

Nutanix Certified Professional Database Automation () v6.5

Nutanix NCP-DB

Version Demo

Total Demo Questions: 32

Total Premium Questions: 328

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Database Automation Concepts and Architecture	136
Topic 2, Database Server VM Management	40
Topic 3, Database Management	53
Topic 4, Database Operations	99
Total	328

QUESTION NO: 1

Which step 's optional when completing the NDE Wizard?

- A. Network Segmentation
- B. NDB Cluster Name
- C. Storage Container
- D. Services

ANSWER: A

Explanation:

Network Segmentation is the optional step when completing the Nutanix Deployment Engine (NDE) wizard. NDE can deploy a functional Nutanix cluster using the required cluster identity, storage, and core service configuration information. Network segmentation is used only when the deployment needs a more advanced network design, such as separating management, storage, virtual-machine, or other traffic onto distinct VLANs or network paths. This capability is valuable for environments with security, compliance, traffic-isolation, or organizational network requirements, but it is not necessary for a basic deployment using a shared network configuration. Therefore, an administrator can complete deployment without configuring segmentation when the intended topology does not require it. The NDE documentation describes the deployment workflow and the network configuration choices available while creating a cluster; the Nutanix documentation portal is the authoritative source for the version-specific Deployment Engine guide. See the [Nutanix Deployment Engine Guide](#) and the [Nutanix Support Portal documentation library](#) for the applicable AOS and NDE release documentation.

QUESTION NO: 2

An operations team wants to use the NDB Alerts Dashboard to show that one administrator has taken ownership of an alert and later record that the condition has been remediated.

Which two alert status values support this workflow? (Choose two.)

- A. Data Resiliency
- B. Acknowledged
- C. Auto
- D. Resolved

ANSWER: B D

Explanation:

Acknowledged indicates that an alert has been reviewed and is being handled. Resolved records that the underlying condition has been addressed. These are alert lifecycle statuses in NDB; Data Resiliency and Auto are not statuses that an administrator applies to an alert.

QUESTION NO: 3

An administrator is configuring NDB to send alert notifications through the organisation's SMTP relay. The mail team requires encrypted transport and confirms that the relay supports the encryption choices exposed by NDB.

Which two SMTP encryption methods can be selected? (Choose two.)

- A. TLS
- B. S/MIME

C. DKIM

D. SSL

ANSWER: A D

Explanation:

TLS and **SSL** are the SMTP encryption methods available in NDB. The selected method, SMTP port, host, and credentials must match the configuration of the mail relay.

S/MIME protects email message content and signing, while DKIM is a domain-level message-authentication mechanism. Neither is an SMTP transport-encryption selection in NDB.

QUESTION NO: 4

What two reference resources are available for Nutanix Era APIs? (Choose two.)

A. Nutanix Support Portal

B. Nutanix.dev API Reference

C. Era API Explorer

D. Era GitHub API Repository

ANSWER: B C

Explanation:

Nutanix.dev API Reference and **Era API Explorer** are the two available reference resources for Nutanix Era APIs. The Nutanix.dev API Reference is the published developer-facing catalog for Nutanix REST APIs. It provides endpoint-level information needed to build integrations, including request methods, resource paths, parameters, request bodies, response schemas, and authentication details. This makes it the primary reference when developers need to understand the supported programmatic operations for Era, now known as Nutanix Database Service (NDB).

The Era API Explorer complements the static reference by providing an interactive interface for discovering the Era API surface and working with API operations directly. It enables administrators and developers to inspect available endpoints and their inputs in the context of an Era deployment, which is particularly useful for validating requests and understanding the API workflow before incorporating calls into automation.

Used together, these resources support both API design and practical API exploration: the API Reference supplies the formal contract, while API Explorer provides a hands-on view of the exposed services. Nutanix maintains its API references at [Nutanix.dev API References](#); NDB documentation and product resources are available through the [Nutanix Database Service User Guide](#).

QUESTION NO: 5

Within NDB, an administrator needs to register a PostgreSQL database server VM with a database.

Where should the database data and log files be located in order to successfully complete this task?

A. On a separate mount point

B. On the same mount point

C. On the same disk as the binary installation

D. On the same disk as the operating system

ANSWER: A

Explanation:

On a separate mount point is correct because Nutanix Database Service (NDB) registration requires the PostgreSQL database data and log location to be isolated from the operating-system and database-binary filesystems. During registration, NDB discovers the database installation and its storage layout, then uses the dedicated database mount point to identify and manage the database consistently. A separate mount point gives NDB a clear boundary for database data, allowing it to perform supported database-management operations and maintain an unambiguous association between the registered database server VM and its PostgreSQL database.

This layout also follows the operational design expected for NDB-managed database servers: operating-system files, PostgreSQL software binaries, and database data have distinct storage roles. Keeping the data and logs on their own mount point makes capacity, lifecycle, and data-management handling more predictable when the VM is brought under NDB control. The administrator should therefore ensure that the PostgreSQL data directory and log files reside on a dedicated mounted filesystem before initiating registration. Refer to the Nutanix documentation portal for the NDB release documentation and database-server registration guidance: [Nutanix Database Service for NDB documentation](#) and [Nutanix Database Service overview](#).

