Increase the RPO frequency setting in order to generate additional recovery points.
Answer: (SHOW ANSWER)
Network segmentation is a powerful Nutanix feature used to isolate replication and storage traffic from standard management traffic. However, implementing this in a " brownfield " environment-where replication is already active-requires careful coordination. In this scenario, the replication failure coincides with the implementation of network segmentation and new firewall rules. The administrator ' s first action must be to verify if segmentation was enabled while replication was actively running. If segmentation is enabled on one site without a matching configuration on the remote cluster, or if the services (Cerebro
/Stargate) were not properly restarted to bind to the new interfaces, the replication stream will " hang " .

Furthermore, enabling segmentation changes the source and destination IP addresses used for replication traffic. If these new segmented IP addresses were not whitelisted in the recent firewall hardening effort, all " Snapshot-on-wire " traffic will be dropped. By verifying the segmentation state first, the administrator can determine if the network plumbing is consistent across both sites. Recreating policies (Option A) or triggering failover (Option C) would not resolve an underlying network isolation issue and would likely result in further failures until the segmentation and firewall rules are properly aligned.

NEW QUESTION: 39
How does the system behave if the network cannot sustain the data change rate required to maintain a 1- minute RPO between the primary and recovery AZs?
A. The VMs continue to be protected in 1-minute RPO, but write I/O is degraded.
B. The system automatically switches to Asynchronous replication mode.
C. The protected VMs are failed over.
D. All write I/O to the protected VMs is paused.
Answer: (SHOW ANSWER)
Nutanix NearSync replication is designed to provide very low RPOs (1 to 15 minutes) by using " Lightweight Snapshots " (LWS). To maintain a 1-minute RPO, the system requires sufficient network bandwidth to transfer the delta changes before the next snapshot interval begins.
If the network bandwidth is insufficient or if the VM ' s data change rate (churn) suddenly spikes beyond the network ' s capacity to keep up, the system is designed to prioritize protection over RPO strictness. In this scenario, the Nutanix Cerebro service will automatically " downshift " the replication from NearSync to standard Asynchronous replication mode. This transition allows the system to use more traditional, robust snapshots that can handle longer transfer times, ensuring that the VM remains protected, even if the RPO temporarily increases to an hourly interval. Once the network conditions improve or the data churn decreases, the system will automatically attempt to " upshift " back to the 1-minute NearSync schedule. This behavior prevents application performance issues like write-pausing (Option D) or I/O degradation (Option A) and avoids unnecessary failovers (Option C), providing a resilient and adaptive data protection mechanism.

NEW QUESTION: 40
An administrator at a retail company is preparing to apply a critical OS security patch to a production SQL Server VM. Before starting, the administrator ensures the VM is part of a Protection Domain and manually creates a local snapshot (Recovery Point). Ten minutes after the patch is applied, the SQL service fails to start, and the database logs indicate corruption in the system registry. The administrator needs to return the VM to its functional state from ten minutes ago as quickly as possible. The administrator decides to perform a Restore from the local snapshot. Under which condition would an In-place Restore (Revert) fail, forcing the administrator to use the Clone (Out-of-place) option instead?
Freecram.net
A. NGT (Nutanix Guest Tools) version was updated immediately after the snapshot.
B. The VM ' s memory was increased from 16GB to 32GB to accommodate the patch installation.
C. The VM is currently powered on and cannot be shut down.
D. A new 50GB data disk was added to the VM configuration after the snapshot was taken.

Answer: (SHOW ANSWER)

An In-place Restore (Revert) is the fastest way to roll back a VM because it overwrites the current vDisks with the data from the snapshot. However, this operation has strict requirements regarding the consistency of the VM ' s hardware configuration. The most common technical blocker for an in-place revert is a change in the disk topology between the time the snapshot was taken and the time of the restore.

