

C1000-085^{Q&As}

IBM Netezza Performance Server V11.x Administrator

Pass IBM C1000-085 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/c1000-085.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by IBM Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

What happens to NFS mounts after a container failover?

- A. NFS is handled by Red Hat OpenShift and does not need to be mounted.
- B. Sysmgr will remount it as part of the IBM Netezza Performance Server environment
- C. A containers do not allow for NFS mounts.
- D. It will be automatically mounted by the ipsnfs service.

Correct Answer: A

QUESTION 2

Which level of transaction isolation does IBM Netezza Performance Server implement?

- A. uncommitted read
- B. serializable
- C. repeatable read D. committed read

Correct Answer: B

QUESTION 3

In IBM Netezza Performance Server, what storage file system is shared across the host containers?

- A. IBM Spectrum Scale or General Parallel File System (GPFS)
- B. NT File System (NTFS)
- C. Glusterfs
- D. 2 File System (ZFS)

Correct Answer: B

QUESTION 4

What is a valid selection from the home page of the IBM Netezza Performance Server web console?

- A. create IBM Object Storage
- B. query performance
- C. create new service

D. monitor snippets

Correct Answer: A

QUESTION 5

Which two command line support tools are used to determine the order of data for a given table?

A. nz_physical_table_layout

B. nz_zonemap

C. nz_pagemap

D. nz_sort_order

E. nz_get

Correct Answer: BC

QUESTION 6

How does IBM Netezza Performance Server store data?

A. Tables are distributed across all data slices compressed

B. Tables are distributed to a single data slice compressed.

C. Tables are distributed across all data slices uncompressed.

D. Tables are distributed across a subset of data slices uncompressed.

Correct Answer: A

QUESTION 7

When computing the row size for a table row, what overhead should be accounted for?

A. 48 bytes for ROWID, CREATEXID and DELETXID only

B. 32 bytes for ROWID, CREATEXID, DELETXID and the null vector

C. 32 bytes for ROWID, CREATEXID and DELETXID and a variable number of bytes for the null vector if there are nullable columns

D. 24 bytes for ROWID, CREATEXID and DELETXID plus null vector if there are nullable columns and record header depending on string columns

Correct Answer: A

QUESTION 8

Which two commands can be used to get a plan file for a query?

- A. nzsqli -c "show planfile "
- B. nzsqli -c "select plan_text from _v_plan"
- C. nzsqli -c "get planfile "
- D. nz_get
- E. nz_plan

Correct Answer: AD

QUESTION 9

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

Correct Answer: C

QUESTION 10

What are two primary purposes of UDX?

- A. it is a required component of all Netezza SQL commands
- B. to force execution only on the SPU nodes
- C. to offer new and unique types of analyses through customized code
- D. to showcase one's capability to write UDX code to duplicate existing SQL functions
- E. to accelerate analysis by processing data in-database

Correct Answer: DE

[C1000-085 Practice Test](#)

[C1000-085 Study Guide](#)

[C1000-085 Braindumps](#)