QUESTION NO: 6

An NDB administrator is looking to free up some space consumption in the container that NDB uses to manage a database landscape.

Where are two of the best locations within NDB that the administrator could view to see where the most storage is being consumed, in order to research for reclamation? (Choose two.)

- A. NDB Dashboard and view the Clone Data Age in days.
- B. View the Source Database grid to determine the size and status of registered databases.
- C. View the Clone Database table to determine the Size and Last Refresh of clone databases.
- D. View the Snapshot table in Time Machine to determine Reclaimable Space and Age of snapshots that have been taken.

ANSWER: C D

Explanation:

Viewing the Clone Database table to determine the Size and Last Refresh of clone databases is a direct way to identify clone instances that may be consuming substantial storage. The size information helps prioritize the highest-impact reclamation opportunities, while the last-refresh detail helps the administrator assess whether a clone is still actively maintained or is a candidate for cleanup. NDB database clones are created from Time Machine restore points and are managed as database resources, so their lifecycle should be reviewed as part of storage housekeeping.

Viewing the Snapshot table in Time Machine to determine Reclaimable Space and Age of snapshots that have been taken provides the most actionable snapshot-level information for reclaiming container capacity. Reclaimable Space identifies the potential capacity benefit associated with removing eligible snapshots, and age helps administrators apply retention and operational requirements before deleting older restore points. Together, these views let the administrator investigate the two major Time Machine-related consumers—clones and snapshots—using storage and lifecycle data rather than relying only on high-level indicators. Nutanix documents Time Machine as the NDB capability for creating, managing, and restoring database snapshots, including clone operations from those restore points. See the [Nutanix NDB Time Machine documentation](#) and the [Nutanix Database Service product overview](#).

QUESTION NO: 7

An administrator needs to deploy a database resilient cluster engine across three clusters. Which database engine can be used for this scenario?

- A. MariaDB

- B. MySQL
- C. PostgreSQL
- D. Oracle

ANSWER: C

Explanation:

PostgreSQL is the supported choice for this resilient database-cluster scenario. Nutanix Database Service (NDB) can deploy and manage a PostgreSQL high-availability topology consisting of a primary database server and replica database servers. Placing the members across three Nutanix clusters provides infrastructure-level fault-domain separation: a failure affecting the cluster hosting one database member does not inherently remove the remaining members from service. NDB automates the database-server provisioning and applies the PostgreSQL configuration needed for replication and role management, reducing the manual work required to establish a consistent resilient deployment.

The underlying PostgreSQL capability is streaming replication, in which a standby continuously receives write-ahead log records from the primary and can be promoted when the primary is unavailable. A three-member design provides a practical resilient topology because it can retain a primary/standby arrangement while distributing database roles across independent Nutanix clusters. NDB uses this supported PostgreSQL architecture to deliver database availability while preserving lifecycle-management and operational controls through its database-management plane. Nutanix describes NDB as providing automated database operations and availability capabilities at [Nutanix Database Service](#). PostgreSQL's replication and standby-promotion behavior is documented in the [PostgreSQL warm standby and streaming replication documentation](#).

QUESTION NO: 8

Which two processes should an administrator take to create a software profile version for a newly-released Oracle ISc patch? (Choose two.)

- A. 1. Download the required patch from Microsoft Support. 2. Apply the patch to the Oracle Home/Grid Home on a cloned VM of the gold image using OPatchAuto. 3. Ensure successful patch application.
- B. 1. Create a new software profile version from the original software profile using the cloned VM where OPatchAuto ran successfully in the Era GUI. 2. Mark the original patch profile as deprecated.
- C. 1. Download the required patch from My Oracle Support. 2. Apply the patch to the Oracle Home/Grid Home on a cloned VM of the gold image using OPatchAuto. 3. Ensure successful patch application.
- D. 1. Create a new software profile version from the original software profile using the cloned VM where OPatchAuto ran successfully in the Era GUI. 2. Mark the original patch profile as hidden.

ANSWER: C D

Explanation:

Creating an Oracle software profile version in Nutanix Database Service begins by obtaining the applicable Oracle patch from My Oracle Support. The patch should then be applied to the Oracle Home or Grid Home on a cloned VM created from the gold image. Using Oracle OPatchAuto is appropriate because it automates patch deployment for supported Oracle environments. Confirming that the patch operation completed successfully is essential before capturing the VM state as a profile version; the resulting software profile must represent a known-good, patched Oracle installation.

After validation, the administrator uses the NDB/Era interface to create a new software profile version from the original software profile, selecting the cloned VM on which OPatchAuto completed successfully. This preserves the profile lineage while making the newly patched binaries available for future database provisioning and maintenance operations. Marking the original patch profile as hidden keeps older versions from being selected for new operations while retaining them for traceability or rollback requirements. Nutanix documents software profile and version management in its [Nutanix Database Service Administration Guide](#); Oracle documents the patching utility in the [Oracle OPatch and OPatchAuto User's Guide](#).