If the administrator added a new 50GB data disk (Option D) after taking the snapshot, the VM ' s current configuration now includes a vDisk that does not exist in the snapshot ' s metadata. Nutanix AOS cannot " revert " a disk that has no point-in-time reference in the recovery point. In this situation, the " In-place Restore

" button is typically disabled or will throw an error, forcing the administrator to use the " Clone " (Out-of- place) option. Cloning creates a new VM based exactly on the snapshot ' s metadata, effectively ignoring the newly added disk. While changing memory (Option B) or NGT versions (Option A) might affect performance or guest features, they do not usually prevent the storage-level revert operation. Being forced into a " Clone " restore increases the RTO slightly, as the administrator must then delete the corrupted VM and ensure the new clone has the correct network settings.

NEW QUESTION: 41

What is the expected outcome that the administrator must manage?

A. The primary site will continue to serve I/O locally to prevent downtime, while the secondary site remains in standby.

B. The secondary site will acquire the lock after a built-in delay, and the primary site will halt all I/O.

C. The Witness will automatically shut down the secondary site to protect data integrity on the primary site.

D. Both sites will attempt to acquire the lock, but the primary site will acquire the lock and continue to serve I/O locally to prevent downtime.

Answer: (SHOW ANSWER)

This question refers to the behavior of Nutanix Metro Availability in conjunction with a Witness VM during a site-to-site communication failure. Metro Availability provides a zero-RPO, synchronous replication solution.

To prevent " split-brain " (where both sites try to serve the same data independently), the Witness VM acts as a tie-breaker. When the two clusters lose connectivity with each other, both will immediately attempt to contact the Witness to " acquire the lock " on the protection domain. The Witness is programmed to prioritize the original " Active " site. If the primary cluster can reach the Witness, it will successfully acquire the lock and continue serving I/O locally to the guest VMs, ensuring zero downtime for the applications. The secondary site, unable to reach the primary or acquire the lock from the Witness, will transition its storage container to an " Inactive " state to protect data integrity. The administrator must manage this state by ensuring the Witness is always reachable from both clusters via independent network paths. If the primary site cannot

reach the Witness, it will stop serving I/O to prevent inconsistent data states, highlighting the critical importance of Witness availability in a high-availability BCDR design.

NEW QUESTION: 42
A Veeam backup job is configured to use Nutanix application-consistent snapshots for several Windows VMs. Crash-consistent snapshots succeed, but Veeam reports that application-consistent snapshots fail with errors related to VSS processing. Upon investigation, the administrator confirms:
* The VMs are powered on.
* Network connectivity between the backup server and cluster is healthy.
Which configuration should the administrator verify to resolve the issue?
A. Enable incremental backups on the VMs.
B. Ensure NGT is installed on the VMs.
C. Increase the snapshot retention count.
D. Configure synchronous replication.
Answer: (SHOW ANSWER)
Application-consistent snapshots are critical for ensuring data integrity for transactional workloads, such as databases. In a Nutanix AHV or ESXi environment, the cluster achieves application consistency by communicating with the guest operating system ' s Volume Shadow Copy Service (VSS). This communication is facilitated exclusively through Nutanix Guest Tools (NGT). When NGT is installed and enabled, it provides a VSS hardware provider that Nutanix uses to " quiesce " (pause/stabilize) the application
' s I/O and flush pending writes to the disk before the snapshot is taken. If crash-consistent snapshots are working but application-consistent ones are failing, it is a primary indicator that the VSS hooks provided by NGT are either missing, outdated, or not communicating correctly with the Nutanix cluster services. Without NGT, the system cannot verify the internal state of the guest OS, leading to VSS processing errors in third- party backup solutions like Veeam or Commvault. Ensuring NGT is installed and the NGT agent is running within the guest VM is the mandatory first step to resolving any snapshot failure involving application-level consistency.

