

IBM Netezza Performance Server V11.x Administrator

IBM C1000-085

Version Demo

Total Demo Questions: 8

Total Premium Questions: 86

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture	22
Topic 2, Installation and Configuration	6
Topic 3, Security	6
Topic 4, Administration	28
Topic 5, Monitoring and Performance	17
Topic 6, Backup and Restore	7
Total	86

QUESTION NO: 1

Which three scripts are included in the IBM Netezza Performance Server Software Support

Tools?

- A. nzsql
- B. nz_ddl
- C. nzsqa
- D. nz_stats
- E. nzbackup
- F. nz tables

ANSWER: B D F

Explanation:

The IBM Netezza Performance Server Software Support Tools include **nz_ddl**, **nz_stats**, and **nz tables**. These scripts are intended to gather system and database information that is useful for diagnosing appliance, workload, schema, and data-distribution issues. The **nz_ddl** script collects data-definition language information, allowing support personnel to review database object definitions and related metadata. The **nz_stats** script captures statistics information that can be relevant when investigating query-planning and performance behavior. The **nz tables** script gathers table-oriented information, helping establish details about the tables involved in a reported issue. Together, these support-collection utilities provide a structured set of diagnostic artifacts that can be supplied during problem analysis, rather than requiring administrators to manually assemble equivalent information. IBM documents Netezza administrative utilities and support-oriented tools in the Netezza Performance Server documentation; see the [IBM Netezza Performance Server documentation](#) and IBM's [Netezza Performance Server support page](#).

QUESTION NO: 2

What is a purpose of the command nzsystem?

- A. to show and change the system state
- B. to show and change nzroot use
- C. to show and change system users
- D. to show and change system access

ANSWER: A

Explanation:

The `nzsystem` command is used by IBM Netezza Performance Server administrators to display and control the appliance's overall system state. It supports operational state-management actions such as starting, stopping, pausing, resuming, and rebooting the Netezza system, as well as reporting its current state. These actions are fundamental administrative controls because they affect the availability and processing status of the entire appliance rather than an individual database, user account, or object.

For example, an administrator can use `nzsystem showState` to determine whether the system is online, paused, stopped, or in another operational condition. Commands such as `nzsystem start`, `nzsystem stop`, and `nzsystem pause` allow an authorized administrator to transition the appliance to the required state for normal operations or planned maintenance. IBM documents `nzsystem` as a system command for managing these appliance-level states and emphasizes that the applicable actions require appropriate administrative authority. See the IBM command reference for [nzsystem](#) and the IBM [Netezza Performance Server documentation](#).

QUESTION NO: 3

Which two external methods can IBM Netezza Performance Server use to authenticate database users?

- A. LDAP
- B. SSH
- C. Radius
- D. PKI-Based
- E. Kerberos

ANSWER: A E

Explanation:

IBM Netezza Performance Server can authenticate database users through LDAP and Kerberos external authentication. LDAP authentication lets Netezza Performance Server validate a user's supplied credentials against an enterprise LDAP directory. This enables centralized account administration and permits database access policies to align with an organization's established directory-service identities and password-management practices. The Netezza administrator configures the appliance's LDAP settings and associates applicable users or groups with the external authentication configuration.

Kerberos authentication supports ticket-based single sign-on. A client obtains a Kerberos service ticket from the organization's Key Distribution Center and presents that ticket when connecting to Netezza Performance Server. The server validates the ticket and maps the authenticated Kerberos principal to the relevant database identity. This approach is particularly useful in environments that use an established Kerberos realm, such as integrated enterprise identity infrastructures, because users can access the database without separately entering a database password after their Kerberos login is established.

IBM documents LDAP and Kerberos as the supported external authentication mechanisms for Netezza Performance Server. See the [IBM Netezza Performance Server security and authentication documentation](#) and the [IBM Kerberos authentication documentation](#).

QUESTION NO: 4 - (SIMULATION)

What is the order of the operations of a typical snippet in a query plan?



ANSWER: See the explanation for the answer

Explanation:

Order the operations as:

The plan scans the source data first, applies restriction/filter predicates, sorts the qualifying rows if required, and finally returns the result set.

QUESTION NO: 5

When integrating IBM Spectrum Protect into IBM Netezza Performance Server (NPS) for backup and restore, how must the NPS host be defined?

