**Vendor:**Mulesoft

**Exam Code:**MCIA-LEVEL-1-MAINTENANCE

**Exam Name:**MuleSoft Certified Integration Architect - Level 1 MAINTENANCE

**Version:**Demo

**QUESTION 1**

A Mule application is being designed for deployment to a single CloudHub worker. The Mule application will have a flow that connects to a SaaS system to perform some operations each time the flow is invoked.

The SaaS system connector has operations that can be configured to request a short-lived token (fifteen minutes) that can be reused for subsequent connections within the fifteen minute time window. After the token expires, a new token

must be requested and stored.

What is the most performant and idiomatic (used for its intended purpose) Anypoint Platform component or service to use to support persisting and reusing tokens in the Mule application to help speed up reconnecting the Mule application to the SaaS application?

A. Nonpersistent object store

B. Persistent object store

C. Variable

D. Database

Correct Answer: D

Reference: https://docs.mulesoft.com/mule-runtime/4.4/reconnection-strategy-about

---

**QUESTION 2**

An organization has deployed runtime fabric on an eight note cluster with performance profile. An API uses and non persistent object store for maintaining some of its state data. What will be the impact to the stale data if server crashes?

A. State data is preserved

B. State data is rolled back to a previously saved version

C. State data is lost

D. State data is preserved as long as more than one more is unaffected by the crash

Correct Answer: D

---

**QUESTION 3**

An organization uses one specific CloudHub (AWS) region for all CloudHub deployments. How are CloudHub workers assigned to availability zones (AZs) when the organization\\\'s Mule applications are deployed to CloudHub in that region?

A. Workers belonging to a given environment are assigned to the same AZ within that region.

B. AZs are selected as part of the Mule application\\\'s deployment configuration.

C. Workers are randomly distributed across available AZs within that region.

D. An AZ is randomly selected for a Mule application, and all the Mule application\\'s CloudHub workers are assigned to that one AZ

Correct Answer: C

Correct answer is Workers are randomly distributed across available AZs within that region. This ensure high availability for deployed mule applications Mulesoft documentation reference : https://docs.mulesoft.com/runtime-manager/cloudhub-hadr

---

**QUESTION 4**

Customer has deployed mule applications to different customer hosted mule run times. Mule applications are managed from Anypoint platform.

What needs to be configured to monitor these Mule applications from Anypoint monitoring and what sends monitoring data to Anypoint monitoring?

A. Enable monitoring of individual applications from runtime manager application settings Runtime manager agent sends monitoring data from the mule applications to Anypoint monitoring

B. Install runtime manager agent on each mule runtime Runtime manager agent since monitoring data from the mule applications to Anypoint monitoring

C. Anypoint monitoring agent on each mule runtime Anypoint monitoring agent sends monitoring data from the mule applications to Anypoint monitoring

D. By default, Anypoint monitoring agent will be installed on each Mule run time Anypoint Monitoring agent automatically sends monitoring data from the Mule applications to Anypoint monitoring