NEW QUESTION: 43
A remote office deployment consists of a two-node Nutanix hybrid cluster. An administrator attempts to configure a protection domain with a 5-minute RPO (Nearsync) replicating to a central datacenter.
Why is the administrator unable to successfully configure this Nearsync schedule?
A. Nearsync replication requires a minimum of three nodes in the cluster.
B. The minimum RPO supported for a two-node cluster is 15 minutes.
C. Nearsync is only supported on All-Flash clusters, not hybrid models.
D. Two-node clusters require a Witness VM to enable Nearsync replication
Answer: (SHOW ANSWER)

Nutanix NearSync replication provides a middle ground between traditional asynchronous replication and synchronous mirroring, offering RPOs as low as 1 minute using Lightweight Snapshots (LWS). However, the LWS mechanism and the high-frequency metadata operations required to maintain a 1-to-15 minute RPO have specific hardware and cluster-level requirements.
One of the strict prerequisites for enabling NearSync is the cluster size. Nutanix requires a minimum of a three-node cluster for NearSync replication. This is because the system must have enough resources to distribute the LWS metadata and handle the increased I/O overhead without compromising the cluster ' s availability or performance. In a two-node cluster (which is common for small ROBO deployments), the system does not meet the minimum redundancy and resource thresholds required for the NearSync engine to operate reliably. While two-node clusters support standard Asynchronous replication (with a 60-minute RPO or higher), they are restricted from using the NearSync lightweight snapshot engine. An administrator attempting to set a 5-minute RPO on a two-node cluster will find the option either grayed out or the task will fail validation. To achieve a 5-minute RPO, the organization would need to expand the remote office cluster to at least three nodes or settle for a higher RPO supported by standard asynchronous replication on the existing hardware.

NEW QUESTION: 44
Which statement describes the most efficient and least disruptive method to achieve a file-level recovery without rolling back the entire VM?
A. Leverage the Nutanix Files " Previous Versions " (FLR) integration. By accessing the .snapshot hidden directory from the root of the Linux file system.
B. Perform an " In-Place Restore " Prism Central, but ensure the " File-Level Recovery " checkbox is selected to prevent overwriting the entire boot vDisk.
C. Clone the snapshot into a new VM, boot it into single-user mode, and use WinSCP to transfer the missing files back.
D. Utilize the Self-Service Restore (SSR) functionality provided by Nutanix Guest Tools (NGT) and restore only deleted files.
Answer: (SHOW ANSWER)
Nutanix offers several data recovery methods, but the objective of " least disruptive " and " most efficient " is specifically addressed by the Self-Service Restore (SSR) feature. SSR is powered by Nutanix Guest Tools (NGT), which must be installed on the virtual machine. This functionality allows a user or an administrator to mount a disk from a historical recovery point directly to the active, running virtual machine as a temporary drive.
Once the snapshot is mounted, the guest operating system treats it as a standard local disk, allowing the user to browse the folder structure and perform a surgical restore of only the specific files that were deleted or corrupted. This process avoids the significant overhead and downtime associated with rolling back the entire VM (Option B) or cloning the VM to a new entity (Option C), which would consume additional compute and storage resources. Option A refers to Nutanix Files (NAS) functionality, which is separate from the block- level VM snapshot recovery discussed

here. By using SSR/NGT, the production environment remains online and undisturbed, and the recovery process is delegated to the level of individual files, providing the highest level of granularity and efficiency in typical day-to-day data loss scenarios.

NEW QUESTION: 45
An administrator performed an unplanned failover from AZ1 (primary) to AZ2 (recovery). Later, AZ1 is restored and the administrator wants to fail back the entities to AZ1. Which two statements are correct?
(Choose two.)
A. Failback uses the same recovery plan, but the failover operations for failback must be performed at the primary AZ.
B. Unplanned failover to AZ1 can use recovery points that were created locally in the past at AZ1, even if no recovery points replicate back from AZ2.
C. After unplanned failover, replication begins in the reverse direction, and failback success depends on recovery points being replicated back from AZ2 to AZ1 after failover.
D. After unplanned failover, failback success depends on recovery points being created locally on AZ1 before outage.
Answer: (SHOW ANSWER)
Nutanix Disaster Recovery is designed to be bi-directional, supporting both failover and failback . In an unplanned failover, the primary site (AZ1) is assumed to have crashed. Once AZ1 is brought back online, the system must synchronize any data changes that occurred while the VMs were running at AZ2.
Statement C is correct because once the VMs are powered on at AZ2, the Nutanix Cerebro service automatically reverses the protection policy. New recovery points are created at AZ2 and replicated back to AZ1. For a successful, modern failback with zero data loss from the cloud session, the administrator must wait for these recovery points to replicate back to AZ1 before initiating the failback.
Statement B is also correct and describes a specific " recovery of last resort " scenario. If AZ1 is restored but connectivity to AZ2 is permanently lost (e.g., the cloud provider has an issue), an administrator can still " Unplanned Failover " back to AZ1 using the historical snapshots that existed locally on AZ1 before the original outage. This would result in data loss equal to the time the VMs were running at AZ2, but it allows for service restoration when the recovery site itself fails. These two statements highlight the flexibility of the Nutanix DR architecture in handling both controlled failback and emergency local restoration.