QUESTION NO: 9

A request is received to refresh a database clone from a new manual snapshot. When the administrator attempts to create the new snapshot from the Time Machine, it is in a Frozen state.

What causes a Time Machine to enter the Frozen state and what are the administrator's options to complete the request?

- A.** Cause: The Time Machine enters a Frozen state when the database is de-registered without removing the Time Machine. Resolution: The Time Machine must be paused, then the database re-registered with a new Time Machine and a new clone created.
- B.** Cause: The Time Machine enters a Frozen state when the database is de-registered without removing the Time Machine. Resolution: To complete the request the database must be re-registered in NDB, create a new Time Machine and create a new clone.
- C.** Cause: The Time Machine enters a Frozen state when the database is de-registered without removing the Time Machine. Resolution: To complete the request the database must be re-registered in NDB, this thaws the Time Machine and it resumes operation so the clone can be refreshed.
- D.** Cause: The Time Machine enters a Frozen state after too many snapshot or log catchup failures. Resolution: Remediate the snapshot or log catchup failures, this thaws the Time Machine, allowing normal operations to continue and the request can be completed.

ANSWER: C

Explanation:

Cause: The Time Machine enters a Frozen state when the database is de-registered without removing the Time Machine. Resolution: To complete the request the database must be re-registered in NDB, this thaws the Time Machine and it resumes operation so the clone can be refreshed. This is correct because a Time Machine remains associated with its protected source database. If that database is de-registered while its Time Machine is retained, NDB freezes the Time Machine because it no longer has a registered source from which to safely coordinate protection operations. Re-registering the same database restores that association and enables NDB to thaw and resume the existing Time Machine rather than requiring a replacement Time Machine or clone. Once the Time Machine is operational again, the administrator can create the required manual snapshot and use that snapshot as the refresh point for the existing clone. This preserves the established Time Machine history and supports the requested refresh workflow with the least disruptive recovery action. Time Machine states and recovery actions are part of NDB data protection administration; see the [Nutanix Database Service Administration Guide](#) and Nutanix's [NDB learning resources](#).

QUESTION NO: 10

Where are two locations an administrator would find the size of a Source Database? (Choose two.)

- A.** On the Era Dashboard Database Summary Table under Source Database.
- B.** Under the Capacity widget on the Database Server VMS > List database summary page.
- C.** On the Databases > Sources table.
- D.** On the Profiles > Database Parameters list for the assigned DB Parameter Profile.

ANSWER: A C

Explanation:

The source database size is available from both the Era Dashboard Database Summary Table under Source Database and the Databases > Sources table. The Dashboard provides a high-level operational overview of managed database resources, and its Database Summary area includes source-database information so an administrator can quickly review database size alongside other inventory details. This is useful when assessing the overall footprint of protected or managed source databases without opening individual database records.

The Databases > Sources table is also an authoritative inventory view for registered source databases. It presents source-level attributes, including size, enabling administrators to compare source databases and identify their capacity consumption directly from the Sources workflow. These two views serve complementary purposes: the dashboard supports at-a-glance monitoring, while the Sources table supports detailed source-database administration and inventory review. Nutanix Database Service was previously branded as Era; therefore, Era references in the interface and training material refer to the Database Service administration experience. Consult the [Nutanix Database Service documentation](#) and the [Nutanix Database Service product page](#) for current interface and feature documentation.

QUESTION NO: 11

Which two network prerequisites are required to deploy an Era instance in production? (Choose two.)

- A. NTP, DNS servers, gateway, and subnet mask must be available and visible to Era.
- B. One static IP address must be available and can be changed using CLI after VM deployment.
- C. A storage container must be available as a prerequisite.
- D. An SMTP server and details must be provided for notifications.

ANSWER: A B

Explanation:

NTP, DNS servers, gateway, and subnet mask must be available and visible to Era because the Era (now Nutanix Database Service) appliance must receive valid network configuration and be able to resolve names, reach required services, and maintain accurate time. DNS supports name resolution for clusters, database servers, and integrated services. A correctly configured gateway and subnet mask provide the appliance with appropriate local-network and routed-network connectivity. Time synchronization through NTP is essential for consistent timestamps, secure communications, logging, and reliable coordination with managed resources.

One static IP address must be available and can be changed using CLI after VM deployment because the Era appliance requires a stable, reachable management address. This address is used to access the appliance and enables ongoing communication between Era, Prism infrastructure, and the database environments that Era manages. A static assignment prevents management connectivity from changing unexpectedly after a lease renewal or network restart. Nutanix deployment guidance treats the appliance network values and an available static address as information that must be planned before production deployment. See the [Nutanix Database Service documentation](#) and the [Nutanix Database Service product page](#).

QUESTION NO: 12

An administrator needs the ability to clone a source database to a point in time.

What is the lowest log frequency in minutes that can be configured in the Era IJL for log catchup operation?

- A. 30
- B. 60
- C. 5
- D. 15

ANSWER: C

Explanation:

5 is the lowest configurable log-catchup frequency for this point-in-time cloning use case in Nutanix Era, now known as Nutanix Database Service. Log catchup collects and applies database transaction-log changes between the protected database copy and the requested recovery or clone time. Configuring the interval at five minutes gives the service the most

granular supported log-collection cadence, which supports a clone that is as close as possible to a selected point in time within the available log history.

