

Nutanix Certified Systems Engineer-Core (NCSE-Core)

Nutanix NCSE-Core

Version Demo

Total Demo Questions: 25

Total Premium Questions: 255

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Discover	114
Topic 2, Design	89
Topic 3, Demonstrate	52
Total	255

QUESTION NO: 1

73. Which two VMs require further data analysis to accurately size resources? (Choose two) ??

- A. prodsrv02
- B. prodsrv05
- C. prodsrv06
- D. prodsrv07

ANSWER: B D

Explanation:

prodsrv05 and **prodsrv07** require further data analysis before their resource requirements can be sized accurately. A defensible VM-sizing recommendation must be based on sufficient performance telemetry rather than a single observed value or an incomplete inventory record. The additional analysis should establish CPU utilization and ready time, memory active/consumed behavior, storage capacity growth, I/O rate, I/O size, read/write mix, latency, and workload peaks over a representative business cycle. It should also account for application characteristics such as databases, batch jobs, backups, and maintenance windows, because these can materially alter the required compute and storage profile.

Nutanix sizing is intended to use workload-specific inputs to model the appropriate cluster configuration, including compute, capacity, and performance needs. Gathering more representative data for prodsrv05 and prodsrv07 therefore reduces the risk of either under-sizing the environment or allocating unnecessary resources. The resulting data can be used as input to Nutanix Sizer and validated against the organization's growth, availability, and protection requirements. See [Nutanix Sizer](#) and the [Nutanix Cloud Platform](#) overview for Nutanix sizing and platform context.

QUESTION NO: 2

An SE is working to size a prospects environment that includes production Oracle databases.

Which Oracle output should the SE request from the prospect to collect accurate sizing information?

- A. STATSPACK
- B. Active Session History Report
- C. Oracle Export Schema
- D. Automatic Workload Repository Report

ANSWER: D

Explanation:

Automatic Workload Repository Report is the appropriate output to request because it provides a time-based, production-derived view of Oracle database workload characteristics needed for infrastructure sizing. An AWR report compares database snapshots over a chosen interval and captures essential indicators such as DB time, average active sessions, CPU utilization, logical and physical I/O rates, I/O latency, throughput, wait events, SQL workload, and instance activity. This allows the SE to evaluate both normal operating behavior and peak-demand periods rather than sizing from a static database footprint alone.

For an accurate assessment, the prospect should generate reports over representative business intervals, including known busy windows and, when relevant, batch or reporting cycles. Reviewing these intervals helps establish compute, memory, and storage-performance requirements while identifying whether the workload is CPU-, latency-, or concurrency-sensitive. Oracle documents that AWR retains performance statistics collected through periodic snapshots and that AWR reports present the performance data for a selected snapshot range. This makes the report a reliable source for understanding the production workload profile used in a sizing exercise. See Oracle's [guidance on generating AWR reports](#) and its [Automatic Workload Repository overview](#).

QUESTION NO: 3

A prospect is evaluating Nutanix for a large number of Oracle RAC deployments. Time to deployment is a top priority, and the prospect has developed a self-service tool in-house to expedite these deployments.

Which Nutanix feature should be discussed?

- A. Calm
- B. Xtract for Databases
- C. ABS
- D. REST API

ANSWER: A

Explanation:

Calm is the appropriate feature to discuss because it provides application automation and self-service provisioning for complex, multi-VM applications such as Oracle RAC. A Calm blueprint defines the application architecture, VM configuration, installation steps, dependencies, and lifecycle actions in a repeatable format. This allows Oracle RAC environments to be deployed consistently at scale rather than requiring manual build procedures for each cluster.

Calm also aligns with the prospect's existing self-service investment. Its REST APIs and integration capabilities enable an in-house portal or automation tool to request and manage deployments programmatically, while Calm orchestrates the underlying provisioning workflow. IT teams can standardize approved Oracle RAC blueprints, apply governance through projects and role-based access, and allow authorized users to deploy from a catalog. This combination directly addresses the priority of reducing deployment time while maintaining consistency and operational control across a large deployment estate.

For Nutanix's overview of application automation and self-service capabilities, see [Nutanix Calm](#). Additional product documentation is available through the [Nutanix Support Portal documentation library](#).

QUESTION NO: 4

A prospective customer wants to build a private cloud in addition to provisioning workloads across multiple public cloud providers.

