PassQueen
a.k.a dumpsqueen.com

IBM Cloud Pak for Integration V2021.2 Administration

IBM C1000-130

Version Demo

Total Demo Questions: 15

Total Premium Questions: 150

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Planning	41
Topic 2, Installation	30
Topic 3, Configuration	39
Topic 4, Administration	30
Topic 5, Troubleshooting	10
Total	150

QUESTION NO: 1

Which two statements are true about using node labels and node selectors for Cloud Pak for Integration workloads?

A. Node labels can identify nodes intended for particular workloads.

B. A node selector restricts scheduling to nodes with matching labels.

C. A node selector automatically adds missing labels to worker nodes.

D. Node labels replace the need for OpenShift RBAC.

E. A node selector guarantees that pods are distributed across availability zones.

ANSWER: A B

Explanation:

Node labels can identify nodes intended for particular workloads and **a node selector restricts scheduling to nodes with matching labels** are correct. Administrators can label worker nodes to identify infrastructure characteristics such as storage availability, worker role, or a workload-isolation boundary. A workload that includes a matching node selector is eligible for scheduling only on nodes that have the specified labels.

A node selector is a scheduling constraint, not a mechanism for creating nodes or granting users access to nodes. If no eligible labeled node is available, the pod remains pending until the constraint can be satisfied. See the OpenShift documentation on controlling Pod placement onto nodes.

QUESTION NO: 2

An administrator is checking that all components and software in their estate are licensed. They have only purchased Cloud Pak for Integration (CP41) li-censes.

How are the OpenShift master nodes licensed?

A. CP41 licenses include entitlement for the entire OpenShift cluster that they run on, and the administrator can count against the master nodes.

B. OpenShift master nodes do not consume OpenShift license entitlement, so no license is needed.

C. The administrator will need to purchase additional OpenShift licenses to cover the master nodes.

D. CP41 licenses include entitlement for 3 cores of OpenShift per core of CP41.

ANSWER: A

Explanation:

CP41 licenses include entitlement for the entire OpenShift cluster that they run on, and the administrator can count against the master nodes. Cloud Pak for Integration includes a restricted-use entitlement for Red Hat OpenShift Container Platform to support the Cloud Pak deployment. This entitlement is not limited solely to application worker nodes: it applies to the OpenShift cluster used to run the Cloud Pak, including the control-plane (master) nodes. Consequently, when determining the OpenShift capacity covered by the Cloud Pak entitlement, the administrator may apply the purchased CP41 license capacity to the master-node cores as part of the cluster's overall licensing calculation. The entitlement is restricted to use in support of the Cloud Pak environment, so the cluster and its capacity must be used consistently with the Cloud Pak license terms. This approach lets an organization that has purchased the appropriate CP41 capacity license both Cloud Pak for Integration and the OpenShift infrastructure on which it runs, without treating the master nodes as automatically excluded from coverage. IBM's Cloud Pak for Integration planning documentation describes the included OpenShift entitlement and its applicable deployment scope; consult the product documentation and the applicable license terms when sizing a production cluster. See IBM Cloud Pak for Integration licensing documentation and IBM Passport Advantage license information.

QUESTION NO: 3

What is a prerequisite when configuring foundational services IAM for single-sign-on?

A. Access to the OpenShift Container Platform console as kubeadmin.

B. Access to IBM Cloud Pak for Integration as kubeadmin.

C. Access to OpenShift cluster as root.

D. Access to IAM service as administrator.

ANSWER: A

Explanation:

Access to the OpenShift Container Platform console as kubeadmin is the required prerequisite because the initial configuration of foundational services Identity and Access Management single sign-on requires cluster-level administrative access. The `kubeadmin` account is the bootstrap OpenShift administrator account and can log in to the OpenShift web console with the permissions needed to access the foundational-services management interface and establish the initial IAM configuration. This initial privileged access is important because SSO configuration affects how users authenticate to Cloud Pak capabilities and how identity-provider integration is administered. After IAM and an identity provider are configured, organizations can grant appropriate administrative and user roles through the configured authentication and authorization model rather than relying on the bootstrap account for ordinary access. IBM's foundational services documentation describes Identity and Access Management as the component used to manage authentication and SSO-related access for Cloud Pak deployments. OpenShift documentation also identifies `kubeadmin` as the initial cluster-administrator user created for a new cluster. See the IBM Cloud Pak foundational services documentation and Red Hat OpenShift authentication documentation.

QUESTION NO: 4 - (SIMULATION)

Select all that apply

What is the correct order of the Operations Dashboard upgrade?