This setting matters because a point-in-time clone requires both a suitable database snapshot or baseline copy and the transaction logs produced after that baseline. Database Service uses those logs to advance the cloned database to the desired time. A shorter catchup interval reduces the potential gap between log-collection events and therefore provides the finest recovery-point objective available through this configuration. Administrators should still ensure that retention, log availability, database workload, and infrastructure capacity align with their organization's recovery requirements when choosing the interval.

For Nutanix product and documentation resources, see [Nutanix Database Service](#) and the [Nutanix Documentation Portal](#).

QUESTION NO: 13

Which Era component collects logs and snapshots from the database?

- A. Profiles
- B. Era Database Agent
- C. SLAs
- D. Time Machine

ANSWER: D

Explanation:

Time Machine is the Era (now Nutanix Database Service, or NDB) component that provides copy-data management for protected databases. A Time Machine maintains the recoverable history of a database by coordinating periodic database snapshots together with the database transaction logs needed to restore to a selected point in time. This retained history enables NDB operations such as restoring a database, provisioning a clone from a specific snapshot, and refreshing a database copy while preserving the consistency required by the supported database engine. The collection and retention behavior is governed through the Time Machine's service-level agreement settings, which define items such as snapshot frequency and log retention. Consequently, when the question asks for the component that collects logs and snapshots, Time Machine is the precise component because it is the object that manages these recovery artifacts and exposes them for point-in-time recovery and copy workflows. Nutanix describes this copy-data-management capability as a core part of its database automation platform; see the [Nutanix Database Service product page](#) and the [Nutanix Database Service documentation portal](#).

QUESTION NO: 14

Which disk is omitted from a Software Profile?

- A. System Page File
- B. SQL Installation
- C. User Database
- D. System Database

ANSWER: A

Explanation:

System Page File is omitted from a Nutanix Database Service (NDB) Software Profile because NDB manages this disk automatically as part of database VM provisioning. A Software Profile defines the database-engine software configuration and the database-related storage layout that is needed to install and operate the selected engine. The page-file requirement, however, is a Windows operating-system concern rather than a database software-layout choice that an administrator needs to model in the profile. NDB creates and sizes the system page-file disk according to the provisioned database server VM's

memory configuration, ensuring the operating system receives an appropriate paging resource without requiring profile-level disk definition. This automatic handling also helps keep Software Profiles reusable across deployments that may use different VM sizes or memory allocations. When building or reviewing a profile for a Microsoft SQL Server database, administrators should therefore expect the database installation and database storage components to be represented as applicable, while the Windows system page-file disk is handled by the NDB provisioning workflow. Nutanix documents Software Profiles as reusable definitions for database software and configuration in the [Nutanix Database Service User Guide](#). Additional NDB documentation and version-specific guidance are available from the [Nutanix Database Service documentation portal](#).

QUESTION NO: 15

When provisioning a new database, which two options are available for the target? (Choose two.)

- A. Provision database to an existing registered
- B. Provision database to an existing container.
- C. Provision databases to a new container.
- D. Provision database to a new database server\cluster.
- E. Provision database to a new database server.

ANSWER: A E

Explanation:

Nutanix Database Service (formerly Era) offers two target choices during database provisioning: an existing registered database server or a new database server. "Provision database to an existing registered" represents the existing-server path: NDB deploys the requested database onto a database server that has already been registered and is managed by NDB. This is appropriate when a suitable database-server VM is already available and has the required engine support, connectivity, and capacity.

"Provision database to a new database server" represents the new-server path. NDB provisions a new database-server VM and then deploys the database onto that managed target. This approach provides a dedicated target when no suitable registered server is available or when isolation is desired. The target selection is therefore about whether the database is deployed to an existing NDB-registered database server or to a newly provisioned database server. Nutanix documentation describes NDB database provisioning and its managed database-server architecture in the [Nutanix Database Service documentation](#); product documentation for the 6.5 release is also available through the [Nutanix Support Portal](#).

QUESTION NO: 16

A database administrator logged into their NDB dashboard notices the Savings over traditional cloning metric in the Data Usage widget.

What does this number represent?

- A. The amount of disk space saved with near-zero byte clones
- B. The amount of disk space saved from compression
- C. The amount of disk space saved with shadow clones
- D. The amount of disk space saved from deduplication

ANSWER: A

Explanation:

The amount of disk space saved with near-zero byte clones represents the storage efficiency achieved when Nutanix Database Service (NDB) provisions database clones without making a full physical copy of the source database data. NDB

uses Nutanix storage cloning capabilities to create copy data management environments efficiently: a clone initially references the source data and consumes additional capacity primarily for blocks that diverge as changes are made. In contrast, a traditional clone requires a separate full-sized copy of the database, immediately consuming capacity comparable to the source. The Data Usage widget's Savings over traditional cloning value therefore expresses the capacity avoided by using NDB's space-efficient clone approach rather than conventional full-copy cloning. This dashboard metric helps administrators understand the tangible storage benefit of NDB automation when database copies are created for development, test, reporting, or similar purposes. It is specifically a comparison of clone-related consumption against the expected consumption of traditional full-copy database cloning. Nutanix describes NDB as providing database lifecycle management and copy data management capabilities, including rapid database cloning and provisioning. See the [Nutanix Database Service product page](#) and the [Nutanix NDB documentation portal](#) for product and documentation resources.