Presently, the staff lacks the time and skillset for on-premises infrastructure management.

Which two features should qualify this project as a Nutanix opportunity?

(choose two)

- A. Flow
- B. Calm
- C. Cloud Connect
- D. All-flash nodes
- E. Prism 1-Click Operations

ANSWER: B E

Explanation:

Calm is appropriate because it provides cloud automation and orchestration for hybrid multicloud environments. It enables teams to create reusable blueprints, offer self-service provisioning through a marketplace, and apply governance to application deployments. This aligns with the customer's need to provision workloads across private-cloud infrastructure and multiple public-cloud providers without requiring staff to perform every deployment manually. Nutanix Cloud Manager, which

includes Calm capabilities, is designed to help organizations automate and govern multicloud operations; see [Nutanix Cloud Manager](#).

Prism 1-Click Operations is also a strong fit for the stated lack of on-premises infrastructure-management time and skills. Prism centralizes management and simplifies common lifecycle tasks, including nondisruptive software and firmware upgrades, cluster expansion, and health-focused operational workflows. These capabilities reduce administrative effort and operational complexity for the private-cloud foundation, allowing a smaller or less specialized team to manage the environment more confidently. Nutanix describes Prism as its management solution for infrastructure operations and lifecycle management at [Nutanix Prism](#). Together, Calm and Prism 1-Click Operations address both multicloud workload provisioning and simplified private-cloud operations.

QUESTION NO: 5

An online retailer currently support over 200 separate PostgreSQL database instances, often with more than 10 copies of each individual database. Their Database Administrators have scripted much of the deployment process for creating dev/test copies, but their infrastructure team struggles to keep up with storage capacity demands. Which Nutanix product should the systems engineer position ?

- A. Beam
- B. Karbon
- C. Calm
- D. Era

ANSWER: D

Explanation:

Era is the appropriate product because it is Nutanix's database automation and lifecycle-management solution, now known as Nutanix Database Service (NDB). It is designed for organizations operating many database instances and needing rapid, self-service provisioning of development and test environments without repeatedly consuming the full storage footprint of production databases. For supported database engines, including PostgreSQL, Era/NDB uses database-aware automation to provision, clone, patch, back up, and recover databases consistently.

The key requirement is the large number of copies required for dev/test alongside constrained storage capacity. Era/NDB addresses this through its Time Machine capability, which captures restore points and enables space-efficient clones. These clones can be created quickly from a source database or point in time while sharing underlying data where appropriate, substantially reducing the capacity required compared with maintaining many independent full copies. This allows DBAs to retain governance and repeatable database operations while giving developers faster access to usable database copies. Nutanix documents NDB's database lifecycle automation and copy-data-management capabilities at [Nutanix Database Service](#) and provides technical documentation at [Nutanix Database Service User Guide](#).

QUESTION NO: 6

What is an important consideration when using RVTools output for sizing a solution?

- A. Only powered on VMs are included in the data.
- B. it only captures provisioned storage
- C. No information on the physical hosts is collected
- D. It is a point in time snapshot of the environment.

ANSWER: D

Explanation:

It is a point in time snapshot of the environment. RVTools collects VMware vCenter inventory and configuration information at the time the report is generated. Consequently, its CPU, memory, storage, virtual machine, and host details represent observed values from that single collection rather than a sustained record of workload behavior. For solution sizing, this means the RVTools export is an important starting point for establishing inventory, capacity allocation, and configuration context, but it should be reviewed alongside performance history and projected business requirements. A snapshot taken during a quiet period, for example, may not represent normal or peak resource demand; similarly, planned growth, application changes, resiliency requirements, and expected consolidation ratios can materially affect the final design. The RVTools documentation describes the tool as collecting information from a vCenter Server instance and exporting it to an Excel workbook, which reinforces that the output is a collected report rather than continuous monitoring data. Nutanix sizing and design work should therefore validate the snapshot with appropriate utilization evidence and workload requirements before finalizing node and capacity recommendations. See the [RVTools product page](#) and the [Nutanix Support Portal documentation page](#) for product and design documentation resources.