Correct Answer: C

---

**QUESTION 5**

An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?
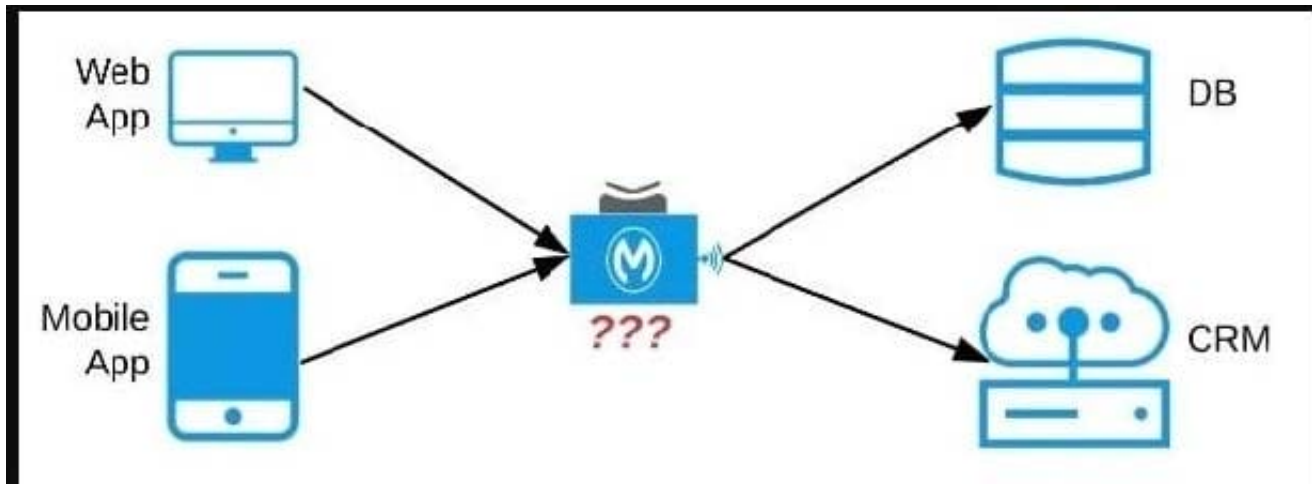
A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
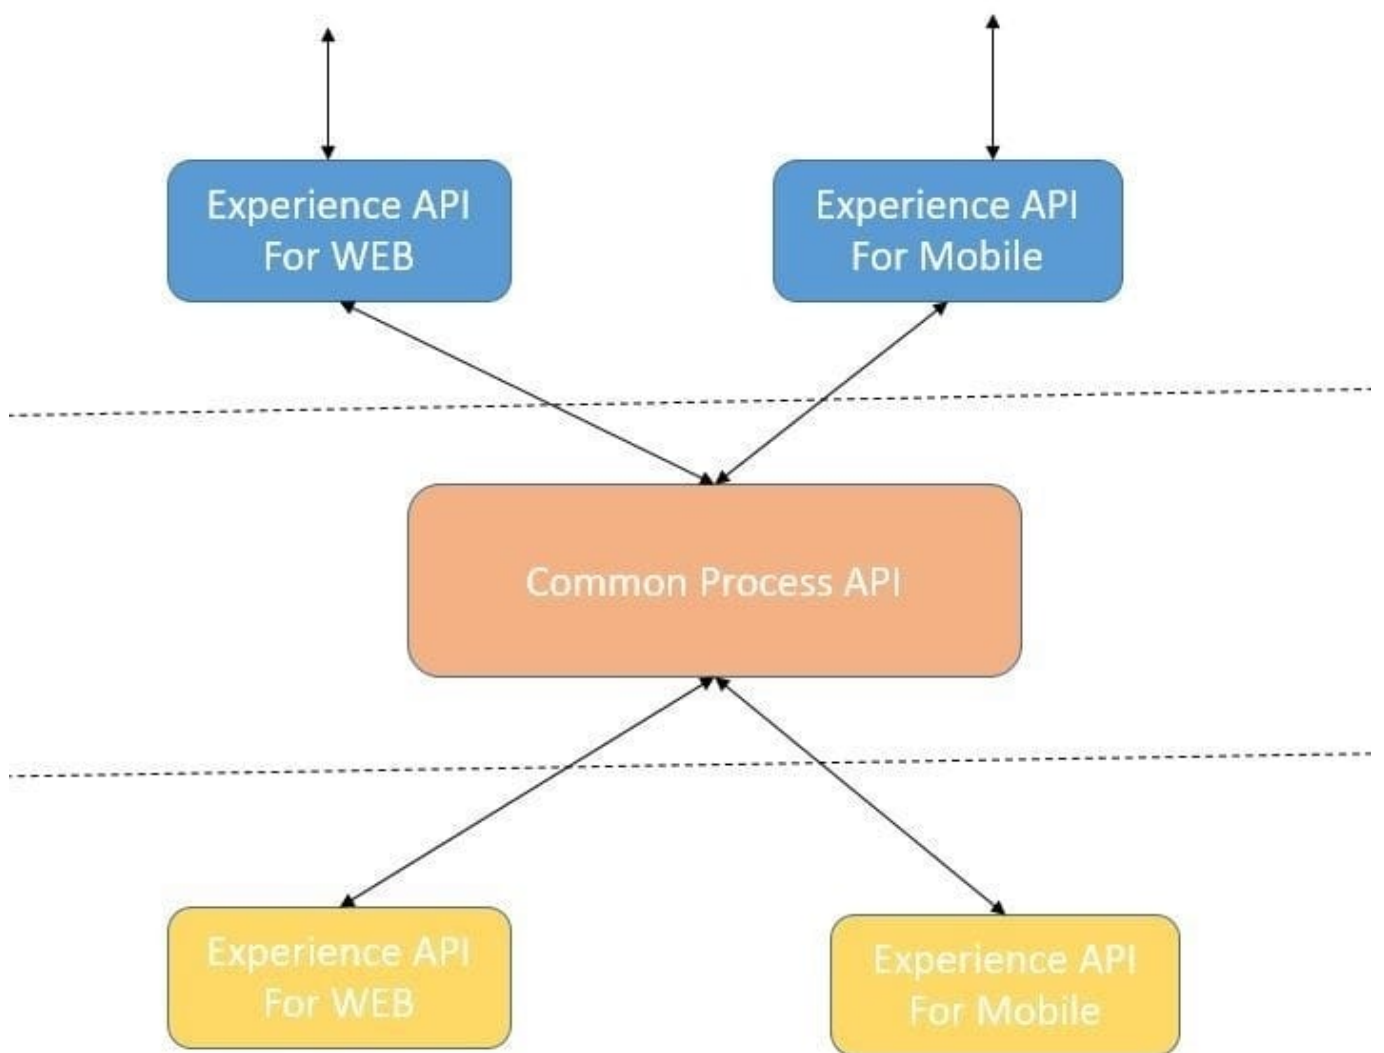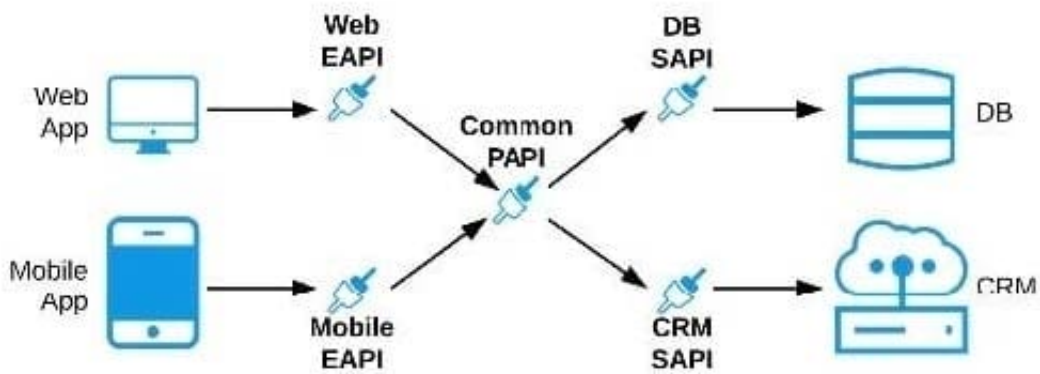
B. One set of APIs (Experience API, Process API, and System API) for the web app, and another set for the mobile app.

C. Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system

D. A common Experience API used by both the web and mobile apps, but separate Process APIs for the web and mobile apps that interact with the database and the CRM System.

Correct Answer: C

Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach. Option : Separate set of API\\'s for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code. Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system Description automatically generated with low confidence Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach. Option : Separate set of API\\'s for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code. Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system

---

**QUESTION 6**

A system API EmployeeSAPI is used to fetch employee\'s data from an underlying SQL database.

The architect must design a caching strategy to query the database only when there is an update to the employees

stable or else return a cached response in order to minimize the number of redundant transactions being handled by the

database.

What must the architect do to achieve the caching objective?

A. Use an On Table Row on employees table and call invalidate cache Use an object store caching strategy and expiration interval to empty

B. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and expiration interval to empty

C. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and set expiration interval to 1-hour

D. Use an on table rule on employees table call invalidate cache and said new employees data to cache Use an object store caching strategy and set expiration interval to 1-hour

Correct Answer: A

---

**QUESTION 7**

Refer to the exhibit.
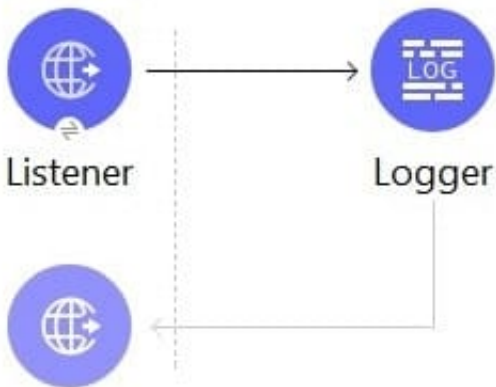


The HTTP Listener and the Logger are being handled from which thread pools respectively?

A. CPU_INTENSIVE and Dedicated Selector pool

B. UBER and NONBLOCKING

C. Shared Selector Pool and CPU LITE

D. BLOCKING _IO and UBER

Correct Answer: C

---

**QUESTION 8**

An organization is designing a mule application to support an all or nothing transaction between serval database operations and some other connectors so that they all roll back if there is a problem with any of the connectors. Besides the database connector , what other connector can be used in the transaction.

A. VM

B. Anypoint MQ

C. SFTP

D. ObjectStore

Correct Answer: A

Correct answer is VM VM support Transactional Type. When an exception occur, The transaction rolls back to its original state for reprocessing. This feature is not supported by other connectors. Here is additional information about Transaction management:

| | Shared Load Balancer | Dedicated Load Balancer |
|---|---|---|
| VPC | Shared VPC (Mulesoft) | VPC (Customer) |
| Default Load Balancer | Cloudhub provides Deault Shared Load Balancer available in All Environment | Need to Purchase |
| Organization Use | Multiple Oragnization | Specific to Organization |
| Certificate | Mulesoft Certificate | Organization Certificate |
| TLS Support | Yes | Yes |
| URL Mapping | Fixed URL Mapping | Customer URL Mapping |
| Timeout | 30 Sec Session Timeout | Custom Timeout |
| Ports | Public Port {80 : 8081, 443 : 8082} | Private Port {80 : 8091, 443 : 8092} |
| Fashion | Round Robin | Round Robin |
| Supports HTTPS Protocol | Yes | Yes |
| Worker Assignment | No | Yes |
| IP Blacklisting/ Whitelisting | No | Yes |
| | https://docs.mulesoft.com/runtime-manager/lb-whitelists | |
| Configure Custom Domain | No | Yes |
| Custom Certificate | No | Yes |
| Rate Limit | Lower Rate Limit and applied According to Region | Higher Rate Limit Threshold |
| VPC | Anypoint VPC optional | Can't Use DLB without Anypoint VPC |