QUESTION NO: 17

An administrator needs to create and refresh clones to a point in time.

NDB allows which two methods for satisfying this requirement? (Choose two.)

- A. Full backups
- B. Snapshots
- C. VM snapshots
- D. Transaction logs

ANSWER: B D

Explanation:

Snapshots and **Transaction logs** are the NDB mechanisms that enable point-in-time clone operations. NDB protects a registered database through a Time Machine, which maintains database-consistent snapshots at configured intervals. A snapshot supplies the baseline copy from which NDB can provision a clone or return a clone to a protected state. This database-aware snapshot is central to providing a usable, consistent copy rather than merely a crash-consistent infrastructure image.

For a requested time that falls between snapshot intervals, NDB uses the database's transaction-log backups to roll the selected snapshot forward to the requested recovery point. The resulting clone reflects the data state at that selected timestamp, subject to the recovery granularity and the available retention window. During a clone refresh, the administrator can therefore select an appropriate Time Machine recovery point, with NDB coordinating the snapshot and log-based recovery workflow for the supported database engine.

This combination supports both rapid copy provisioning from restore points and more precise recovery to a specific moment, which is the purpose of NDB's Time Machine and clone lifecycle capabilities. See the [Nutanix Database Service product overview](#) and the [Nutanix Support Portal documentation](#) for NDB Time Machine, backup, recovery, and clone administration guidance.

QUESTION NO: 18

In NDB, which two Time Machine options are available when creating a database clone? (Choose two.)

- A. Point in time
- B. Restore point
- C. Backup
- D. Snapshot

ANSWER: A D

Explanation:

When an NDB database clone is created from a Time Machine, the available recovery selections are **Point in time** and **Snapshot**. Point in time supports creation of a clone using a specified moment within the Time Machine's recoverable timeline. NDB uses the Time Machine's protected database history to provision the clone at that selected time, enabling users to obtain a usable copy that reflects the database state at the required moment.

Snapshot supports cloning from a particular Time Machine snapshot. A snapshot is a named, discrete recovery state maintained by the Time Machine, so selecting it gives the clone a known and consistent source state. This is useful when a team needs a development, test, reporting, or troubleshooting copy corresponding to a specific protected checkpoint. Both choices are part of the Time Machine-based cloning workflow: one identifies an exact time in the available recovery window, while the other identifies a captured snapshot checkpoint. Nutanix Database Service documentation describes Time Machine as the NDB capability for database protection, recovery, and clone operations. See the [Nutanix Database Service product page](#) and the [Nutanix product documentation portal](#) for NDB Time Machine and cloning guidance.

QUESTION NO: 19

An administrator enables NDB Multi-Cluster on Cluster

A.

Cluster B is then registered with ND

B.

What are the different NDB Service VMs present in each Nutanix cluster?

A. Cluster A: 1 NDB ServerCluster B: 1 NDB Agent

B. Cluster A: 1 NDB AgentCluster B: 1 NDB Server, 1 NDB Agent

C. Cluster A: 1 NDB Server, 1 NDB AgentCluster B: 1 NDB Agent

D. Cluster A: 1 NDB AgentCluster B: 1 NDB Server

ANSWER: C

Explanation:

Cluster A: 1 NDB Server, 1 NDB AgentCluster B: 1 NDB Agent is correct because the cluster on which NDB Multi-Cluster is enabled acts as the primary NDB cluster. NDB deploys its central NDB Server VM on that primary cluster. The NDB Server provides the management-plane services used to administer database operations, policies, provisioning workflows, and the overall NDB environment. An NDB Agent is also present on the primary cluster so that the NDB Server can communicate with and orchestrate database resources hosted there.

When Cluster B is registered to the NDB deployment, it becomes a managed or registered cluster rather than another primary NDB Server location. NDB deploys an NDB Agent on Cluster B. This agent establishes the required connection to the NDB Server on Cluster A and enables the central NDB service to discover, manage, and execute supported database lifecycle tasks for resources on Cluster B. This architecture centralizes NDB management while allowing database workloads to be managed across multiple Nutanix clusters. Nutanix documents this primary-cluster NDB Server and registered-cluster NDB Agent model in its [NDB Administration Guide](#) and [Nutanix Database Service overview](#).

QUESTION NO: 20

Which NDB HA control plane VM is upgraded first during an NDB upgrade?