ANSWER: See the explanation for the answer

Explanation:

The correct Operations Dashboard upgrade order is:

The operator must be upgraded first because it manages the Operations Dashboard operand. Where manual approval is configured, the OLM Install Plan must be approved before the operator upgrade can complete. Upgrade the Operations Dashboard instance (operand) next, then upgrade the integration capabilities that are traced by Operations Dashboard.

QUESTION NO: 5 - (HOTSPOT)

Before upgrading the Foundational Services installer version, the installer catalog source image must have the correct tag. To always use the latest catalog click on where the text 'latest' should be inserted into the image below?

Answer Area

ANSWER:

Answer Area

Explanation:

The required change is made to the tag portion of the installer catalog image reference. Container image references use the form `registry-or-path/image-name:tag`. In the displayed CatalogSource specification, the image name is followed by a colon and the current version value, shown as `3.8.0`. That final value is the tag. Replacing that value with `latest` makes the CatalogSource use the latest published installer catalog image, yielding an image reference that ends with `ibm-common-service-catalog:latest`.

This matters during a Foundational Services installer upgrade because the CatalogSource is the OpenShift resource that tells Operator Lifecycle Manager where to retrieve the operator catalog. Pointing its image at the current catalog tag ensures that the catalog metadata made available to the cluster corresponds to the intended installer catalog release. The other fields in the snippet define CatalogSource behavior and polling configuration, but the actual catalog image version is controlled specifically by the suffix after the colon on the `image:` line. IBM's upgrade guidance explicitly shows the required catalog image using the `latest` tag. See the IBM documentation for upgrading Foundational Services from an operator release.

QUESTION NO: 6

Given the high availability requirements for a Cloud Pak for Integration deployment, which two components require a quorum for high availability?

A. Multi-instance Queue Manager

B. API Management (API Connect)

C. Application Integration (App Connect)

D. Event Gateway Service

E. Automation Assets

ANSWER: A B

Explanation:

Multi-instance Queue Manager and **API Management (API Connect)** are the components that rely on quorum-based decisions in a highly available Cloud Pak for Integration design. For IBM MQ high availability, the queue manager's HA arrangement uses multiple members so that a surviving majority can safely determine which member is permitted to own and run the queue manager. This quorum mechanism is essential to prevent split-brain conditions, where more than one member might otherwise attempt to become active after a network or infrastructure failure. IBM MQ documents the use of

quorum-based availability for native HA queue managers, which is the Kubernetes-oriented HA model used in Cloud Pak deployments.

API Connect also uses clustered subsystem services and persistent distributed data stores. Its highly available topology requires a sufficient number of healthy members to retain a majority and continue safely through a member or node failure. Maintaining quorum ensures that the API management control-plane data remains consistent and that the deployment does not accept conflicting state changes during a partition. Consequently, production HA designs use the required replica counts and failure domains so the majority remains available. See IBM's native high availability queue manager documentation and API Connect high-availability overview.

QUESTION NO: 7

Which two storage access modes are valid Kubernetes persistent-volume access modes?

A. ReadWriteOnce (RWO)

B. ReadWriteMany (RWX)

C. ReadWriteAll (RWA)

D. ReadOnlyManyWrite (ROMW)

E. SharedReadWrite (SRW)

ANSWER: A B

Explanation:

ReadWriteOnce (RWO) and **ReadWriteMany (RWX)** are correct Kubernetes persistent-volume access modes. RWO permits the volume to be mounted as read-write by a single node. RWX permits the volume to be mounted as read-write by many nodes. Cloud Pak for Integration capabilities select storage access requirements according to their workload design, and the storage provider must support the required mode.

ReadWriteAll and ReadOnlyManyWrite are not Kubernetes persistent-volume access modes. Kubernetes also defines ReadOnlyMany (ROX) for volumes that can be mounted read-only by many nodes. See the Kubernetes documentation for persistent-volume access modes.

QUESTION NO: 8

What are two capabilities of the IBM Cloud Pak foundational services operator?

A. Messaging service to get robust and reliable messaging services.

B. Automation assets service to store, manage, and retrieve integration assets.

C. License Service that reports the license use of the product and its underlying product details that are deployed in the containerized environment.

D. API management service for managing the APIs created on API Connect.

E. IAM services for authentication and authorization.

ANSWER: C E

Explanation:

The IBM Cloud Pak foundational services operator installs and manages shared platform services used by IBM Cloud Paks. **License Service that reports the license use of the product and its underlying product details that are deployed in the containerized environment.** is a foundational-services capability because IBM License Service collects and makes available license-metering information for containerized IBM software. This enables administrators to view and report usage for products deployed in the cluster, supporting licensing compliance and audit requirements. **IAM services for**

authentication and authorization. is also a core capability because foundational services provide common identity and access-management components, including authentication and role-based access control for Cloud Pak users and platform services. These shared IAM capabilities allow products installed on the same OpenShift environment to use consistent user access and authorization controls rather than implementing identity management independently. IBM documentation identifies both the licensing service and Identity and Access Management as foundational services available to Cloud Pak deployments. See the IBM License Service documentation and the IBM Identity and Access Management documentation.