QUESTION NO: 7

Which two items are considered by Sizer when calculating the required cluster resources with only the Raw input workload type for a multi-VM workload? (Choose two)

- A. The vCPU:pCore ratio applied for specific VMs
- B. Large VM NUMA requirements
- C. The required RAM for the CVM
- D. The vCPU:pCore ratio applied to all VMs

ANSWER: C D

Explanation:

For a multi-VM workload entered using only the Raw input type, Nutanix Sizer uses aggregate workload values rather than characteristics of individual virtual machines. It applies the selected vCPU:pCore ratio uniformly to the aggregate VM workload. This ratio is a key CPU-sizing input because it determines how much physical-core capacity is needed to support the specified total vCPU demand, subject to the selected oversubscription policy.

Sizer also includes Controller VM resource overhead in its cluster calculation. The required RAM for the CVM is reserved as part of the platform footprint on each node, so it affects the usable memory available for user workloads and, consequently, the number and configuration of nodes Sizer recommends. Including this overhead ensures the design has sufficient capacity for both the customer workload and Nutanix's distributed storage and control-plane services.

Because Raw input represents a consolidated multi-VM workload, Sizer does not have per-VM placement details from which to model individual VM-specific CPU ratios or large-VM NUMA placement requirements. For additional Sizer information, see the [Nutanix Sizer page](#) and the [Nutanix Support Portal documentation](#).

QUESTION NO: 8

An SE is in the discovery phase with a new customer. The customer is frustrated with the lack of transparency into their environment with the current monitoring tools. The customer has difficulty correlating the degradation of the end user experience with point-in-time events.

Which Prism feature should the SE demonstrate to show relevant value?

- A. Analysis page with custom charts and event reporting
- B. A specific virtual machine's Performance Summary
- C. Event tab on the Alerts page
- D. A specific virtual machine's virtual disk Additional Stats

ANSWER: A

Explanation:

Analysis page with custom charts and event reporting is the most relevant Prism capability because it gives the customer a consolidated, time-correlated view of infrastructure behavior and operational events. An SE can use Prism's analysis capabilities to create charts for the metrics that matter to the reported end-user issue, such as virtual machine, host, storage, CPU, memory, latency, or IOPS performance. Viewing those metric trends over a selected time range makes it possible to identify exactly when service degradation began and determine whether it aligns with an event, alert, configuration change, or other environmental condition. This directly addresses the customer's need for transparency rather than presenting isolated, point-in-time measurements. Custom charts also support a discovery conversation because the SE can tailor the analysis to the customer's applications, workloads, and stated user-impact symptoms, turning raw operational data into an understandable timeline. Nutanix documents Prism as the management interface for monitoring cluster and VM performance, and its reporting and analysis functions help administrators visualize performance trends and investigate operational conditions. See the [Nutanix Prism Central Guide: Monitoring Performance](#) and the [Nutanix Prism overview](#).

QUESTION NO: 9

An administrator has an existing cluster and needs to improve performance Which two items are required to achieve even greater performance?

- A. Workloads are multi-threaded
- B. VMs have one vCPU configured
- C. VMs have multi-queue disable
- D. VMs have more than one vCPU

ANSWER: A D

Explanation:

"Workloads are multi-threaded" and "VMs have more than one vCPU" are required together to obtain additional performance from virtual CPU resources. A VM can execute work concurrently across multiple virtual CPUs only when the guest operating system and the application have independent threads that can run in parallel. Multi-threading allows the application to divide processing into concurrent execution paths; assigning more than one vCPU gives those paths CPU resources on which to execute simultaneously. This can reduce processing time and improve throughput when sufficient physical CPU capacity is available on the cluster.

In Nutanix AHV, vCPU sizing should reflect the workload's demonstrated parallelism and demand. The administrator should therefore verify that the application benefits from multiple execution threads before increasing the VM's vCPU count, then monitor CPU usage and latency to validate the result. This approach aligns VM configuration with the application's ability to consume added compute resources and supports effective cluster performance planning. Nutanix documentation for AHV administration and VM resource configuration is available at [Nutanix Documentation](#), and Nutanix's virtualization guidance is available at [Nutanix Virtualization](#).

QUESTION NO: 10

Which two items represent measurable POC success criteria? (Choose two.)