A. NDB Server

B. Leader NDB API Server

C. Leader HAProxy

D. NDB Agent VMs

ANSWER: B

Explanation:

Leader NDB API Server is upgraded first in the NDB high-availability control-plane upgrade workflow. The NDB API Server is the control-plane endpoint that handles API requests and participates in the coordination required for NDB operations. In an HA deployment, beginning with the leader API Server establishes the upgraded control-plane leadership component before the remaining API Server and HAProxy virtual machines are processed. This ordered procedure is designed to preserve a managed, predictable transition of control-plane services while the NDB software version is updated across the deployment. Administrators should initiate upgrades from the NDB interface and allow each workflow stage to complete, rather than manually changing the prescribed VM order. They should also confirm that the NDB deployment is healthy and that required backups and compatibility prerequisites have been met before starting the update. Nutanix documents NDB lifecycle and upgrade procedures in its product documentation portal; consult the version-specific administration documentation because the supported upgrade paths and prerequisite checks can vary by NDB release. See the [Nutanix NDB documentation portal](#) and Nutanix's [Database Service product page](#) for official NDB information and documentation access.

QUESTION NO: 21

An IT manager is onboarding a new Nutanix systems administrator to manage their NDB deployment. The manager would like to assign the NDB role that provides the minimum amount of permissions to allow provisioning of new databases.

Which role should be assigned to accomplish this task?

- A. Infrastructure administrator
- B. Super administrator
- C. Database infrastructure administrator
- D. Database administrator

ANSWER: D

Explanation:

Database administrator is the appropriate least-privileged NDB role for a user who must provision new databases. NDB role-based access control separates database lifecycle activities from NDB platform and infrastructure administration, allowing permissions to be assigned according to the user's operational responsibility. The Database administrator role is intended for database-focused work, including provisioning database instances from approved database software profiles and using the NDB service to manage the databases assigned to that administrator's scope. This lets a newly onboarded administrator perform the requested provisioning task without requiring broad control of NDB configuration, user and role administration, or underlying infrastructure resources. Applying this role supports the principle of least privilege: the user receives the permissions necessary for database provisioning and ongoing database administration, while higher-level NDB and infrastructure controls remain delegated to the appropriate administrative personnel. Nutanix documents the available NDB roles and their permissions in its administration documentation; see the [Nutanix Database Service Administration Guide—User Roles](#) and the [Nutanix Product Support and Documentation portal](#).

QUESTION NO: 22

A developer requests a new development database environment. The Era administrator identifies that there are insufficient resources on the

development cluster.

Which two profiles should the administrator adjust to get the copy to fit within the resources available? (Choose two.)

- A. Network
- B. Software
- C. Compute

ANSWER: C D

Explanation:

Compute and **Database Parameters** are the appropriate profiles to review when a development copy must fit into a cluster with constrained capacity. A Compute profile defines the infrastructure sizing applied to the database server VM, including its CPU and memory allocation. Selecting or modifying a smaller Compute profile directly reduces the virtual infrastructure resources required by the copied database environment.

A Database Parameters profile supplies database-engine configuration settings for the deployed copy. For development workloads, administrators can use parameters appropriate to a lower-capacity, non-production instance, reducing the database memory footprint and related runtime resource consumption. This allows the database configuration to align with the smaller infrastructure allocation while still creating a usable development environment.

Using these profiles together supports right-sizing at both levels: VM-level capacity through the Compute profile and database-instance resource behavior through the Database Parameters profile. Nutanix Database Service uses profiles to standardize and consistently apply database deployment and copy settings, enabling administrators to provide environments matched to the capacity and purpose of the target cluster. See the [Nutanix Database Service documentation on profiles](#) and the [Nutanix Database Service product overview](#).

QUESTION NO: 23

Which statement is true regarding the NDB alert retention feature?

- A. NDB supports alert retention globally.
- B. NDB supports retention at the individual policy.
- C. NDB supports retention at the grouped policy.
- D. NDB supports retention at the registered cluster.

ANSWER: A

Explanation:

NDB supports alert retention globally. Alert retention is an appliance-wide configuration that controls how long alert records are preserved before NDB removes older entries. Because the setting is global, it establishes one consistent retention period for alerts generated across the NDB environment, helping administrators manage the size of retained alert history and maintain a predictable operational-audit window. Administrators configure this behavior through the NDB administration interface rather than defining a separate retention interval for each database operation, policy grouping, or registered cluster. This design centralizes alert-history lifecycle management: an administrator can adjust the retention setting once and have it apply consistently to the alert data maintained by the NDB appliance. Retention affects the availability of historical alert information for troubleshooting, operational review, and monitoring investigations, so the configured period should balance the organization's diagnostic and audit needs against the storage consumed by accumulated alert records. Nutanix documents NDB administrative configuration and alert-management capabilities in its product documentation. See the [Nutanix Database Service documentation](#) and the [Nutanix documentation portal](#) for the release-specific alert configuration guidance.

QUESTION NO: 24

A legal team wants to audit a quarter end (July 31st) banking transaction event and asks the administrator to create a SQL database from the timeframe.

What action should the administrator perform within Era?

- A. Provision a database and attach the July 31st backup from Time Machine.

- B. Restore a backup from July 31st, Recover the database to bring the database up.
- C. Provision a database, restore the backup from July 31st, and Recover the database to bring the database up.
- D. Create a clone from July 31st snapshot in Time Machine.

