



## **MuleSoft**

### **Exam Questions MCIA-Level-1**

MuleSoft Certified Integration Architect - Level 1

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

#### NEW QUESTION 1

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

**Answer:** A

#### NEW QUESTION 2

An organization has implemented a continuous integration (CI) lifecycle that promotes Mule applications through code, build, and test stages. To standardize the organization's CI journey, a new dependency control approach is being designed to store artifacts that include information such as dependencies, versioning, and build promotions.

To implement these process improvements, the organization will now require developers to maintain all dependencies related to Mule application code in a shared location.

What is the most idiomatic (used for its intended purpose) type of system the organization should use in a shared location to standardize all dependencies related to Mule application code?

- A. A MuleSoft-managed repository at repository.mulesoft.org
- B. A binary artifact repository
- C. API Community Manager
- D. The Anypoint Object Store service at cloudhub.io

**Answer:** C

#### NEW QUESTION 3

A trading company handles millions of requests a day. Due to nature of its business, it requires excellent performance and reliability within its application.

For this purpose, company uses a number of event-based API's hosted on various mule clusters that communicate across a shared message queue sitting within its network.

Which method should be used to meet the company's requirement for its system?

- A. XA transactions and XA connected components
- B. JMS transactions
- C. JMS manual acknowledgements with a reliability pattern
- D. VM queues with reliability pattern

**Answer:** C

#### NEW QUESTION 4

A global organization operates datacenters in many countries. There are private network links between these datacenters because all business data (but NOT metadata) must be exchanged over these private network connections.

The organization does not currently use AWS in any way.

The strategic decision has Just been made to rigorously minimize IT operations effort and investment going forward.

What combination of deployment options of the Anypoint Platform control plane and runtime plane(s) best serves this organization at the start of this strategic journey?

- A. MuleSoft-hosted Anypoint Platform control plane CloudHub Shared Worker Cloud in multiple AWS regions
- B. Anypoint Platform - Private Cloud Edition Customer-hosted runtime plane in each datacenter
- C. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in multiple AWS regions
- D. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in each datacenter

**Answer:** D

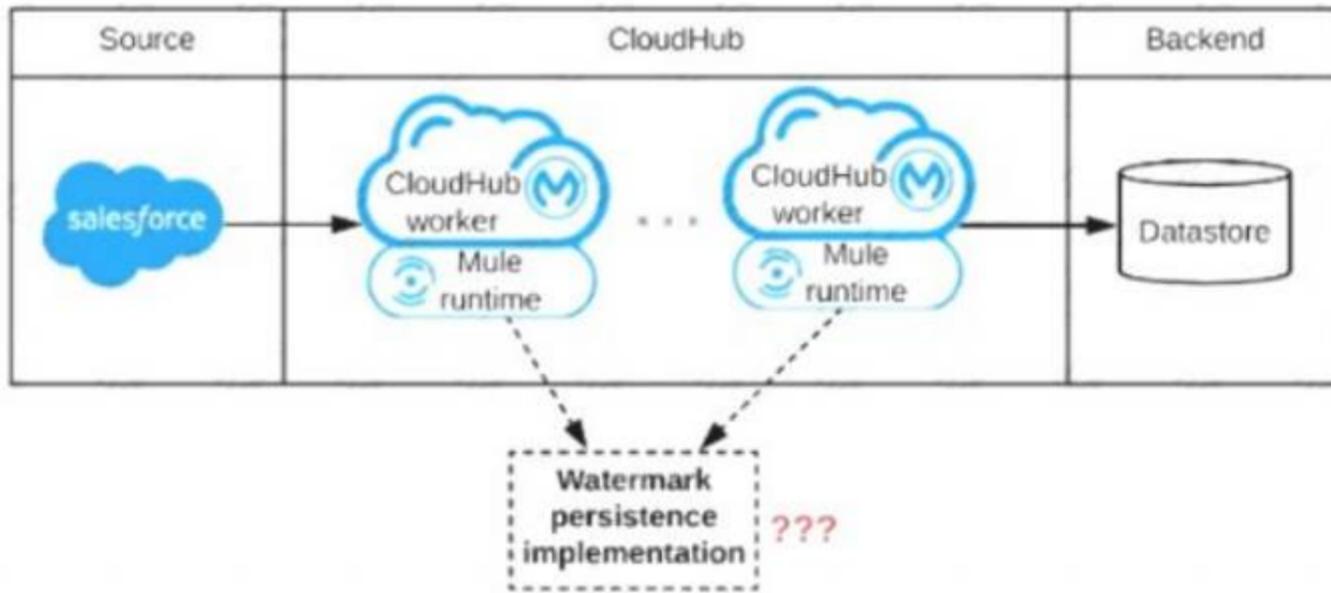
#### Explanation:

