

Practice Exam Questions



Professional
Certification



C1000-150

IBM Cloud Pak for Business Automation Administration



EXAMKILLER

Help Pass Your Exam At First Try



Exam C1000-150

IBM Cloud Pak for Business Automation v21.0.3 Administration

Version: 6.0

[Total Questions: 60]

Question No : 1

Which type of event reflects an ongoing activity instance in a workflow process?

- A. BAW event
- B. Time series
- C. Tracking event
- D. Active summary

Answer: C

Explanation: A tracking event reflects an ongoing activity instance in a workflow process. It is used to capture and track the progress of the process and can be used to create an audit trail.

References: [1] <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-monitor-and-manage-logs> [2] <http://www.bawtech.com/what-are-business-activity-workflows-baw/>

Question No : 2

Prior to deploying the Cloud Pak for Business Automation operator, which two common prerequisites exist for all Cloud Pak for Business Automation capabilities (excluding Business Automation Insights)?

- A. LDAP
- B. Database
- C. Persistent Volumes
- D. Network Policies
- E. Routes

Answer: B,C

Explanation: Before deploying the Cloud Pak for Business Automation operator, it is important to ensure that the necessary database and persistent volume are in place. This will allow the operator to store and persist data as needed. The other options are not prerequisites for all Cloud Pak for Business Automation capabilities.

References:

[1] https://www.ibm.com/support/knowledgecenter/SSFTN5_9.7.1/com.ibm.wbpm.inst.doc/topics/t_k8s_prereq.html [2] https://www.ibm.com/support/knowledgecenter/SSFTN5_9.7.1/com.ibm.wbpm.inst.doc/topics/t_k8s_install_operator.html

Question No : 3

Where do the images reside for an air-gapped Cloud Pak for Business Automation upgrade?

- A. IBM registry
- B. RedHat quay.io registry
- C. Local registry
- D. Docker Hub

Answer: C

Explanation: When performing an air-gapped upgrade of Cloud Pak for Business Automation, the images used for the upgrade reside in a local registry. An air-gapped environment is one in which there is no external network access, so the images cannot be pulled from a remote registry such as IBM registry, RedHat quay.io registry, or Docker Hub. Instead, the images must be pre-pulled and stored in a local registry that is accessible to the OpenShift cluster.

Reference:

https://www.ibm.com/support/knowledgecenter/en/SSYHZ8_20.0.x/com.ibm.dba.baw.install/topics/install_airgap_prereq.html

https://www.ibm.com/support/knowledgecenter/en/SSYHZ8_20.0.x/com.ibm.dba.baw.install/topics/install_airgap_procedure.html

https://www.ibm.com/support/knowledgecenter/en/SSYHZ8_20.0.x/com.ibm.dba.baw.install/topics/install_airgap_registry.html

Question No : 4

What is the best approach to determine the cause of a performance issues with Operational Decision Manager Rule Execution Server?

- A. Enable an audit trace on the Decision Server Console.
- B. Collect browser network trace.
- C. Specify duration threshold and set log level to FINE.
- D. Capture a JVM heap dump.

Answer: A

Explanation: The best approach to determine the cause of a performance issues with Operational Decision Manager Rule Execution Server is to enable an audit trace on the Decision Server Console. This will allow you to view detailed information about rule

execution, including the time taken for each rule, the parameters passed to each rule, and the results returned.

References:

[1] https://www.ibm.com/support/knowledgecenter/SSFTN5_9.7.1/com.ibm.wbpm.inst.doc/topics/t_audit_trace.html [2] https://www.ibm.com/support/knowledgecenter/SSFTN5_9.7.1/com.ibm.wbpm.inst.doc/topics/t_ds_console_overview.html

Question No : 5

Which type of log collector uses input and output plug-ins to collect data from multiple sources and to distribute or send data to multiple destinations?