ANSWER: D

Explanation:

Create a clone from July 31st snapshot in Time Machine. Nutanix Era, now named Nutanix Database Service, uses a Time Machine to maintain application-consistent database recovery points, including snapshots and, where configured, transaction-log-based recovery capability. Creating a clone from the recovery point representing July 31st provides a separate, usable SQL database copy containing the data as it existed at that selected point in time. This is particularly appropriate for a legal or audit request because it preserves the production database while allowing the legal team to inspect, query, and validate the requested quarter-end transactions in an isolated environment. The clone is storage efficient because it uses Nutanix snapshot technology rather than requiring a complete independent copy of all database data at creation time. It also supports the operational need to provide access to historical data promptly without disrupting the source database or changing its current state. Nutanix documents Time Machine as the mechanism for creating and managing database snapshots and using them to create database clones. See the [Nutanix Database Service User Guide](#) and the [Nutanix Database Service product overview](#).

QUESTION NO: 25

An administrator has been asked to provision a new Oracle single instance database, but cannot proceed with the first step of the wizard. Which statement best describes the current situation?

- A. Time Machine is not configured.
- B. Software profile is missing.
- C. No remote clusters are configured.
- D. A snapshot must be taken on the original DB.

ANSWER: B

Explanation:

Software profile is missing. Nutanix Database Service (NDB) uses a software profile as a required provisioning input for Oracle databases. The profile supplies the database-engine software and version-specific configuration information that NDB needs to present a valid Oracle provisioning workflow. In practical terms, it represents the Oracle database installation media and associated configuration/patch level that NDB will apply when it deploys the database server and creates the requested single-instance database. Therefore, if no eligible Oracle software profile exists for the selected database type and version, the administrator cannot advance through the initial provisioning selections because NDB has no software definition to associate with the new database.

An administrator must first create, import, or otherwise make available an appropriate Oracle software profile, ensuring it is compatible with the target environment and intended Oracle version. After the profile is available, it can be selected in the provisioning wizard and the database deployment can continue. Software profiles are a core NDB lifecycle-management construct, enabling repeatable and standardized database provisioning rather than relying on a manually configured database installation for each request. Nutanix documentation for NDB is available through the [Nutanix NDB documentation portal](#); additional NDB learning material is available at [Nutanix.dev NDB labs](#).

QUESTION NO: 26

NDB can send alerts to which system?

- A. ServiceNow

- B. Email
- C. Microsoft Teams
- D. PagerDuty

ANSWER: B

Explanation:

Email is the native alert-notification destination in Nutanix Database Service (NDB). An NDB administrator configures an SMTP server in the NDB administration settings, including the SMTP host and sender details. NDB can then use that mail service to deliver alert notifications to the configured recipient addresses. This allows database and platform operators to receive operational information, such as warning, failure, and other alert conditions, without needing to remain logged in to the NDB console.

The notification mechanism depends on a reachable SMTP service and appropriately configured mail-routing and authentication settings in the environment. After SMTP configuration is complete, the NDB alerting workflow uses email as the supported built-in channel for distributing those notifications. This aligns with the NDB product documentation's administration guidance for configuring SMTP-based email notifications and with Nutanix's description of NDB as a database-management platform that provides operational visibility and automation. See the [Nutanix Database Service documentation](#) and the [Nutanix Database Service product page](#).

QUESTION NO: 27

Which two profiles are required to provision a new Oracle database through Era? (Choose two.)

- A. Memory
- B. Compute
- C. Network
- D. Storage

ANSWER: B C

Explanation:

Compute and **Network** are required profiles when provisioning an Oracle database with Nutanix Database Service (formerly Era). A compute profile defines the virtual-machine sizing used for the database server, including the allocated vCPU and memory resources. Selecting this profile ensures that the Oracle database VM is created with an approved, repeatable resource configuration appropriate for the intended workload.

A network profile supplies the network configuration that the new database server requires, such as the managed network, IP-address assignment method, DNS settings, gateway, and related connectivity details. NDB uses this profile to connect the provisioned database VM to the required application and management network environment. Together, these profiles provide the infrastructure sizing and connectivity inputs needed for a standardized Oracle provisioning workflow.

NDB uses profiles to make database provisioning consistent and policy-driven across deployments. The available profile types and the information collected during provisioning can vary with the database engine and NDB release, but compute and network configuration are foundational inputs for deploying a database server. See the [Nutanix Database Service documentation](#) and the [Nutanix Database Service product overview](#).

QUESTION NO: 28

Which two options can NDB leverage to refresh a database clone? (Choose two.)

- A. Cerebro logs
- B. Snapshots

C. Transaction logs

D. Templates

ANSWER: B C

Explanation:

Snapshots and **Transaction logs** are the two mechanisms NDB can leverage when refreshing a database clone. A snapshot supplies a crash-consistent, point-in-time copy of the source database data that NDB can use as the baseline for the refresh. This enables the clone to be reset efficiently to a known source state without requiring a complete new database deployment.

Transaction logs provide the changes recorded after that baseline. When applicable to the database engine and the clone's configuration, NDB can apply the relevant log information so that the refreshed clone is brought forward toward the desired state of its source. Using a snapshot baseline together with database transaction logs supports efficient data movement while preserving the database-engine recovery and consistency process.