NEW QUESTION: 46
What must an administrator verify before starting a migration of entities from a protection domain to a protection policy?
A. The entities must be added to a category in Prism Central.
B. The recovery plan must be active.
C. The entities must have no ongoing replication tasks.

D. A full backup of the protection domain must be completed.

Answer: (SHOW ANSWER)

Migrating from legacy Protection Domain (PD) based management to the modern Prism Central (PC) based Protection Policies is a recommended best practice for scaling Nutanix Disaster Recovery. This migration involves moving VMs from a " Prism Element " centric management model to a centralized, category-based model in Prism Central.

To ensure that no data is lost during this transition, the most critical prerequisite is to verify that a successful, full replication (backup) of the Protection Domain has recently completed. This ensures that the recovery site already possesses a complete and consistent set of data for the VMs before they are unlinked from the old policy and attached to the new one. While adding VMs to a category (Option A) is a step in the process, it is not the first verification step required for data safety. If an administrator attempts to migrate while a replication is incomplete or failing, the transition to the new policy might leave the VMs unprotected for an extended period, or recovery points might be missing if a disaster occurs mid-migration. By completing a full backup of the PD first, the administrator creates a " safety net " of recovery points that can be used if the migration itself encounters issues. This ensures that the transition to modern DR is performed with minimal risk to the business ' s data protection objectives.

Valid NCP-BC-7.5 Dumps shared by EduDump.com for Helping Passing NCP-BC-7.5 Exam! EduDump.com now offer the **newest NCP-BC-7.5 exam dumps**, the EduDump.com NCP-BC-7.5 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com NCP-BC-7.5 dumps with Test Engine here: https://www.edudump.com/exams/Nutanix/NCP-BC-7.5/premium/ (**112** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 47

An administrator plans on performing a failover of a VM from a source cluster to a DR cluster. The administrator needs the VM ' s IP address preserved after failover. In what scenario is the VM ' s IP address preserved?

A. Use unique networks for source and destination networks.

B. The IP address is always preserved, no additional configuration required.

C. The gateway and subnet mask are left empty in the network mapping.

D. The source and destination network subnet must be stretched.

Answer: (SHOW ANSWER)

Preserving IP addresses during a failover is a common requirement for applications with hardcoded IP references or complex inter-service dependencies. In a standard multi-site DR setup, the source and destination sites reside on different Layer 3 subnets, meaning a VM would naturally require a new IP address to be routable at the recovery site. To achieve " IP Preservation, " the underlying network infrastructure must support Layer 2 (L2) stretching.

L2 stretching (often implemented via VXLAN or specialized VPN tunnels like Nutanix Flow Virtual Networking) allows a single logical subnet to span multiple physical locations. When a VM fails over to a cluster on a stretched subnet, it effectively stays on the same network broadcast domain. As a result, its IP address, default gateway, and subnet mask remain valid and unchanged. If the subnets are not stretched, the administrator must rely on the Recovery Plan ' s " IP Mapping " feature to assign a new, valid IP at the destination-which does not " preserve " the original address. While preserving IPs simplifies application recovery, it requires more complex network engineering to ensure that traffic is correctly routed to the active site during a disaster, highlighting the trade-off between application simplicity and infrastructure complexity in BCDR design.

