



Get Free IBM C1000-173 Dumps PDF Questions

Why risk failure? Download updated IBM Cloud Pak for Data V4.7 Architect exam PDF questions today. Practice with real C1000-173 dumps and verified answers designed to help you ace your certification quickly using PrepBolt C1000-173 exam pdf questions and answers.

Thank you for Downloading C1000-173 exam PDF Demo

<https://prepbolt.com/C1000-173.html>

QUESTIONS & ANSWERS
DEMO VERSION
(LIMITED CONTENT)

Question 1

Question Type: MultipleChoice

Insurance industry datasets frequently include personally identifiable information (PII) and many data analysts need access to datasets but not to PII.

Which Cloud Pak for Data services leverage Data Protection Rules?

Options:

- A- Watson Discovery and Watson Assistant
- B- Data Refinery, Watson Pipeline, and Watson Studio
- C- IBM Data Virtualization, Data Privacy, and IBM Knowledge Catalog
- D- DataStage, SPSS Modeler, and Python Notebooks

Answer:

C

Explanation:

IBM Cloud Pak for Data includes built-in Data Protection Rules to enforce access control on sensitive data, such as PII. These rules are integrated directly into services like IBM Data Virtualization, Data Privacy, and IBM Knowledge Catalog. When analysts or applications access data through these services, the platform automatically masks, obfuscates, or restricts access to sensitive fields based on the defined policies. This ensures compliance with data privacy regulations and organizational security policies without manual intervention.

Question 2

Question Type: MultipleChoice

What registry permissions does OpenShift cluster node require?

Options:

- A- Only the master nodes ever require registry access.
- B- Only the bastion node should be able to reach the registry.

- C- While the bastion node pushes to the private registry, only the worker nodes pull images.
- D- All nodes must be able to push to and pull from the registry.

Answer:

D

Explanation:

In an OpenShift environment that hosts IBM Cloud Pak for Data, all cluster nodes---including master and worker nodes---must have access to the container registry to pull required images during deployment and runtime. In scenarios involving custom images, some nodes may also need to push to the registry. While the bastion node may initiate the setup or mirror images, it is not the only node involved. Therefore, all nodes should be configured with both pull and, where applicable, push access to the registry to ensure consistent deployment and operations.

Question 3

Question Type: MultipleChoice

When creating a Db2 Big SQL service instance, which two service resource items should be taken into account when sizing the cluster?

Options:

- A- Number of physical cores
- B- Amount of memory
- C- Number of virtual cores
- D- Number of workers
- E- Maximum expected throughput

Answer:

B, D

Explanation:

When provisioning a Db2 Big SQL service instance in IBM Cloud Pak for Data, key considerations include the amount of memory available per worker node and the total number of worker nodes. These factors directly affect query execution parallelism and system capacity. The system does not distinguish between physical and virtual cores at the configuration level. Throughput is an outcome of

sizing, not a direct sizing parameter. Therefore, memory and number of workers are the two most critical sizing metrics.

Question 4

Question Type: MultipleChoice

How many service instances can be provisioned for Watson Discovery at one time?

Options:

- A- 15
- B- 20
- C- 5
- D- 10

Answer:

D

Explanation:

'You can create a maximum of 10 instances per deployment. After you reach the maximum number, the New instance button is not displayed in IBM Cloud Pak for Data.'

Question 5

Question Type: MultipleChoice

Which type of search allows Watson Discovery to find the most relevant material in documents?

Options:

- A- Keyword
- B- Faceted
- C- Total

Answer:

A

Explanation:

Watson Discovery employs advanced natural language processing and ranking algorithms that begin with keyword-based search. This type of search enables the tool to locate and return the most relevant document passages based on the presence and context of keywords in user queries. While it may also use metadata filters (facets), the core retrieval method remains keyword-driven to find relevance at both document and passage levels.

Question 6

Question Type: MultipleChoice

What endpoint will an application use to interact with Db2 Big SQL?

Options:

- A- Representative State Transfer (REST) Endpoint
- B- System Local Efficient Endpoint Pathways (SLEEP)
- C- Simple Normalized Optimum Representative Endpoint (SNORE)
- D- Dynamic Representative Endpoint Activation Mobility (DREAM)

Answer:

A

Explanation:

Applications interact with Db2 Big SQL using industry-standard protocols. The most common and supported interface is through a REST (Representative State Transfer) API endpoint. REST endpoints allow for external applications to query, manage, and manipulate data within Big SQL using simple HTTP calls. None of the other options---SLEEP, SNORE, or DREAM---are valid or recognized interfaces in IBM Cloud Pak for Data or Db2 Big SQL documentation.

Question 7

Question Type: MultipleChoice

Which statement describes MPP (Massively Parallel Processing) Database architecture?

Options:

- A- A data warehouse that needs all compute nodes to access a shared data store simultaneously.
- B- Two or more databases kept in sync via two-phase commit.
- C- An analytics environment that improves performance by dividing the data across many nodes.
- D- A transactional system which uses multiple nodes for maximum availability.

Answer:

C

Explanation:

MPP, or Massively Parallel Processing, is a database architecture model where data is divided and processed across multiple compute nodes in parallel. Each node works independently on a portion of the data, dramatically improving query performance and throughput for analytics workloads. This model is ideal for big data and analytical queries, not transactional workloads. It differs from shared-disk models or replication strategies like two-phase commit. The correct definition involves distributed data and parallel query execution, as described in option C.

Question 8

Question Type: MultipleChoice

Which Watson Pipeline component puts a value in columns so it can be consumed by DataStage?

Options:

- A- Instantiate User Columns
- B- Prepare User Parameters
- C- Set User Variables
- D- Initialize User Values

Answer:

C

Explanation:

In Watson Pipelines, the component that enables users to define and assign values that can be referenced later in the pipeline---including by downstream components like DataStage---is Set User Variables. This component allows the user to create name-value pairs and store them as environment variables, which are accessible to DataStage and other execution blocks. This ensures dynamic parameter passing and enhances pipeline reusability. The other options listed do not correspond to valid Watson Pipeline components as defined in the official Cloud Pak for Data 4.7 release.

Thank You for trying C1000-173 PDF Demo

To try our C1000-173 practice exam software
visit link below

<https://prepbolt.com/C1000-173.html>

Start Your C1000-173 Preparation

Use Coupon **"SAVE50"** for extra 50% discount on the purchase of Practice Test Software. Test your C1000-173 preparation with actual exam questions.