QUESTION NO: 9

Which of the following would contain mqsc commands for queue definitions to be executed when new MQ containers are deployed?

A. MORegistry

B. CCDTJSON

C. OperatorImage

D. ConfigMap

ANSWER: D

Explanation:

ConfigMap is correct because the IBM MQ Operator can obtain MQSC configuration from a Kubernetes ConfigMap when it creates or starts a queue manager in a container. MQSC is the administrative command language used to define IBM MQ objects, including local queues, remote queues, aliases, topics, channels, listeners, and related queue-manager settings. By placing the required commands in a ConfigMap and referencing that resource in the QueueManager custom resource, an administrator makes the configuration declarative and available during deployment. The operator mounts or supplies that configuration to the queue-manager container so that the commands are applied as part of the queue manager's initialization process. This approach is particularly useful in Cloud Pak for Integration because queue definitions can be version-controlled alongside deployment manifests and consistently reapplied when environments are built or recreated. ConfigMaps are Kubernetes resources intended for non-sensitive configuration data, which is appropriate for ordinary MQSC object-definition commands. IBM documents the use of an MQSC ConfigMap with the MQ Operator in its MQSC ConfigMap documentation. General Kubernetes behavior for ConfigMaps is also described in the Kubernetes ConfigMap documentation.

QUESTION NO: 10

Starling with Common Services 3.6, which two monitoring service modes are available?

A. OCP Monitoring

B. OpenShift Common Monitoring
C CP4I Monitoring

C. CS Monitoring

D. Grafana Monitoring

ANSWER: A C

Explanation:

The available monitoring modes are **OCP Monitoring** and **CS Monitoring**. OCP Monitoring uses the monitoring capabilities supplied by the Red Hat OpenShift Container Platform. This mode is appropriate when the cluster's platform monitoring stack is used to collect, retain, and query metrics for workloads and infrastructure in the Cloud Pak environment.

CS Monitoring uses the monitoring service supplied by IBM Cloud Platform Common Services. It provides the Common Services monitoring deployment and its associated components for collecting and exposing metrics in environments where that service mode is selected. The selection of a mode determines which monitoring stack is deployed and used by the

Cloud Pak installation; these are the two supported monitoring-service modes documented for Common Services version 3.6.

IBM's cluster-monitoring documentation for Common Services describes the configuration and operation of these monitoring approaches. See IBM Cloud Pak documentation: Cluster monitoring for Common Services 3.6.x and IBM Cloud Pak for Integration documentation: Monitoring.

QUESTION NO: 11

If the App Connect Operator is installed in a restricted network, which statement is true?

A. Simple Storage Service (S3) can not be used for BAR files.

B. Ephemeral storage can not be used for BAR files.

C. Only Ephemeral storage is supported for BAR files.

D. Persistent Claim storage can not be used for BAR files.

ANSWER: A

Explanation:

Simple Storage Service (S3) can not be used for BAR files. An App Connect Enterprise integration server needs access to the BAR file while it is being created or updated by the App Connect Operator. In a restricted-network deployment, outbound connectivity to external services is intentionally unavailable or tightly controlled. Therefore, using an S3-hosted BAR file is not supported as a BAR source in this scenario, because retrieving it requires access to the relevant object-storage endpoint and associated authentication services. Instead, BAR content must be made available from storage that is accessible within the restricted OpenShift environment, such as ephemeral storage for the integration server lifecycle or a persistent volume claim when the BAR files must survive pod replacement and be retained independently. IBM's App Connect Operator documentation describes the supported BAR-file storage choices and the limitations that apply to restricted-network installations. See IntegrationServer custom resource reference and Installing in a restricted network.

QUESTION NO: 12

Which statement describes the Aspera High Speed Transfer Server (HSTS) within IBM Cloud Pak for Integration?

A. HSTS allows an unlimited number of concurrent users to transfer files of up to 500GB at high speed using an Aspera client.

B. HSTS allows an unlimited number of concurrent users to transfer files of up to 100GB at high speed using an Aspera client.

C. HSTS allows an unlimited number of concurrent users to transfer files of up to 1TB at highs peed using an Aspera client.

D. HSTS allows an unlimited number of concurrent users to transfer files of any size at high speed using an Aspera client.