Correct answer is MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in each datacenter There are two things to note about the question which can help us figure out correct answer.. \* Business data must be exchanged over these private network connections which means we can not use MuleSoft provided Cloudhub option. So we are left with either customer hosted runtime in external cloud provider or customer hosted runtime in their own premises. As customer does not use AWS at the moment. Hence that don't have the immediate option of using Customer-hosted runtime plane in multiple AWS regions. hence the most suitable option for runtime plane is Customer-hosted runtime plane in each datacenter \* Metadata has no limitation to reside in organization premises. Hence for control plane MuleSoft hosted Anypoint platform can be used as a strategic solution.

Hybrid is the best choice to start. Mule hosted Control plane and Customer hosted Runtime to start with. Once they mature in cloud migration, everything can be in Mule hosted.

#### NEW QUESTION 5

Refer to the exhibit.



A Mule application is being designed to be deployed to several CloudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.

A watermark will be used to only retrieve those Satesforce Accounts that have been modified since the last time the integration logic ran.

What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?

- A. Persistent Anypoint MQ Queue
- B. Persistent Object Store
- C. Persistent Cache Scope
- D. Persistent VM Queue

**Answer: B**

**Explanation:**

\* An object store is a facility for storing objects in or across Mule applications. Mule uses object stores to persist data for eventual retrieval.

\* Mule provides two types of object stores:

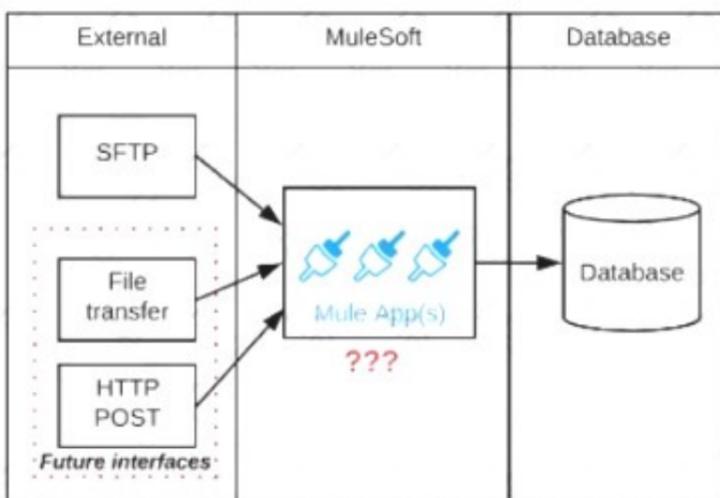
- 1) In-memory store – stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime.
- 2) Persistent store – Mule persists data when an object store is explicitly configured to be persistent.

In a standalone Mule runtime, Mule creates a default persistent store in the file system. If you do not specify an object store, the default persistent object store is used.

MuleSoft Reference: <https://docs.mulesoft.com/mule-runtime/3.9/mule-object-stores>

**NEW QUESTION 6**

Refer to the exhibit.



A business process involves the receipt of a file from an external vendor over SFTP. The file needs to be parsed and its content processed, validated, and ultimately persisted to a database. The delivery mechanism is expected to change in the future as more vendors send similar files using other mechanisms such as file transfer or HTTP POST.

What is the most effective way to design for these requirements in order to minimize the impact of future change?

- A. Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- B. Create a Process API to receive the file and process it using a MuleSoft Batch Job while delegating the data save process to a System API
- C. Create an API that receives the file and invokes a Process API with the data contained In the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed
- D. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing

**Answer: C**

**Explanation:**

\* Scatter-Gather is used for parallel processing, to improve performance. In this scenario, input files are coming from different vendors so mostly at different times. Goal here is to minimize the impact of future change. So scatter Gather is not the correct choice.