NEW QUESTION: 48
An administrator needs to provide leadership with a report showing the timeframe of each step during Recovery Plan failover testing. How should the administrator produce this report?
A. Filter the Events log in Prism Central for Recovery Plan activity.
B. Review the built-in execution report from the test failover results.
C. Generate a custom report within Prism Central Intelligent Operations.
D. Use ncli on a CVM to query Recovery Plan execution history.
Answer: (SHOW ANSWER)
Validation is a core pillar of a robust Business Continuity and Disaster Recovery (BCDR) strategy. Nutanix provides built-in orchestration through Recovery Plans, which automate the failover process of virtual machines between sites. When a " Test Failover " is executed, the system performs all the necessary steps, such as booting VMs in an isolated network, applying network mappings, and executing any custom scripts.
To provide visibility into this process, Nutanix generates an automated execution report upon completion of any Recovery Plan task. This report provides a detailed chronological breakdown of every action performed by the orchestrator, including the exact start and end times for each stage (e.g., " Replication Check, " " Power On, " " IP Assignment "). This built-in reporting is superior to manual event log filtering because it aggregates all related activities into a single, digestible summary that specifically tracks the performance of the DR plan. For leadership and compliance teams, these execution reports serve as documented proof that the RTO (Recovery Time Objective) is achievable and that every orchestrated step was validated during the test window.

NEW QUESTION: 49
An administrator is concerned about the operational overhead of manually creating and managing categories across both the primary and recovery sites to ensure Recovery Plans function correctly during failover and failback. What Nutanix Disaster Recovery behavior minimizes this concern?
A. Categories only need to be configured at the recovery site since that is where failover operations are executed.

B. Categories must be manually replicated using the Prism Central API between sites to ensure consistency.
C. Recovery Plans replicate categories only during an active failover event to minimize replication traffic.
D. Categories are automatically synchronized to the recovery site AZ when they are specified in a Recovery Plan.
Answer: (SHOW ANSWER)
Nutanix Disaster Recovery (formerly known as Leap) utilizes Prism Central-based " Categories " to group virtual machines for protection. These categories are used within Protection Policies to determine replication schedules and within Recovery Plans to determine power-on sequences and network mappings. A common operational concern is the need to manually replicate these category definitions across different availability zones (AZs).
To reduce this operational overhead, Nutanix has implemented an automatic synchronization mechanism.
When an administrator creates a Recovery Plan in Prism Central and specifies certain categories for failover, the system automatically synchronizes the definition and metadata of those categories to the paired recovery AZ. This ensures that if the primary Prism Central instance becomes unavailable, the recovery site already possesses the necessary grouping logic to identify and recover the correct virtual machines . This " policy- driven " approach ensures consistency across sites and eliminates the risk of human error associated with manual configuration (Option B). By automating the lifecycle of categories between paired sites, Nutanix ensures that the disaster recovery environment is always aligned with the production environment, allowing for seamless failover and failback operations without the need for repetitive administrative tasks.

NEW QUESTION: 50
An administrator is tasked with ensuring the VMs do not experience downtime during an upcoming network maintenance on the primary cluster. The VMs are protected by a Protection Policy and are configured under a Recovery Plan. What failover mechanism should the administrator use to ensure the VMs are available on the target cluster before the maintenance window?
Freecram.net
A. Unplanned Failover
B. Planned Failover
C. Validate the Recovery Plan
D. Test Failover
Answer: (SHOW ANSWER)
In a Business Continuity strategy, " downtime " can be categorized as either unexpected (disaster) or controlled (maintenance). When an administrator knows that a disruption is coming, such as scheduled network maintenance on a primary cluster, they should utilize the " Planned Failover " mechanism within the Nutanix Recovery Plan.
A Planned Failover is a graceful, orchestrated migration. It first shuts down the virtual machines at the primary site to ensure all data is flushed and no new writes occur. It then performs a final

synchronization to the recovery site to ensure zero data loss (Zero RPO), and finally powers the VMs on at the destination site according to the plan ' s sequence. This differs significantly from an " Unplanned Failover " (Option A), which assumes the primary site is already offline and may result in minor data loss depending on the last successful replication. A " Test Failover " (Option D) merely validates the process in an isolated environment and does not move the actual production workload. By executing a Planned Failover before the maintenance begins, the administrator ensures that the services are safely running on the target cluster, maintaining application availability and meeting the objective of zero downtime during the maintenance window.