- A. Journald
- B. Rsyslog Sidecar
- C. Fluentd
- D. Audit Container

Answer: C

Explanation: Fluentd is a log collector that uses input and output plug-ins to collect data from multiple sources and to distribute or send data to multiple destinations. This allows Fluentd to collect and process data from various sources and send it to various destinations with minimal effort.

References: [1] <https://docs.fluentd.org/> [2] <https://www.fluentd.org/>

Question No : 6

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

- A. READY
- B. PROGRESS
- C. WAITING
- D. COMPLETED
- E. DESIRED

Answer: D,E

Explanation: After scaling a Process Mining deployment, the replica set values for COMPLETED and DESIRED are updated.

References: [1] <https://docs.camunda.org/manual/7.11/process-mining/technical-setup/> [2] <https://docs.camunda.org/manual/7.11/process-mining/configuration/>

Question No : 7

What should be supplied as part of the custom resource prior to deployment if it is desired to use a root CA signer certificate that is signed by a recognized certificate authority?

- A. root_ca_certificate
- B. root_ca_key
- C. root_ca_store
- D. root_ca_secret

Answer: A

Explanation: If it is desired to use a root CA signer certificate that is signed by a recognized certificate authority, the rootcacertificate should be supplied as part of the custom resource prior to deployment. This is necessary in order for the root CA signer certificate to be validated.

References: [1] <https://kubernetes.io/docs/tasks/configure-pod-container/configure-pod-certificates/#running-an-https-server> [2] <https://kubernetes.io/docs/concepts/cluster-administration/certificates/>

Question No : 8

What does IBM Cloud Pak foundational services monitoring require?

- A. Role-based access control (RBAC) to monitor APIs and data
- B. Red Hat OpenShift Container Platform monitoring to be installed
- C. Kibana as the datasource
- D. Adopter customization to query and visualize application metrics

Answer: A

Explanation: IBM Cloud Pak foundational services monitoring requires Role-based access control (RBAC) to monitor APIs and data. This ensures that only authorized users have access to the data and APIs that are being monitored. It also ensures that data is only being accessed by users with the appropriate permissions. Kibana is used as the data source for the Cloud Pak foundational services monitoring. Adopter customization is only

necessary to query and visualize application metrics. Red Hat OpenShift Container Platform monitoring is not required for Cloud Pak foundational services monitoring.

References:

[1] https://www.ibm.com/support/knowledgecenter/SSFTN5_2.2.2/monitoring/overview.html

[2] https://www.ibm.com/support/knowledgecenter/SSFTN5_2.2.2/monitoring/rbac.html

Question No : 9

Which statement is true concerning secrets on Kubernetes master (API server)?

- A. Secrets are by default encrypted at rest.
- B. Extra steps are needed to encrypt secrets at rest.
- C. Secrets only transit through Kubernetes master; encryption at rest is out of its responsibilities.
- D. Secrets are self-secured needing no encryption at rest.

Answer: B

Explanation: Secrets on Kubernetes master (API server) are not encrypted by default at rest, extra steps are needed to encrypt them. Kubernetes does not automatically encrypt secrets at rest, it is the responsibility of the user to ensure that the secrets are encrypted. There are several ways to encrypt secrets at rest, such as using Kubernetes native encryption, third-party encryption tools or using a cloud provider's encryption services.

Question No : 10

What is one consideration for data backup within a multiple-zone cluster?

- A. Make regular backups of each environment.
- B. Enable the multiple-zone cluster sync feature.
- C. Leverage the OpenShift 'replica sync' feature.
- D. Make regular backups for the active zone only.

Answer: D

Explanation: When backing up data within a multiple-zone cluster, it is important to make regular backups for the active zone only. This ensures that the data is up-to-date and any changes made in the other zones are not lost if something goes wrong.

Question No : 11

When using the product Certificate manager service, what might one consider in regards to the default setting of refreshing certificates?

- A. Pods that use the refreshed certificate are automatically restarted.
- B. Certificates cannot be configured to automatically refresh.
- C. Certificate Refresh is disabled by default.
- D. Pods that use the refreshed certificate need to be manually restarted.