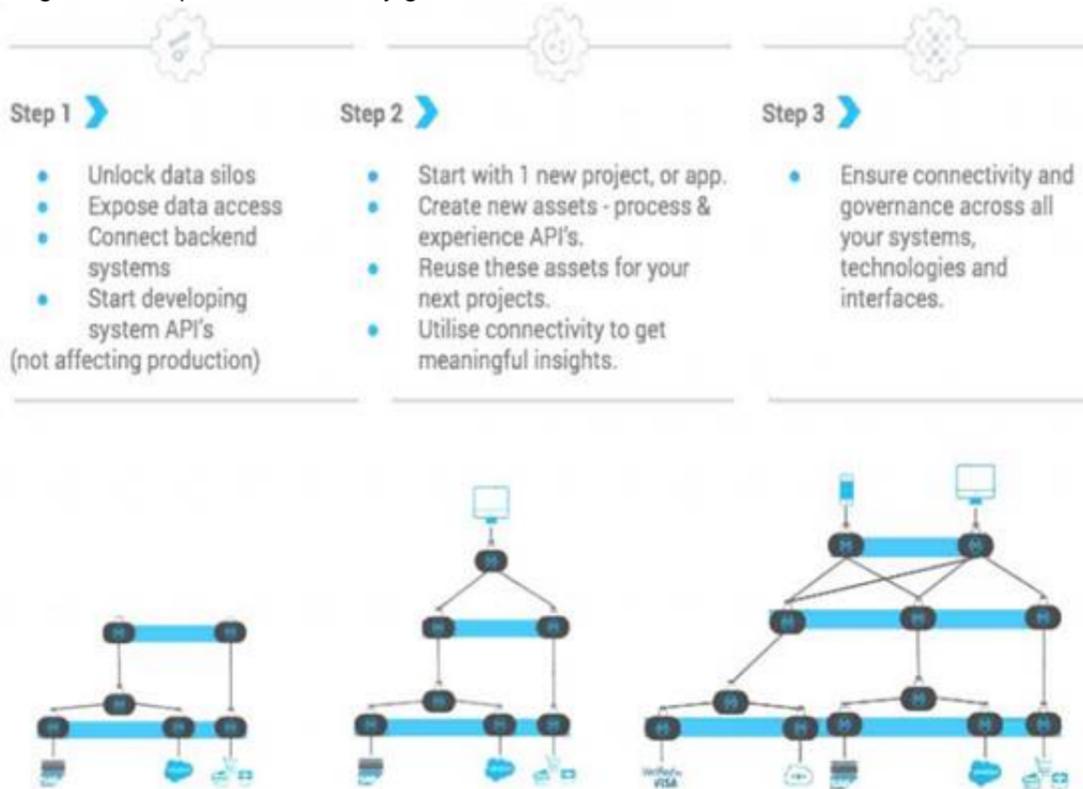
\* If we use 1 API to receive all files from different Vendors, any new vendor addition will need changes to that 1 API to accommodate new requirements. So Option A and C are also ruled out.

\* Correct answer is Create an API that receives the file and invokes a Process API with the data contained in the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed. Answer to this question lies in the API led connectivity approach.

\* API-led connectivity is a methodical way to connect data to applications through a series of reusable and purposeful modern APIs that are each developed to play a specific role – unlock data from systems, compose data into processes, or deliver an experience. System API : System API tier, which provides consistent, managed, and secure access to backend systems. Process APIs : Process APIs take core assets and combines them with some business logic to create a higher level of value. Experience APIs : These are designed specifically for consumption by a specific end-user app or device.

So in case of any future plans, organization can only add experience API on addition of new Vendors, which reuse the already existing process API. It will keep impact minimal.

Diagram Description automatically generated



**NEW QUESTION 7**

In a Mule Application, a flow contains two (2) JMS consume operations that are used to connect to a JMS broker and consume messages from two(2) JMS destination. The Mule application then joins the two JMS messages together.

The JMS broker does not implement high availability (HA) and periodically experiences scheduled outages of upto 10 mins for routine maintenance.