- A. as a client of the storage manager server
- B. as the primary storage manager server
- C. as the primary backup target
- D. as a cloud storage repository

ANSWER: A

Explanation:

The NPS host must be defined **as a client of the storage manager server**. In this integration, IBM Spectrum Protect (formerly IBM Tivoli Storage Manager) supplies the centralized backup-server service, while NPS is the system that initiates backup or restore activity and transfers its backup data through the installed/configured client components. Defining the NPS host as a Spectrum Protect client gives it a client-node identity and the required communication settings, authentication, storage policy association, and permissions to send backup objects to the Spectrum Protect server and retrieve them during recovery.

This client/server arrangement is essential because the Spectrum Protect server owns the storage-management catalog and controls the destination storage pools, retention, and recovery data. NPS therefore participates as a protected workload rather than taking the role of the backup-management server itself. Administrators must configure the Spectrum Protect client details on the NPS host and ensure that the corresponding client node is registered on the Spectrum Protect server before NPS backup and restore operations can use the integration. IBM's NPS documentation is available in the [IBM Netezza Performance Server documentation](#); related client/server configuration concepts are documented in the [IBM Storage Protect documentation](#).

QUESTION NO: 6

What should be the state of the IBM Netezza Performance Server system in order to run the `nz_manual_vacuum` command?

- A. online
- B. running
- C. stopped
- D. paused

ANSWER: C

Explanation:

The required state is **stopped**. The `nz_manual_vacuum` utility is a maintenance command used to manually perform vacuum processing for the Netezza Performance Server system. IBM specifies that this command is run only when the appliance is stopped, ensuring that no active database processing or data activity can conflict with the maintenance operation. Stopping the system provides the exclusive and stable system condition needed for the utility to work safely against the underlying managed data structures. Administrators should therefore plan a maintenance window, stop NPS by using the appropriate system-management procedure, run `nz_manual_vacuum` as an authorized administrative user, and then return the system to service after the operation has completed. This requirement is important operationally because the command is not intended to run concurrently with normal query workload, loads, or other database activity. IBM's command documentation and product documentation can be consulted for the applicable syntax, prerequisites, and operational procedures: [IBM Documentation search: nz_manual_vacuum](#) and [IBM Netezza Performance Server documentation](#).

QUESTION NO: 7 - (DRAG DROP)

DRAG DROP

What is the order of the operations of a typical snippet in a query plan?

Unordered Options

ScanNode
SortNode
RestrictNode
ReturnNode

Ordered Options

ScanNode
SortNode
RestrictNode
ReturnNode

ANSWER:**Unordered Options**

ScanNode
SortNode
RestrictNode
ReturnNode

Ordered Options

ScanNode
RestrictNode
SortNode
ReturnNode

Explanation:

A typical Netezza query-plan snippet follows the data as it moves through the relational operations. Processing starts at **ScanNode**, which reads the required table data on the system's processing units. This is the source of the rows used by the remainder of the snippet. The next operation is **RestrictNode**, where row-selection predicates are evaluated so that only qualifying rows continue through the plan. Applying the restriction early is important because it reduces the amount of data that later operations need to handle.

Once the relevant rows have been scanned and filtered, **SortNode** performs any ordering needed by the query plan. Sorting must operate on the filtered result set, since the sort order applies to rows that remain after the query's restrictions have been evaluated. The final operation is **ReturnNode**. It returns the completed rows produced by this snippet to the next consumer in the plan or to the final result path.

Thus, the required order is **ScanNode** → **RestrictNode** → **SortNode** → **ReturnNode**. This ordering reflects the normal flow of data from source access, through predicate evaluation and ordering, and then to result delivery. IBM's Netezza documentation describes query planning and the EXPLAIN facility used to inspect these plan operations; see the [IBM Netezza documentation](#) for the product documentation and query-plan references.

QUESTION NO: 8

Which two query-plan characteristics can help an administrator investigate inefficient join processing?

- A. Redistribution of rows before a join
- B. Large table scan activity
- C. The client workstation operating system version
- D. The password expiration date of the query user
- E. The backup retention policy

ANSWER: A B**Explanation:**

A query plan can show that rows must be redistributed before a join. Redistribution moves data between processing resources so that matching rows can be joined, and excessive movement can add significant overhead for large tables. This information can help identify opportunities to review table distribution choices.

A query plan can also show scans of large tables. When a query scans more data than expected, an administrator can investigate predicate selectivity, data organization, zone-map effectiveness, and statistics freshness. Reviewing both data movement and scan activity is a useful part of Netezza query-performance analysis. See the [IBM Netezza query tuning documentation](#).