- A. Tested VDI Broker
- B. Deployed 100 desktops on AHV using Citrix MCS plugin
- C. Restored VMs from Protection Domain
- D. Compare Nutanix to existing infrastructure
- E. Demonstrated Prism

ANSWER: B C

Explanation:

“Deployed 100 desktops on AHV using Citrix MCS plugin” is a measurable POC success criterion because it defines a specific, countable workload, identifies the target hypervisor, and requires validation of the Citrix Machine Creation Services integration. The POC team can objectively verify whether all 100 desktops were provisioned and function as intended, making the outcome repeatable and unambiguous. AHV is Nutanix’s integrated virtualization platform and is intended to support virtualized workloads without requiring a separate hypervisor license; see [Nutanix AHV](#).

“Restored VMs from Protection Domain” is also measurable because it requires an observable recovery operation to complete successfully. A successful test can be documented through the selected recovery point, restored virtual-machine state, recovery time, and validation that the recovered workload starts and provides its expected services. Nutanix data-protection workflows use protection domains and recovery capabilities to safeguard and recover application data, so performing a VM restoration directly demonstrates a concrete resilience requirement. Nutanix describes its disaster-recovery and recovery capabilities at [Nutanix Disaster Recovery](#). Together, these criteria test both a defined end-user workload deployment outcome and a verifiable data-protection recovery outcome.

QUESTION NO: 11

A customer needs to increase availability in their Nutanix environment by enabling rack awareness. The current configuration consists of : 1 RF2 cluster consisting of 7 identical NX-3170 1U nodes , All Storage efficiencies are disabled , Which two changes should the systems engineer propose to meet the customer requirements ? (Choose two)

- A. Add 1 additional NX-3170 1U node.
- B. Add 2 additional NX-3170 1U nodes.
- C. Distribute nodes across 2 separate racks.
- D. Distribute nodes across 3 separate racks.

ANSWER: B D

Explanation:

For an RF2 Nutanix cluster, rack awareness provides an additional failure-domain boundary by ensuring that replica placement can span separate racks. To enable rack awareness with the desired resilience, the cluster should be arranged across three racks, with sufficient nodes in each rack to maintain availability and placement requirements if an entire rack becomes unavailable. A seven-node cluster cannot be distributed evenly as three nodes per rack. Adding two identical NX-3170 1U nodes creates a nine-node cluster, allowing an even three-node-per-rack layout across three racks.

Therefore, the appropriate proposal is to add two additional NX-3170 1U nodes and distribute the resulting nine nodes across three separate racks. This layout gives Nutanix the physical topology needed to treat each rack as a distinct fault domain and place data replicas accordingly, improving resilience against a rack-level failure while retaining RF2 data protection. Storage-efficiency settings do not change the node-count and rack-distribution prerequisites for this topology. See Nutanix documentation on [configuring rack awareness](#) and the [Nutanix architecture overview](#).

QUESTION NO: 12

prospect uses Vmware, Oracle, and Microsoft products for a critical application. The prospect prefers an all flash array because they have significantly invested in servers for their Oracle environment to save on licensing costs.

which feature should be discussed to address this objection?

- A. AHV
- B. Hyper-V support
- C. ABS

D. All flash nodes

ANSWER: C

Explanation:

ABS is the appropriate feature to discuss because Acropolis Block Services enables a Nutanix cluster to present storage as iSCSI block storage to external, non-Nutanix compute hosts. This lets the prospect retain the existing Oracle server estate, which is important when Oracle licensing has been optimized around a specific number of physical servers, while still consuming Nutanix all-flash storage services. The Oracle servers do not need to be migrated onto Nutanix nodes merely to use the storage; they can access provisioned block-storage volumes across the network. This approach also supports an environment containing VMware, Oracle, and Microsoft workloads, allowing storage modernization without forcing an immediate hypervisor or server-platform change for the critical application. Nutanix describes its unified storage offering as providing block storage alongside other data services, and its product portfolio is designed to support flexible workload placement and external consumption models. For product context, see [Nutanix Unified Storage](#) and [Nutanix Cloud Infrastructure](#).

QUESTION NO: 13

An SE builds a sizer scenario for a prospect's remote site running on an NX-1365S-G5.

Using only the sizer BOM, which two actions can be performed? (Choose two.)

- A. Review budgetary pricing
- B. Review workload sizing details
- C. Review power consumption requirements
- D. Review applicable discounts from deal registration
- E. Reference the HCL for additional SKUs relevant to the configuration

ANSWER: B C

Explanation:

The Sizer bill of materials provides a deployable hardware configuration derived from the scenario's workload inputs. Therefore, it can be used to review workload sizing details, including the resource requirements and the selected node configuration needed to satisfy the modeled CPU, memory, storage, and resiliency needs. This allows an SE to validate that the proposed NX platform configuration aligns with the requirements captured in the Sizer scenario.