ANSWER: D

Explanation:

HSTS allows an unlimited number of concurrent users to transfer files of any size at high speed using an Aspera client. This describes the role of IBM Aspera High Speed Transfer Server: it is an enterprise transfer-server capability that uses Aspera's FASP transport technology to move data at high speed over wide-area networks. Unlike conventional file-transfer approaches that can become inefficient when latency or packet loss is present, FASP is designed to fully use the available bandwidth while maintaining reliable, secure transfers. HSTS acts as the endpoint that receives and sends content for Aspera clients, supporting large-scale file movement without imposing a file-size limit as part of the server capability. Its architecture is also intended to scale for many simultaneous transfer users, making it appropriate for integration workflows where high-volume or very large payloads must be exchanged efficiently. The wording "any size" reflects the product capability; practical limits can still arise from the deployed infrastructure, available storage, operating-system constraints,

network capacity, and organizational policies rather than from a stated HSTS file-size ceiling. IBM's Aspera product overview describes the high-speed transfer technology and its enterprise use cases: IBM Aspera. IBM's documentation portal provides product documentation for configuration and operation of Aspera components: IBM Aspera documentation.

QUESTION NO: 13

Which command shows the current cluster version and available updates?

A. update

B. oc adm upgrade

C. adm update

D. upgrade

ANSWER: B

Explanation:

`oc adm upgrade` is the OpenShift administrative command used to inspect the cluster upgrade state. When run without an action subcommand, it reports the cluster's current version, the upgrade channel in use, available target releases, and whether the cluster is upgradeable. This gives an administrator the essential information needed to determine which OpenShift updates can be selected before beginning an upgrade. Cloud Pak for Integration runs on Red Hat OpenShift, so administrators use the OpenShift CLI to assess the underlying cluster version and available upgrade paths as part of planning platform maintenance. The command can also be used with additional arguments to select an upgrade target, but its default status output is specifically what identifies the current version and available updates. IBM Cloud Pak administration should account for supported OpenShift versions and product compatibility before any cluster update is initiated. See Red Hat's Updating a cluster using the CLI documentation and IBM's OpenShift upgrade guidance for Cloud Pak for Integration.

QUESTION NO: 14

Which two App Connect resources enable callable flows to be processed between an integration solution in a cluster and an integration server in an on-premise system?

A. Sync server

B. Connectivity agent

C. Kafka sync

D. Switch server

E. Routing agent

ANSWER: B D

Explanation:

Connectivity agent and **Switch server** are the App Connect resources used to support callable-flow processing across the boundary between a cluster-based integration solution and an on-premises integration server. A connectivity agent establishes an outbound, secure connection from the on-premises environment to the cluster. Because the connection is initiated from inside the on-premises network, it can traverse restrictive firewall environments without requiring inbound access to the local integration server. It provides the communications channel through which callable-flow requests and responses can pass.

A switch server works with that connection to route callable-flow requests to the appropriate reachable integration server. It acts as the switching component for the callable-flow infrastructure, enabling the integration solution in the cluster to locate and invoke a flow that is deployed in the connected on-premises environment. Together, these resources provide the secure

connectivity and request-routing capabilities required for hybrid callable flows. IBM documents callable-flow connectivity and the switch-server role in its App Connect guidance; see Connecting callable flows and Switch servers.

QUESTION NO: 15 - (DRAG DROP)

Select all that apply

What is the correct sequence of steps to delete IBM MQ from IBM Cloud Pak for Integration?

ANSWER:

Explanation:

To remove IBM MQ cleanly from IBM Cloud Pak for Integration, first delete the Queue Managers that are managed by the IBM MQ Operator. Queue Managers are Kubernetes custom resources, and the operator is responsible for reconciling and managing them while it remains installed. Removing those Queue Manager resources first ensures that their associated IBM MQ workloads and managed resources can be shut down and cleaned up in the supported order.

After the Queue Managers have been deleted, log in to the OpenShift cluster web console. In the project that contained the Queue Managers, open the *Installed Operators* area from *Operators*. This is where the cluster administrator locates the installed IBM MQ Operator subscription and its operator entry for that namespace. The final action is to uninstall the Operator. At that point, there are no remaining Queue Manager instances for it to manage, so the operator can be removed without leaving managed instances in place.

This ordering follows the general Kubernetes operator lifecycle principle that operator-managed custom resources should be removed before the controller/operator that manages them is uninstalled. IBM's Cloud Pak for Integration documentation is available at IBM Cloud Pak for Integration documentation, and IBM MQ container and operator documentation is available at IBM MQ documentation. The required sequence therefore begins with Queue Manager deletion and ends with IBM MQ Operator uninstallation.