Clone refresh is therefore based on source database copy and recovery capabilities: NDB uses the snapshot as the data baseline and transaction logs as the incremental change stream. This behavior is part of NDB clone database lifecycle management. See the [Nutanix Database Service User Guide: Clone Database Management](#) and the [Nutanix Database Service product documentation overview](#).

QUESTION NO: 29

An administrator adds a new VLAN for database server VMs on an AHV cluster. The VLAN must be available as a selectable network for NDB provisioning workflows.

Which two actions are required? (Choose two.)

A. Add the subnet and VLAN in Prism Element Network Configuration.

B. Create a new VPC subnet for the database server VMs.

C. Add the VLAN in NDB Networks.

D. Add the network through NDB Multi-Cluster Management.

ANSWER: A C

Explanation:

The subnet and VLAN must first be configured on the Nutanix cluster through Prism Element network configuration. The VLAN must then be added to the NDB network inventory so NDB can use it for database server provisioning. Configuring a VPC or adding a network through Multi-Cluster Management does not register the VLAN for this workflow.

QUESTION NO: 30

Which port should be opened to ensure the NDB server, NDB agent, and NDB database agent can communicate with the registered cluster via REST API calls?

A. 2009

B. 3205

C. 3260

D. 9440

ANSWER: D

Explanation:

9440 is the required port because it is the HTTPS port used by Prism Element for management and REST API access on a Nutanix cluster. NDB components use the registered cluster's Prism Element REST interface to perform management operations needed for database automation, including discovering cluster resources and coordinating infrastructure actions associated with database workloads. Consequently, network connectivity from the NDB server, NDB agent, and NDB database agent to the registered cluster on TCP port 9440 must be permitted by any intervening firewall or security policy. This communication is encrypted with HTTPS and is directed to the cluster's management interface or virtual IP, as configured during cluster registration. Allowing this port supports the control-plane communication NDB requires with the cluster; it is distinct from ports used for storage protocols or internal service communications. Nutanix documents NDB networking requirements and component connectivity in its NDB documentation, while the Nutanix documentation portal provides the applicable version-specific network ports and protocols references. See the [Nutanix Database Service documentation](#) and the [Nutanix Support Portal documentation](#).

QUESTION NO: 31

A clone needs to be created on an existing database server, but there are no Authorized database servers defined to support the cloned database.

In which two ways can Database Server VMs be authorized during a Clone Create operation? (Choose two.)

- A.** Within Time Machine, use the Actions menu and Authorize Database Server VMs to identify which database server VMs to authorize for clones.
- B.** Within the Database Server VM Summary window, update the database sever to be authorized for supporting cloned databases.
- C.** During the Clone creation process, click the plus icon to authorize one or more database server VMs.
- D.** During the database server provisioning process, flag the database server to be authorized for supporting cloned databases.

ANSWER: A C

Explanation:

Nutanix Database Service (NDB) requires database server VMs to be explicitly authorized before they can host clones associated with a Time Machine. This authorization associates eligible existing database server VMs with the Time Machine and makes them available as clone targets while preserving NDB's control over where clone workloads can be placed. An administrator can perform this setup directly from the Time Machine by using its Actions menu and selecting *Authorize Database Server VMs*. That workflow is useful when preparing one or more eligible servers before beginning the clone request.

NDB also supports authorization inline during clone creation. When no suitable authorized server is already listed, the Clone Create workflow provides the plus-icon action to authorize database server VMs without abandoning the operation. The administrator can select one or more appropriate database server VMs, authorize them for the Time Machine, and then continue configuring the clone. Both methods meet the same prerequisite: the selected VM becomes an authorized target for clones created from that Time Machine. This allows clone placement to be controlled at the Time Machine level while keeping the clone workflow efficient. See the [Nutanix Database Service 6.5 documentation](#) and the [Nutanix Product Support portal](#).

QUESTION NO: 32

Which two methods can be used to upgrade Era? (Choose two.)

- A.** Nutanix LCM
- B.** Offline Upgrade
- C.** Rolling Upgrade
- D.** One-click Upgrade

ANSWER: B D

Explanation:

Era, now known as Nutanix Database Service (NDB), supports both connected and disconnected upgrade workflows.

Offline Upgrade is the supported approach when the Era appliance cannot access Nutanix update services directly. An administrator downloads the appropriate upgrade bundle from the Nutanix Support Portal, transfers it to an accessible location, and uploads or applies the bundle through the Era/NDB interface. This method is intended for dark-site and otherwise restricted environments while still allowing the appliance software to be brought to a supported release.

One-click Upgrade refers to the integrated online upgrade workflow available from the Era management interface. When the appliance has the required connectivity, Era can discover the available release, obtain the upgrade package, perform required prechecks, and orchestrate the upgrade from the console. Nutanix documentation identifies online one-click and offline bundle-based upgrades as the available appliance-upgrade paths. Review release-specific prerequisites, compatibility information, and upgrade guidance before beginning either workflow. See the [Nutanix Database Service 6.5 documentation](#) and the [Nutanix Support Portal](#) for software bundles and release documentation.