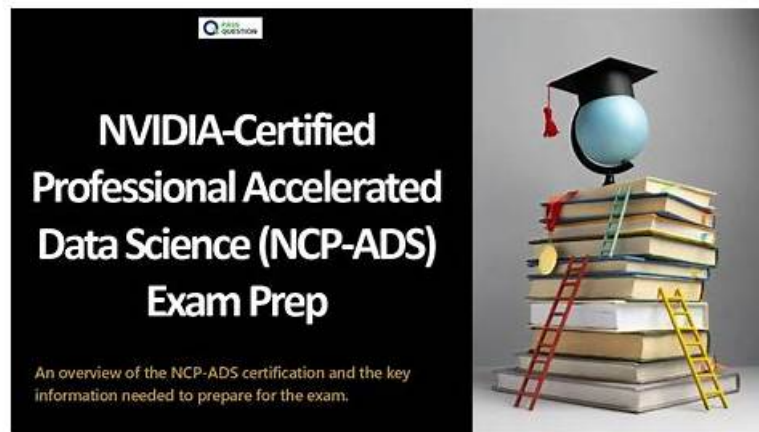


Real NCP-ADS Torrent & Real NCP-ADS Braindumps



We can proudly claim that you can successfully pass the exam just on the condition that you study with our NCP-ADS preparation materials for 20 to 30 hours. And not only you will get the most rewards but also you will get an amazing study experience by our NCP-ADS Exam Questions. For we have three different versions of our NCP-ADS study guide, and you will have different feelings if you have a try on them.

We are glad to receive all your questions on our NCP-ADS learning guide. If you have any questions about our NCP-ADS study questions, you have the right to answer us in anytime. Our online workers will solve your problem immediately after receiving your questions. Because we hope that you can enjoy the best after-sales service. We believe that our NCP-ADS Preparation exam will meet your all needs. Please give us a chance to service you; you will be satisfied with our NCP-ADS study materials.

>> Real NCP-ADS Torrent <<

Real NCP-ADS Braindumps | NCP-ADS Latest Exam Experience

Infinite striving to be the best is man's duty. We have the responsibility to realize our values in the society. Of course, you must have enough ability to assume the tasks. Then our NCP-ADS study materials can give you some help. First of all, you can easily pass the exam and win out from many candidates. The NCP-ADS certificate is hard to get. If you really crave for it, our NCP-ADS study materials are your best choice. We know it is hard for you to make decisions. You will feel sorry if you give up trying.

NVIDIA-Certified-Professional Accelerated Data Science Sample Questions (Q65-Q70):

NEW QUESTION # 65

A financial institution is developing an ETL pipeline to ingest and process large volumes of streaming data from various sources, including stock market feeds, real-time transactions, and economic indicators. The ETL process must be highly efficient to minimize latency while ensuring data integrity.

Which of the following strategies is best suited for implementing a high-performance, GPU-accelerated ETL pipeline?

- A. Use Pandas and Python's built-in threading library to handle concurrent data ingestion and transformation.
- B. Load data directly into an Excel spreadsheet and use VBA macros to clean and transform it.
- C. Store all streaming data in a PostgreSQL database before performing batch transformations.
- D. Utilize NVIDIA Morpheus with RAPIDS to preprocess real-time streaming data using GPU acceleration.

Answer: D

NEW QUESTION # 66

You need to process a dataset containing 10 billion rows, applying complex transformations and aggregations. The dataset does not fit into RAM, and you need an efficient, scalable solution for parallel processing.

Which of the following is the best choice?

- A. Vaex
- B. Pandas
- C. Dask
- D. PySpark

Answer: D

NEW QUESTION # 67

You are working with large datasets in cuDF and have noticed significant performance bottlenecks due to repeated computation and excessive shuffling in your workflow. You want to use data caching to optimize the execution plan and reduce redundant operations. Which of the following is the best way to implement data caching in cuDF to avoid repeated recomputation and excessive shuffling?

- A. Store intermediate DataFrames as temporary CSV files on disk and reload them when needed to simulate a cache.
- B. Use the `.reset_index()` method after performing transformations to force materialization of the dataset and cache results.
- C. Use the `.persist()` method on a cuDF DataFrame to store intermediate results in GPU memory, reducing redundant calculations.
- D. Use the `.to_pandas()` method to convert the cuDF DataFrame into a Pandas DataFrame, leveraging CPU caching instead of GPU caching.

Answer: C

NEW QUESTION # 68

You are a data scientist analyzing a social media network with NVIDIA cuGraph to identify the most influential users using the PageRank algorithm.

Which option best describes how cuGraph PageRank operates on a directed graph?

- A. PageRank in cuGraph uses an iterative power method to update node importance values based on incoming edges, incorporating a damping factor to handle random jumps.
- B. PageRank in cuGraph operates only on undirected graphs and cannot be applied to networks where edges have a direction.
- C. PageRank assigns equal importance to all nodes in the graph initially and updates values only based on outgoing edges, ignoring incoming edges.
- D. PageRank in cuGraph is a label propagation algorithm that clusters nodes into communities rather than ranking their importance.

Answer: A

NEW QUESTION # 69

You are building a large-scale AI training pipeline that requires efficient storage and retrieval of structured and unstructured datasets across multiple GPUs.

Which of the following is the best NVIDIA technology to organize and manage datasets at scale?

- A. NVIDIA Nsight Systems for managing dataset storage and retrieval performance.
- B. NVIDIA Magnum IO for high-performance I/O and dataset storage optimization.
- C. NVIDIA Morpheus for accelerating dataset indexing and retrieval in AI pipelines.
- D. NVIDIA Clara Imaging for storing structured and unstructured datasets efficiently.

Answer: B

NEW QUESTION # 70

.....

Dear candidates, have you thought to participate in any NVIDIA NCP-ADS exam training courses? In fact, you can take steps to pass the certification. UpdateDumps NVIDIA NCP-ADS Exam Training materials bear with a large number of the exam questions you need, which is a good choice. The training materials can help you pass the certification.