Answer: C

Explanation: When using the product Certificate manager service, the default setting of refreshing certificates is disabled. Pods that use the refreshed certificate do not need to be manually restarted, as certificate refresh will be automatically enabled when the certificate expires.

Question No : 12

Which parameter is required to forward audit logging?

- A. AUDIT_CONTENT_BY_PROVIDERS
- B. ENABLE_AUDIT_LOGGING_FORWARDING
- C. AUDIT_ENABLED
- D. SAS API SERVER AUDIT ENABLED

Answer: B

Explanation: To forward audit logging, the ENABLEAUDITLOGGINGFORWARDING parameter is required. This parameter is used to enable the forwarding of audit logs to an external service.

References:

- [1] https://www.ibm.com/support/knowledgecenter/SSFTN5_2.2.2/cpd/administer/audit.html
- [2] https://www.ibm.com/support/knowledgecenter/SSFTN5_2.2.2/cpd/administer/overview.html

Question No : 13

A starter deployment requires which two capabilities to be installed independently?

- A. Operational Decision Manager and Automation Decision Services

- B. Content Platform Engine and Navigator
- C. Process Mining and Robotic Process Automation
- D. Business Automation Insights and Kafka

Answer: A

Explanation: When deploying the Cloud Pak for Business Automation, the starter deployment requires one capability to be installed independently:

Operational Decision Manager (ODM) and Automation Decision Services (ADS): These capabilities provide a set of tools for creating and managing business rules, decision services, and analytics. They are typically used to automate decision-making processes within an organization.

https://www.ibm.com/support/knowledgecenter/en/SSYHZ8_20.0.x/com.ibm.dba.baw.install/topics/install_overview.html

Question No : 14

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: To install Operational Decision Manager Standalone CNCF for production, two steps are required:

- ☞ Download the container images: The container images are necessary for the installation, and they need to be downloaded and made available to the cluster.
- ☞ Download the archive from IBM Passport Advantage (PPA) to get the Helm chart: The Helm chart contains the configuration information necessary to install and deploy the ODM. It can be downloaded from the IBM Passport Advantage website.

It's important to note that other steps such as pruning existing deployment snapshots from the database, deleting the container registry image pull secret, or adding the resAdmin user to the IBM Entitled registry are optional and not required for the installation of Operational Decision Manager Standalone CNCF for production.

Reference:

https://www.ibm.com/support/knowledgecenter/en/SSYHZ8_20.0.x/com.ibm.dba.odm.cncf/

cncf_install_overview.html

https://www.ibm.com/support/knowledgecenter/en/SSYHZ8_20.0.x/com.ibm.dba.odm.cncf/cncf_install_images.html

Question No : 15

What are Kubernetes events?

- A. Objects that create pods and run them to completion.
- B. Objects that determine which jobs run and when.
- C. Objects that provide insight about what is happening inside a cluster.
- D. Objects that determine which tasks are run during deployment.

Answer: C

Explanation: Kubernetes Events are objects that provide insight about what is happening inside a cluster, such as the state of a pod, the status of a deployment, or any other type of activity. They can be used to detect and troubleshoot errors, or to track the progress of a deployment.

References: [1] <https://kubernetes.io/docs/tasks/debug-application-cluster/discover-events/> [2] <https://kubernetes.io/docs/concepts/cluster-administration/logging/>

Question No : 16

Which two foundation pattern components are always deployed with any of the Cloud Pak for Business Automation capabilities?

- A. Kafka and Business Automation Studio
- B. Application Engine (AE) and Platform UI
- C. Business Automation Navigator (BAN) and Resource Registry (RR)
- D. Business Automation Insights (BAI) and Content Platform Engine (CPE)

Answer: C

Explanation: The two foundation pattern components that are always deployed with any of the Cloud Pak for Business Automation capabilities are Business Automation Navigator (BAN) and Resource Registry (RR). BAN provides an entry point for users to access the Business Automation capabilities, while RR is the repository for the resources used in the deployment process. The other options are not related to this process.