**QUESTION 9**

A company wants its users to log in to Anypoint Platform using the company\\'s own internal user credentials. To achieve this, the company needs to integrate an external identity provider (IdP) with the company\\'s Anypoint Platform master organization, but SAML 2.0 CANNOT be used. Besides SAML 2.0, what single-sign-on standard can the company use to integrate the IdP with their Anypoint Platform master organization?

A. SAML 1.0

B. OAuth 2.0

C. Basic Authentication

D. OpenID Connect

Correct Answer: D

As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). Configure identity management using one of the following single sign-on standards: 1) OpenID Connect: End user identity verification by an authorization server including SSO 2) SAML 2.0: Web-based authorization including cross-domain SSO

---

**QUESTION 10**

An organization is implementing a Quote of the Day API that caches today\\'s quote. What scenario can use the CloudHub Object Store connector to persist the cache\\'s state?

A. When there is one deployment of the API implementation to CloudHub and another one to customer hosted mule runtime that must share the cache state.

B. When there are two CloudHub deployments of the API implementation by two Anypoint Platform business groups to the same CloudHub region that must share the cache state.

C. When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.

D. When there are three CloudHub deployments of the API implementation to three separate CloudHub regions that must share the cache state.

Correct Answer: C

Object Store Connector is a Mule component that allows for simple key-value storage. Although it can serve a wide variety of use cases, it is mainly design for: - Storing synchronization information, such as watermarks. - Storing temporal information such as access tokens. - Storing user information. Additionally, Mule Runtime uses Object Stores to support some of its own components, for example: - The Cache module uses an Object Store to maintain all of the cached data.

- The OAuth module (and every OAuth enabled connector) uses Object Stores to store the access and refresh tokens. Object Store data is in the same region as the worker where the app is initially deployed. For example, if you deploy to the Singapore region, the object store persists in the Singapore region. MuleSoft Reference : https://docs.mulesoft.com/object-store-connector/1.1/ Data can be shared between different instances of the Mule application. This is not recommended for Inter Mule app communication. Coming to the question, object store cannot be used to share cached data if it is deployed as separate Mule applications or deployed under separate Business Groups. Hence correct answer is When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.

---

**QUESTION 11**

A company is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTP5 POST and must be

acknowledged immediately.

Once acknowledged the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to the rejections from the back-end system will need to be processed manually (outside the banking system).

The mule application will be deployed to a customer hosted runtime and will be able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization\\'s firewall. The back-end system has a track record of

unreliability due to both minor network connectivity issues and longer outages.

Which combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

A. One or more On Error scopes to assist calling the back-end system An Untill successful scope containing VM components for long retries A persistent dead-letter VM queue configure in Cloud hub

B. An Until Successful scope to call the back-end system One or more ActiveMQ long-retry queues One or more ActiveMQ dead-letter queues for manual processing

C. One or more on-Error scopes to assist calling the back-end system one or more ActiveMQ long-retry queues A persistent dead-letter Object store configuration in the CloudHub object store service

D. A batch job scope to call the back in system An Untill successful scope containing Object Store components for long retries. A dead-letter object store configured in the Mule application

Correct Answer: B

---

**QUESTION 12**

A mule application designed to fulfil two requirements

a) Processing files are synchronously from an FTPS server to a back-end database using VM intermediary queues for load balancing VM events b) Processing a medium rate of records from a source to a target system using batch job scope

Considering the processing reliability requirements for FTPS files, how should VM queues be configured for processing files as well as for the batch job scope if the application is deployed to Cloudhub workers?

A. Use Cloud hub persistent queues for FTPS files processing There is no need to configure VM queues for the batch jobs scope as it uses by default the worker\\'s disc for VM queueing

B. Use Cloud hub persistent VM queue for FTPS file processing There is no need to configure VM queues for the batch jobs scope as it uses by default the worker\\'s JVM memory for VM queueing

C. Use Cloud hub persistent VM queues for FTPS file processing Disable VM queue for the batch job scope

D. Use VM connector persistent queues for FTPS file processing Disable VM queue for the batch job scope

Correct Answer: C