PassQueen
a.k.a dumpsqueen.com

IBM Cloud Pak for Business Automation v21.0.3 Administration

IBM C1000-150

Version Demo

Total Demo Questions: 8

Total Premium Questions: 82

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Planning	18
Topic 2, Installation	8
Topic 3, Configuration	22
Topic 4, Administration	18
Topic 5, Troubleshooting	16
Total	82

QUESTION NO: 1

What are two functions of container resource requests in an OpenShift pod specification?

A. They help the scheduler select a node with sufficient resources.

B. They contribute to namespace resource-quota calculations.

C. They automatically restart failed containers.

D. They define the maximum resource use of a container.

E. They create persistent volumes for the pod.

ANSWER: A B

Explanation:

They help the scheduler select a node with sufficient resources and **they contribute to namespace resource-quota calculations** are correct. CPU and memory requests represent the resources a container needs to run. The scheduler uses these values when determining whether a node can host the pod. Requests are also evaluated by ResourceQuota controls that limit aggregate resource consumption in a namespace.

Resource requests do not automatically restart a failed container and do not define the maximum permitted resource use. Resource limits define the maximum CPU or memory that a container can consume. See the Kubernetes documentation for Resource Management for Pods and Containers.

QUESTION NO: 2

Which two steps are required to install Operational Decision Manager Standalone CNCF for production?

A. Prune existing deployment snapshots from the database.

B. Download the container images.

C. Delete the container registry image pull secret.

D. Add the resAdmin user to the IBM Entitled registry.

E. Download the archive from IBM Passport Advantage (PPA) to get the Helm chart.

ANSWER: B E

Explanation:

Downloading the container images and downloading the archive from IBM Passport Advantage (PPA) to get the Helm chart are required preparation steps for a production installation of Operational Decision Manager Standalone on a CNCF-certified Kubernetes environment. The PPA archive supplies the product deployment artifacts, including the Helm chart used to define and install the ODM Kubernetes resources. Helm uses this chart, together with environment-specific values, to create and configure the ODM deployment in the target cluster.

The ODM container images are also required because the Kubernetes workloads for the Decision Server components must be instantiated from the supported product images. In a production process, the downloaded images are made available to the registry from which the cluster can pull them, allowing the Helm deployment to start the required pods with the correct ODM software version. These two downloads therefore provide both essential parts of the installation: the deployment definition and the executable runtime images. IBM's installation documentation describes obtaining the Helm chart from Passport Advantage and obtaining or mirroring the required container images before deploying the standalone offering. See the IBM Cloud Pak for Business Automation ODM standalone production installation documentation and the IBM Operational Decision Manager documentation.

QUESTION NO: 3

In preparing the IBM Business Automation Studio and Application Engine environments for recovery, the shared_configuration. sc_run_as_user in the CR definition must have what in common between the primary and secondary environments?

A. Identical LDAP configurations

B. Identical router host name

C. Same OpenShift administrator account

D. Same UID

ANSWER: D

Explanation:

Same UID is required. The `shared_configuration.sc_run_as_user` setting identifies the operating-system user under which the relevant IBM Business Automation Studio and Application Engine workloads run. For a disaster-recovery topology, that user must resolve to the same numeric user identifier (UID) in both the primary and secondary environments. This consistency ensures that files, mounted storage content, and process access expectations retain compatible ownership and permissions after workloads are recovered or started in the secondary cluster.

UIDs, rather than merely user names, are significant to Linux and containerized storage permissions. A matching name with a different numeric UID can result in a recovered workload being unable to read, write, or otherwise access data created in the primary environment. Configuring the same UID for `sc_run_as_user` before recovery preparation therefore helps ensure that the recovered environment can use shared or restored persistent data safely and predictably. IBM Cloud Pak for Business Automation documentation provides deployment and operational guidance for the product at IBM Cloud Pak for Business Automation 21.0.3 documentation. For background on why numeric UIDs govern ownership and access on Linux systems, see the Red Hat Linux file permissions overview.

QUESTION NO: 4

What is the primary purpose of a NetworkPolicy in an OpenShift cluster?

A. To control allowed network traffic to and from pods.

B. To create external routes for services.

C. To assign persistent volumes to pods.

D. To provide credentials for pulling container images.

ANSWER: A

Explanation:

To control allowed network traffic to and from pods is correct. A NetworkPolicy defines permitted ingress traffic, egress traffic, or both for pods selected by labels. It can restrict communication between Cloud Pak for Business Automation workloads and other services to only the required ports, namespaces, and peer pods.

NetworkPolicy does not create routes, assign storage, or configure image registry access. Its enforcement depends on the cluster network plugin supporting network policies. See the Kubernetes documentation for Network Policies.

