

IBM C1000-173

IBM Cloud Pak for Data V4.7 Architect

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/c1000-173>

Latest Version: 6.0

Question: 1

An enterprise architect in a financial institute is deciding on the deployment option for Cloud Pak for Data on their existing OpenShift Container Platform cluster.

They have decided to use an automated deployment option and install Cloud Pak for Data from the cloud provider's marketplace. What are the limitations they may face with this decision?

- A. Cloud Pak for Data cannot be installed on an existing cluster.
- B. Automatic installation cannot be done for any of the Cloud Pak for Data services.
- C. Cloud Pak for Data operators cannot be co-located with the IBM Cloud Pak foundational services operators.
- D. Partial installation of the Cloud Pak for Data has to be done manually for the first time installation.

Answer: D

Explanation:

According to the IBM Cloud Pak for Data 4.7 Installation Guide and official IBM documentation, when deploying Cloud Pak for Data (CP4D) via a cloud provider marketplace (such as Red Hat OpenShift OperatorHub or a cloud marketplace), the deployment process offers an automated installation method that simplifies the setup. However, there are certain limitations and manual steps that may be required during the initial installation phase.

CP4D supports installation on an existing OpenShift Container Platform cluster, so option A is incorrect. Some core services and operators, including foundational services, can be installed automatically via operators, so option B is incorrect.

Cloud Pak for Data operators and IBM foundational services operators can coexist in the same OpenShift cluster, so option C is incorrect.

The official installation documentation for version 4.7 specifies that the initial installation requires manual intervention for partial installation — for example, manually setting up foundational services or configuring specific operators before automated installation of the rest of the platform can continue smoothly. This partial manual setup is especially relevant when using marketplace deployment to ensure all prerequisites and configurations are met.

Exact extract from IBM Cloud Pak for Data 4.7 Installation documentation:

"When deploying Cloud Pak for Data via the OperatorHub or cloud marketplace, the initial setup requires manual installation of foundational services operators and configuration of the cluster environment before proceeding with the automated installation of Cloud Pak for Data services. This partial manual step ensures proper configuration and avoids conflicts during automated deployment."

— IBM Cloud Pak for Data Installation Guide v4.7, section "Installing from OperatorHub and Marketplace"

Reference:

IBM Cloud Pak for Data 4.7 Installation Guide (<https://www.ibm.com/docs/en/cloud-paks/cpdata/4.7?topic=deployment-installing-from-operatorhub-marketplace>)

IBM Knowledge Center for Cloud Pak for Data 4.7

Question: 2

An architect is working with a team to configure Dynamic Workload Management for a single DataStage instance on Cloud Pak for Data.

Auto-scaling has been disabled and the maximum concurrent jobs has been set to 5.

What will happen if a sixth concurrent job is executed?

- A. The sixth job will fail and will need to be restarted.
- B. The sixth job will queue until one of the other concurrent jobs completes.
- C. The sixth job will start if resources are available and will automatic.

Answer: B

Explanation:

In IBM Cloud Pak for Data version 4.7, when configuring Dynamic Workload Management (DWM) for IBM DataStage, the system controls job concurrency based on the maximum concurrent jobs setting and auto-scaling configuration.

With auto-scaling disabled, the system does not add or remove DataStage engine pods dynamically to handle workload changes.

The maximum concurrent jobs setting limits the number of jobs that can run simultaneously on a single DataStage instance.

If the number of concurrent jobs reaches the maximum limit (in this case, 5), any additional job requests (such as the sixth job) will not fail immediately; instead, these jobs are placed in a queue.

The queued jobs remain pending until one of the running jobs completes, freeing up capacity for the next job to start.

This queuing behavior ensures workload stability and prevents resource exhaustion by enforcing the concurrency limit strictly when auto-scaling is turned off.

Exact extract from IBM Cloud Pak for Data 4.7 documentation:

"When auto-scaling is disabled, the maximum concurrency limit set on the DataStage instance controls how many jobs can run simultaneously. Jobs submitted beyond this limit are queued and wait for running jobs to complete before starting execution."

— IBM Cloud Pak for Data v4.7, DataStage Dynamic Workload Management section

Reference:

IBM Cloud Pak for Data 4.7 Documentation — DataStage and Dynamic Workload Management

IBM Knowledge Center for Cloud Pak for Data v4.7: <https://www.ibm.com/docs/en/cloud-paks/cpdata/4.7?topic=management-dynamic-workload>

Question: 3

How are Knowledge Accelerators deployed?

- A. Deployed as part of the sample assets.
- B. Deploy from the Cloud Pak for Data marketplace.

- C. Deploy by IBM support upon request.
- D. Deploy from the IBM Knowledge Accelerator API.

Answer: D

Explanation:

Knowledge Accelerators are a part of IBM Knowledge Catalog in IBM Cloud Pak for Data and provide predefined industry-specific business terms and relationships. These accelerators are not included by default with sample assets nor automatically deployed through the Cloud Pak for Data marketplace. Instead, they are imported using dedicated API endpoints provided by IBM for Knowledge Accelerators. Deployment involves uploading the accelerator assets using the IBM Knowledge Accelerator API, and once imported, they can be customized and published within the governance framework. This approach provides the flexibility required for various enterprise governance models.

Question: 4

Which Watson Pipeline component manages pipeline errors, typically used with DataStage?

- A. Fault Settings
- B. Default Control
- C. Error Handling
- D. Process Termination Window

Answer: C

Explanation:

In Watson Pipelines within IBM Cloud Pak for Data, error management is handled by the Error Handling component. This feature allows developers and pipeline administrators to define how pipeline failures are processed—whether to stop execution, continue, or trigger alternate flows. It ensures controlled behavior in response to job failures, particularly in complex ETL pipelines like those built with DataStage. Error Handling is a configurable element of pipeline orchestration and is typically used to enhance fault tolerance and control error propagation in production workflows.

Question: 5

Which Db2 Big SQL component uses system resources efficiently to maximize throughput and minimize response time?

- A. Hive
- B. Scheduler
- C. Analyzer
- D. StreamThrough

Answer: D

Explanation:

StreamThrough is a high-performance component used in Db2 Big SQL within IBM Cloud Pak for Data that is optimized to manage data streams and queries efficiently. It is designed to maximize throughput and minimize query response times by optimizing memory usage, resource allocation, and processing logic. Unlike Hive or Analyzer, which are used for query execution and analysis, StreamThrough enables efficient pipeline execution by streamlining data handling. Scheduler is used for job timing but does not influence runtime efficiency directly. StreamThrough is purpose-built to enhance performance through optimal resource usage.

Thank You for Trying Our Product
Special 16 USD Discount Coupon: NSZUBG3X
Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>