The bill of materials can also be used to review power consumption requirements. Hardware BOM output identifies the platform and component configuration on which power-related planning is based, enabling the SE to use the documented platform power specifications when validating remote-site rack, UPS, and electrical capacity. Power planning is especially relevant for a remote deployment, where physical infrastructure capacity may be constrained. Nutanix Sizer is the Nutanix sizing and configuration tool used to create these scenario-based infrastructure recommendations; see [Nutanix Sizer](#). For platform-level hardware specifications and deployment planning information, consult the [Nutanix hardware platform documentation](#).

QUESTION NO: 14

An SE is performing workload discovery with an SMB prospect running SQL Server. The databases is less than 1 TB, but its performance and uptime are critical to the prospect's business operations. Sizer suggests using 1065S nodes.

Which aspects of the NX-1065S-G5 would cause an SE to filter out that model from Sizer's

Automatic configuration? (choose two.)

- A. Single Power Supply
- B. Limited to 4 nodes
- C. 1Gb Networking
- D. Single HDD
- E. Single Socket Processor

ANSWER: C E

Explanation:

1Gb Networking and **Single Socket Processor** are the relevant characteristics for excluding the NX-1065S-G5 from an automatically generated configuration when the SQL Server workload has business-critical performance and availability requirements. Although the database capacity is modest, sizing must account for sustained transaction activity, latency sensitivity, growth, and the infrastructure headroom needed during failures or maintenance. A 1 Gb network interface offers substantially less bandwidth than modern data-center connectivity and can become a constraint for VM I/O, replication, backup traffic, and cluster resiliency operations. A single-socket processor design also provides less available compute and scaling headroom for a critical SQL Server deployment than platforms intended for more demanding production workloads. The SE should therefore use Sizer filtering to focus on platforms that better match the required performance and uptime profile, rather than selecting solely based on the database's current capacity. Nutanix describes its platform portfolio and workload-oriented hardware selection considerations at [Nutanix Hardware Platforms](#). Additional SQL Server deployment guidance is available at [SQL Server on Nutanix Best Practices](#).

QUESTION NO: 15

A prospect is experiencing disk performance issues in their database environment. The prospect's databases currently run on physical servers and an iSCSI SAN with spinning disks. Which two Nutanix features should an SE discuss with this prospect to address their disk performance issues? (Choose two)

- A. Snapshots
- B. Data Locality
- C. De-duplication
- D. Erasure Coding
- E. All Flash Nodes

ANSWER: B E

Explanation:

Data Locality and All Flash Nodes directly address the latency and I/O limitations commonly encountered when databases run on physical servers backed by a traditional spinning-disk SAN. Nutanix Data Locality keeps the active data for a virtual machine local to the node where that VM is running whenever possible. This reduces reliance on traversing a storage network for reads and writes, lowering latency and avoiding SAN-related bottlenecks. Nutanix's distributed architecture can also use replicas on other nodes to maintain availability while preserving efficient access to local data.

All Flash Nodes provide the media performance needed for latency-sensitive database workloads. Replacing spinning disks with SSD- or NVMe-based capacity enables substantially higher IOPS and lower response times, especially for random reads and writes that are typical of transactional database activity. Combined with Nutanix's software-defined distributed storage fabric, all-flash hardware gives the prospect a clear path away from the constraints of the existing iSCSI SAN while supporting predictable database performance at scale. Nutanix discusses the importance of local data access in its [Data Locality overview](#) and describes flash-based platform capabilities in its [Cloud Infrastructure product information](#).

QUESTION NO: 16

A prospective customer is evaluating solutions to support a project. The systems engineer proposes a hybrid model as recommended bySizer. The prospect is concerned about the performance of a hybrid configuration and is evaluating an all flash array. Which two Nutanix features should be discussed to address the concerns of the prospective customer? (Choose two)

- A. AHV Turbo Mode
- B. Data Locality
- C. Automatic Data Tiering
- D. Acropolis Dynamic Scheduler

ANSWER: B C

Explanation:


Data Locality and Automatic Data Tiering directly address the performance concerns associated with a correctly sized Nutanix hybrid cluster. Data Locality places a virtual machine's active data on the Controller VM of the node where that virtual machine is running whenever possible. This minimizes unnecessary network hops and enables reads and writes to be serviced locally, which reduces latency and makes effective use of the distributed architecture.