NEW QUESTION: 51

An organization is finalizing its Disaster Recovery (DR) plan. The primary objective is to balance cost- efficiency with a target RTO of under 15 minutes. Data currently resides in Object Storage, but the team is debating between a Zero Compute approach and a Pilot Light approach. Why would a Pilot Light infrastructure be selected over a Zero Compute model despite the higher " Moderate " cost?

A. Only model that allows for a " Power-On " recovery trigger.

B. Lowest cost option available for long-term idle infrastructure.

C. Faster recovery due to metadata and cluster already existing.

D. Eliminates the need for any Object Storage (S3/Blob) costs.

Answer: (SHOW ANSWER)

When designing disaster recovery to a cloud environment (such as NC2 on AWS or Azure), organizations must choose a compute model that aligns with their Recovery Time Objective (RTO). The " Zero Compute " model is the most cost-efficient because it stores only the data (snapshots) in low-cost Object Storage (S3 or Blob) and does not maintain a running cluster. However, the RTO for Zero Compute is high because, in a disaster, the organization must first deploy a new Nutanix cluster, configure it, and then begin the process of hydrating data from the Object Storage.

In contrast, the " Pilot Light " model involves keeping a minimal, active Nutanix cluster (e.g., 3 nodes) running at the recovery site. While this carries a moderate ongoing cost, it significantly reduces the RTO.

Because the cluster already exists, the management plane (Prism Central), the storage containers, and all metadata are already active. When a failover is triggered, the system only needs to power on the VMs or perform a small amount of data hydration. This allows the organization to meet aggressive RTO targets of under 15 minutes, which is generally impossible with the Zero Compute model. Pilot Light provides the " ready-to-go " infrastructure needed for mission-critical applications that cannot afford the multiple hours required to provision a fresh cluster from scratch during an emergency.

NEW QUESTION: 52

An administrator is preparing to configure DR between an on-prem AZ and Nutanix Cloud AZ. Replication fails immediately after configuration. Which prerequisite should be verified?

A. The clusters are running identical hypervisors.

B. Both clusters have external IP addresses.

C. Deduplication is disabled.

D. Synchronous replication is configured.

Answer: (SHOW ANSWER)

Establishing disaster recovery between an on-premises data center and a Nutanix Cloud Availability Zone (AZ) requires a robust communication path that can traverse different network boundaries. Unlike local replication within a single site, cross-site replication-especially to a public cloud environment-relies on the ability of the clusters to identify and reach one another over an external network. Nutanix Disaster Recovery requires that both the on-premises cluster and the Nutanix Cloud AZ instance have external IP addresses configured for their respective Controller VMs (CVMs) and virtual interfaces.

These external IP addresses allow the Cerebro service at the source site to establish a secure handshake with the Cerebro service at the cloud site. Without these routable external IPs, the replication traffic is unable to find its destination, leading to the immediate failure of the configuration as observed in this scenario. While identical hypervisors (Option A) are often preferred for simplicity, Nutanix supports Cross-Hypervisor DR (CHDR). Furthermore, deduplication status (Option C) does not prevent the establishment of a replication link. Synchronous replication (Option D) is restricted by latency requirements and is not a fundamental prerequisite for basic connectivity to a cloud AZ. Therefore, verifying the presence of external, reachable IP addresses is the mandatory first step in troubleshooting cross-site connectivity issues between on-premises and cloud environments.

Valid NCP-BC-7.5 Dumps shared by EduDump.com for Helping Passing NCP-BC-7.5 Exam! EduDump.com now offer the **newest NCP-BC-7.5 exam dumps**, the EduDump.com NCP-BC-7.5 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com NCP-BC-7.5 dumps with Test Engine here: https://www.edudump.com/exams/Nutanix/NCP-BC-7.5/premium/ (**112** Q&As Dumps, **35%OFF Special Discount Code: freecram**)