QUESTION NO: 5

To allow a Cloud Pak for Business Automation deployment to pull images from a secured private registry, which two actions are required?

A. Create an image pull secret in the target namespace.

B. Reference the image pull secret in the custom resource configuration.

C. Store the registry password in a ConfigMap.

D. Grant cluster-admin permission to every service account.

E. Create a Route for the registry.

ANSWER: A B

Explanation:

Create an image pull secret in the target namespace and **reference the image pull secret in the custom resource configuration** are correct. The pull secret contains the registry authentication data that OpenShift uses when it retrieves protected images. The Cloud Pak for Business Automation custom resource identifies the pull secret so that the operator can apply the image access configuration to the managed workloads. A ConfigMap does not protect registry credentials, and granting cluster-admin authority does not provide authentication to a container registry. IBM documents image-access configuration in the Cloud Pak for Business Automation 21.0.3 documentation. See also the Red Hat guidance for managing image pull secrets.

QUESTION NO: 6

Once a starter deployment of the Cloud Pak for Business Automation is installed, where can access to the different capability services and applications be found?

A. By opening a terminal to the ibm-cp4a-operator pod and open the /opt/ibm/cp4ba-access.txt file.

B. By opening a config map which contains the route URL to access the components and a secret which contains the credentials to use with the different URLs.

C. By opening a config map which contains the route URL to access the components as well as the username and password to use with the URL in clear text.

D. By opening the cpd-access route, which leads to a page that lists the components URLs, usernames and passwords to use.

ANSWER: B

Explanation:

By opening a config map which contains the route URL to access the components and a secret which contains the credentials to use with the different URLs. is correct because a Cloud Pak for Business Automation starter deployment exposes its installed capabilities through OpenShift routes. The deployment makes the relevant endpoint information available in a Kubernetes ConfigMap, allowing an administrator to identify the URLs for the capability services and applications. Authentication information is stored separately in a Kubernetes Secret, which is the appropriate platform resource for credentials and other sensitive values. An administrator can use these two resources together to obtain the appropriate service URL and the credentials required to sign in.

This arrangement follows the Kubernetes model of separating non-sensitive configuration data from sensitive data. ConfigMaps are intended for configuration values that can be consumed by workloads or inspected by authorized administrators, whereas Secrets provide the mechanism for storing confidential information such as user credentials. IBM's Cloud Pak for Business Automation documentation describes accessing deployed capabilities and their routes, while Red Hat documents the platform distinction between ConfigMaps and Secrets. See the IBM Cloud Pak for Business Automation 21.0.3 documentation and the OpenShift documentation for ConfigMaps and Secrets.

QUESTION NO: 7

After scaling a Process Mining deployment, which two replica set values are updated?

A. READY

B. PROGRESS

C. WAITING

D. COMPLETED

E. DESIRED

ANSWER: A E

Explanation:

DESIRED and **READY** are the relevant ReplicaSet values to observe after a Process Mining deployment is scaled. Scaling changes the target number of pod replicas that the workload controller must maintain, which is represented by the DESIRED count. As the controller creates or terminates Process Mining pods to reach that target, the READY count reflects how many of those replicas have successfully started and passed their readiness checks. Administrators use these values together to confirm both the requested scale level and whether the scaled workload is actually available to receive traffic. When the scale operation and pod startup complete successfully, READY should converge on DESIRED. This operational model applies to the Kubernetes/OpenShift workload resources used by Cloud Pak for Business Automation deployments, including Process Mining. IBM's administration documentation is available in the Cloud Pak for Business Automation 21.0.3 documentation. The standard ReplicaSet status fields and their meanings are documented by Kubernetes ReplicaSet documentation.

QUESTION NO: 8

How is the Business Automation Studio web interface accessed?

A. Via web browser at URL https : //: /BAStudio/

B. Via IBM Cloud Pak process administration console

C. Via Workflow Center Web Console

D. Via IBM Cloud Pak platform UI

ANSWER: A

Explanation:

Business Automation Studio is a browser-based authoring environment, so it is accessed directly through its deployed web URL: `https://<host>:<port>/BAStudio/`. The host and port are the values configured for the Business Automation Studio deployment and identify the server endpoint on which the Studio web application is available. After navigating to this address in a supported web browser, the user authenticates with the configured credentials and can use the Studio interface to create and manage automation artifacts such as processes, decisions, and user interfaces, subject to assigned permissions. This direct URL is the normal entry point because Business Automation Studio is itself a web application, rather than a page launched from an administrative console. IBM documentation describes Business Automation Studio as a web-based environment and provides its access and getting-started information through the Business Automation Studio documentation. See the IBM Business Automation Studio overview and the IBM Cloud Pak for Business Automation 21.0.3 documentation.