Automatic Data Tiering optimizes the use of the flash and HDD tiers in a hybrid node. Nutanix continuously identifies frequently accessed, or "hot," data and promotes it to the SSD tier, while less-active data can reside on the capacity tier. As a result, applications receive flash-like performance for their active working set without requiring that all stored data reside on all-flash media. Together, these capabilities explain why a Sizer-recommended hybrid configuration can provide strong application performance while retaining the cost efficiency of HDD capacity.

For additional technical background, see the [Nutanix Bible](#) discussion of the Nutanix Distributed Storage Fabric and Nutanix's [product documentation portal](#).

QUESTION NO: 17

A two-node ROBO cluster is configured with a witness VM.



What happens when Node A goes down?

- A. The- cluster becomes unavailable and goes into read-only mode.
- B. All operations and services on the Node B are shut down and go into a waiting state.
- C. The cluster is unaffected and no administrator intervention is required.
- D. Node B sends a leadership request to the Witness VM and goes into single-node mode.

ANSWER: D

Explanation:

Node B sends a leadership request to the Witness VM and goes into single-node mode. In a two-node Nutanix ROBO deployment, the Witness VM supplies the independent third-party vote needed to maintain cluster quorum when one physical node becomes unavailable. After Node A fails, Node B communicates with the witness and requests leadership. Once that request is granted, Node B can continue operating as the surviving member of the cluster rather than treating the situation as an unconfirmed split-brain condition.

This design lets the surviving node keep providing cluster services and allows workloads that can run on the remaining node to continue, subject to the resources and protection configuration available at that site. The witness does not host user workloads or data; its purpose is quorum arbitration. When Node A returns, it can rejoin the cluster and synchronize as required. Nutanix documents the witness as the quorum mechanism specifically used to enable continued availability for two-node clusters during a single-node failure. See the Nutanix [Cluster Witness VM documentation](#) and the [Nutanix product documentation portal](#).

QUESTION NO: 18

An administrator needs application-consistent snapshots of a Windows virtual machine running Microsoft SQL Server on AHV.

Which two requirements must be met to use the Nutanix VSS-based snapshot workflow? (Choose two.)

- A. Install Nutanix Guest Tools in the Windows guest
- B. Enable the Nutanix VSS Snapshot Service
- C. Install VMware Tools in the Windows guest
- D. Configure the VM to use a physical RDM disk
- E. Disable the Windows Volume Shadow Copy Service

ANSWER: A B

Explanation:

Nutanix Guest Tools must be installed in the Windows guest because it provides the guest integration required for Nutanix to coordinate application-consistent snapshot operations. The Nutanix VSS Snapshot Service must also be enabled and operational so that the guest can quiesce VSS-aware applications before the snapshot is taken.

Without guest-based VSS coordination, a snapshot can still represent a crash-consistent recovery point, but it does not provide the application-consistent behavior required for this SQL Server workload. See the [Nutanix Documentation portal](#) for current Nutanix Guest Tools and data-protection configuration guidance.

QUESTION NO: 19

Which tool could be used to perform a sizing discovery in a general server virtualization environment running both VMware ESXi and Nutanix AHV?

- A. MAP
- B. Collector
- C. AWR
- D. RVTools

ANSWER: B

Explanation:

Collector is the appropriate tool because Nutanix Collector is designed to gather infrastructure inventory, utilization, and performance data that can be used for a Nutanix sizing assessment. It supports discovery in heterogeneous virtualization environments, including VMware vSphere/ESXi and Nutanix AHV. The collected information provides the inputs needed to understand existing CPU, memory, storage capacity, storage performance, and virtual-machine characteristics before designing a target Nutanix solution.

For a mixed environment, using a discovery mechanism that can collect data across both hypervisor platforms is important so that the sizing exercise reflects the complete workload estate rather than only one portion of it. Collector exports the gathered data for use with Nutanix sizing processes and tools, helping the systems engineer develop a configuration that is aligned with observed workload requirements and expected growth. Nutanix positions its sizing resources around collecting accurate workload data before creating a solution recommendation. See the [Nutanix Sizer overview](#) and the [Nutanix Collector Guide](#) for product and collection guidance.

QUESTION NO: 20

An administrator is creating a Prism Central data protection policy for a group of business-critical VMs.

Which two functions can be configured through the protection policy? (Choose two.)

- A. Snapshot or replication schedule
- B. Recovery point retention settings
- C. Creation of a storage container
- D. Assignment of a VM VLAN
- E. Configuration of physical switch ports

ANSWER: A B

Explanation:

A protection policy defines the schedule used to create recovery points and the retention behavior for those recovery points. This allows the administrator to align VM protection with the required recovery point objective and the period for which snapshots must remain available.

Protection policies can be applied to workloads through category-based assignments in Prism Central. The policy does not create storage containers or configure VM network connectivity; those are separate storage and networking administration tasks. See the [Nutanix Data Protection](#) overview and the [Nutanix Documentation portal](#).

QUESTION NO: 21

Which function does the Curator service perform?

- A. Schedules snapshots
- B. Start and stops other cluster services
- C. Stores cluster configuration
- D. Distributes storage optimization tasks

ANSWER: D

Explanation:

The correct answer is **Distributes storage optimization tasks**. In a Nutanix cluster, Curator is the distributed service responsible for managing and scheduling background storage-maintenance and optimization operations across the cluster. It

coordinates tasks such as data rebalancing when capacity or node membership changes, garbage collection of no-longer-needed data, disk cleanup, and other data-placement or storage-efficiency activities. Rather than performing all work from a single centralized process, Curator distributes work to the appropriate Controller VMs and disks, allowing these operations to scale with the cluster while minimizing disruption to foreground virtual-machine I/O. Curator runs periodic scans and creates work items that are executed throughout the Nutanix Distributed Storage Fabric, helping ensure that data remains appropriately balanced and storage resources are maintained efficiently. This responsibility is a core part of the platform's self-managing, distributed architecture. Nutanix describes Curator as the cluster component that performs background data-management activities in its architecture documentation. For additional details, see the [Nutanix Prism Web Console Guide: Cluster Services](#) and the [Nutanix Bible](#).

QUESTION NO: 22

In the event of a primary datacenter failure, a customer wants to power on virtual machines in a secondary datacenter with an RPO and RTO of no more than 15 minutes. Upon further discussions with the customer, the systems engineer (SE) finds that round trip latency between the 2 sites can reach up to 6ms. Which Nutanix feature should the SE recommend?

- A. Metro Availability
- B. Rack Awareness
- C. NearSync
- D. Fault Tolerance FT1

ANSWER: C

Explanation:

NearSync is the appropriate recommendation because it is Nutanix disaster-recovery replication designed for workloads that need a very small recovery point objective without requiring synchronous, zero-data-loss replication between sites. NearSync can replicate protected virtual machines at intervals as low as one minute, so it can readily satisfy an RPO of 15 minutes. In a site-failure event, the recovery workflow brings the replicated workloads online at the recovery site; with an appropriately sized recovery environment and a tested recovery plan, this supports the stated 15-minute RTO target.

The reported round-trip latency can reach 6 ms, which makes an asynchronous NearSync design especially suitable for the requirement. NearSync is intended to provide low-RPO protection across separate locations while avoiding the synchronous intersite write-acknowledgement dependency associated with metro-style availability architectures. The SE should also validate available bandwidth, replication capacity, boot ordering, network mappings, and recovery runbook timing during a recovery test, since these operational factors determine whether the end-to-end RTO is consistently achieved. Nutanix describes its disaster recovery and orchestration capabilities at [Nutanix Leap](#); implementation documentation is available through the [Nutanix Support Portal documentation](#).

QUESTION NO: 23

A prospect provides sizing details for a new application that will require 15TiB of usable storage. The prospect needs to plan for an additional 20% storage growth as part of the initial purchase. Including capacity to fully re-protect in the event of a node failure, which configuration provides minimal cost while meeting the capacity requirement?

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Cannot be determined from the information provided; the node configurations and their usable/raw capacities are missing.

ANSWER: E

Explanation:

Cannot be determined from the information provided; the node configurations and their usable/raw capacities are missing. The requirement first grows from 15 TiB to 18 TiB of usable capacity by including the 20% planned growth. However, selecting a minimal-cost Nutanix configuration that can fully re-protect after a node failure requires more than this usable-capacity figure. The calculation must account for the configured redundancy level, the number of nodes, each node's storage contribution, metadata and system-reserved capacity, and sufficient free capacity on the surviving nodes to create replacement replicas after a failed node is removed. The required raw-capacity reserve differs according to the cluster's node count and storage layout; therefore, no specific hardware configuration can be selected without the details that were replaced by generic labels. A valid answer would need to identify the actual node count and per-node capacity, then verify that the surviving cluster has enough capacity for the protected 18 TiB workload and re-protection. Nutanix storage resiliency and capacity behavior should be validated using the applicable product documentation and sizing tools: [Nutanix Documentation](#) and [Nutanix Software-Defined Storage](#).

QUESTION NO: 24

Refer to Exhibit.



A systems engineer is creating a configuration in Sizer for a customer that is building an app for use as an eBooks repository the app requires the following two workloads:

1. 50 generic web server VMs, each with 2 vCPUs on a 4:1 vCPU ratio, 8 GB of RAM, and 100GB of storage each
2. A fileserver workload with 80TB of usable storage available, which contains mostly PDF's and pre- compressed data.

Any reduction in performance of the web servers must be avoided. After entering the workload with default options, the Sizer outputs the configuration shown in the exhibit.

Which adjustment does the systems engineer need to make to the scenario to meet the performance requirements?

- A. Switch to manual and change the CPUs of the node configuration to a dual 10 core model
- B. Switch to manual sizing and increase the RAM of the node configuration to 384 GB
- C. increase the vCPU:CPU ratio of the Webserver workload to 6:1
- D. increase the RAM of the web server workload to 10 GB per VM.

ANSWER: A

Explanation:

Switch to manual and change the CPUs of the node configuration to a dual 10 core model. The web tier requires 50 virtual machines with 2 vCPUs each, for a total of 100 vCPUs. Because the stated 4:1 vCPU-to-physical-CPU ratio is a performance requirement rather than merely a consolidation preference, Sizer must have sufficient physical CPU capacity in the selected node configuration to honor that ratio. Moving from the default CPU selection to dual 10-core processors increases the available core count per node and enables the proposed configuration to satisfy the web-server CPU

requirement without requiring additional CPU oversubscription. This directly addresses the customer's requirement that web-server performance not be reduced while retaining the specified workload characteristics. Nutanix Sizer is intended to model workload compute, memory, and storage requirements and produce an appropriately configured Nutanix solution; manual configuration selection is appropriate when the default hardware selection does not meet a workload constraint. See the [Nutanix Sizer product page](#) and the [Nutanix Support Portal documentation library](#) for Sizer documentation and sizing guidance.

QUESTION NO: 25

During a meeting with a fortune 100 prospect, an SE achieves technical buy-in from the architecture team. The prospect shares that their management recently attended an executive briefing with Cisco and was impressed by the vision presented.

Management is pushing the arching the architecture tea to adopt Cisco as their HCI standard based on a beief that Nutanix is not a major player.

Which resource should an SE highlight?

- A. Gartner Magic Quadrant
- B. ESG Performance Analysis
- C. Net Promoter Score
- D. Reference architectures

ANSWER: A

Explanation:

Gartner Magic Quadrant is the most appropriate resource because the stated objection is an executive-level perception that Nutanix is not a significant HCI market participant. A Gartner Magic Quadrant provides independent third-party market research that evaluates vendors against Gartner's defined criteria, including their ability to execute and completeness of vision. This gives the SE a credible, executive-friendly way to demonstrate Nutanix's established presence and strategic position in the hyperconverged infrastructure market, rather than relying solely on vendor claims or technical validation.

For this conversation, the resource supports the architecture team's technical confidence with analyst validation that management can readily understand and use in its standardization decision. Nutanix was named a Leader in Gartner's 2020 Magic Quadrant for Hyperconverged Infrastructure Software, reflecting its market standing in that assessment. The SE can use the report positioning and accompanying Gartner methodology to reframe the discussion around independent market recognition, enterprise adoption, and vision. See Nutanix's announcement at [Nutanix: Leader in the Gartner Magic Quadrant for HCI](#) and Gartner's overview of [Magic Quadrants research methodology](#).