What is the most idiomatic (used for its intended purpose) way to build the mule flow so it can best recover from the expected outages?

- A. Configure a reconnection strategy for the JMS connector
- B. Enclose the two(2) JMS operation in an Until Successful scope
- C. Consider a transaction for the JMS connector
- D. Enclose the two(2) JMS operations in a Try scope with an Error Continue error handler

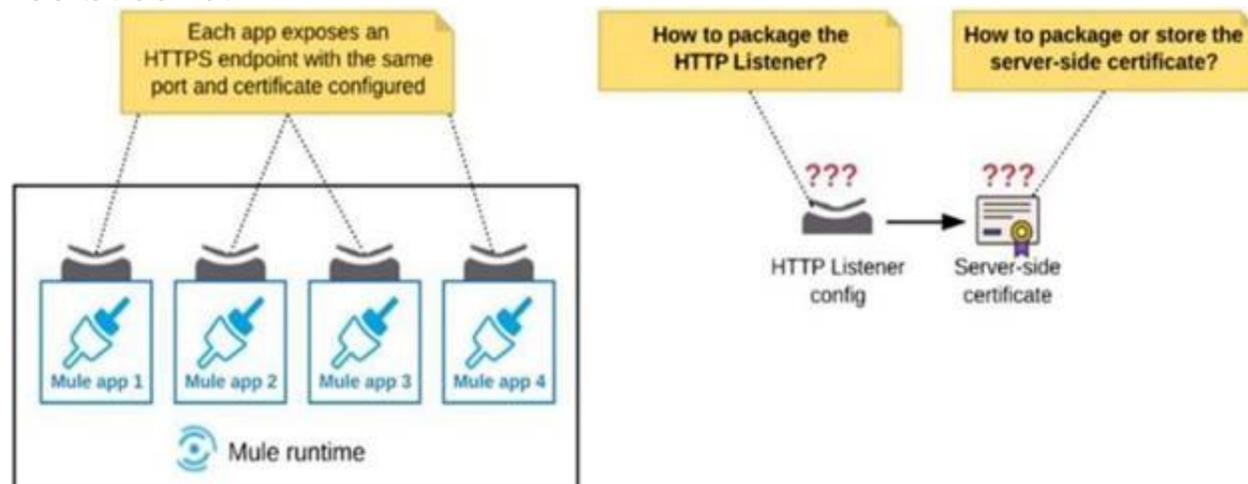
**Answer: A**

**Explanation:**

When an operation in a Mule application fails to connect to an external server, the default behavior is for the operation to fail immediately and return a connectivity error. You can modify this default behavior by configuring a reconnection strategy for the operation. You can configure a reconnection strategy for an operation either by modifying the operation properties or by modifying the configuration of the global element for the operation. The following are the available reconnection strategies and their behaviors: None is the default behavior, which immediately returns a connectivity error if the attempt to connect is unsuccessful Standard (reconnect) Sets the number of reconnection attempts and the interval at which to execute them before returning a connectivity error Forever (reconnect-forever) Attempts to reconnect continually at a given interval

**NEW QUESTION 8**

Refer to the exhibit.



An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a server-side certificate that rotates often.

What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

- A. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint Package the server-side certificate in ALL Mule APPLICATIONS that need to expose an HTTPS endpoint
- B. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint
- C. Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION
- D. Package an HTTPS Listener configuration In all Mule APPLICATIONS that need to expose an HTTPS endpoint Package the server-side certificate in a NEW Mule DOMAIN project
- E. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing It from all Mule applications that need to expose an HTTPS endpoint
- F. Package the server-side certificate in the SAME Mule DOMAIN project Go to Set

Answer: B

**Explanation:**

In this scenario, both A & C will work, but A is better as it does not require repackage to the domain project at all.

Correct answer is Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint. Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION.

What is Mule Domain Project?

\* A Mule Domain Project is implemented to configure the resources that are shared among different projects. These resources can be used by all the projects associated with this domain. Mule applications can be associated with only one domain, but a domain can be associated with multiple projects. Shared resources allow multiple development teams to work in parallel using the same set of reusable connectors. Defining these connectors as shared resources at the domain level allows the team to:

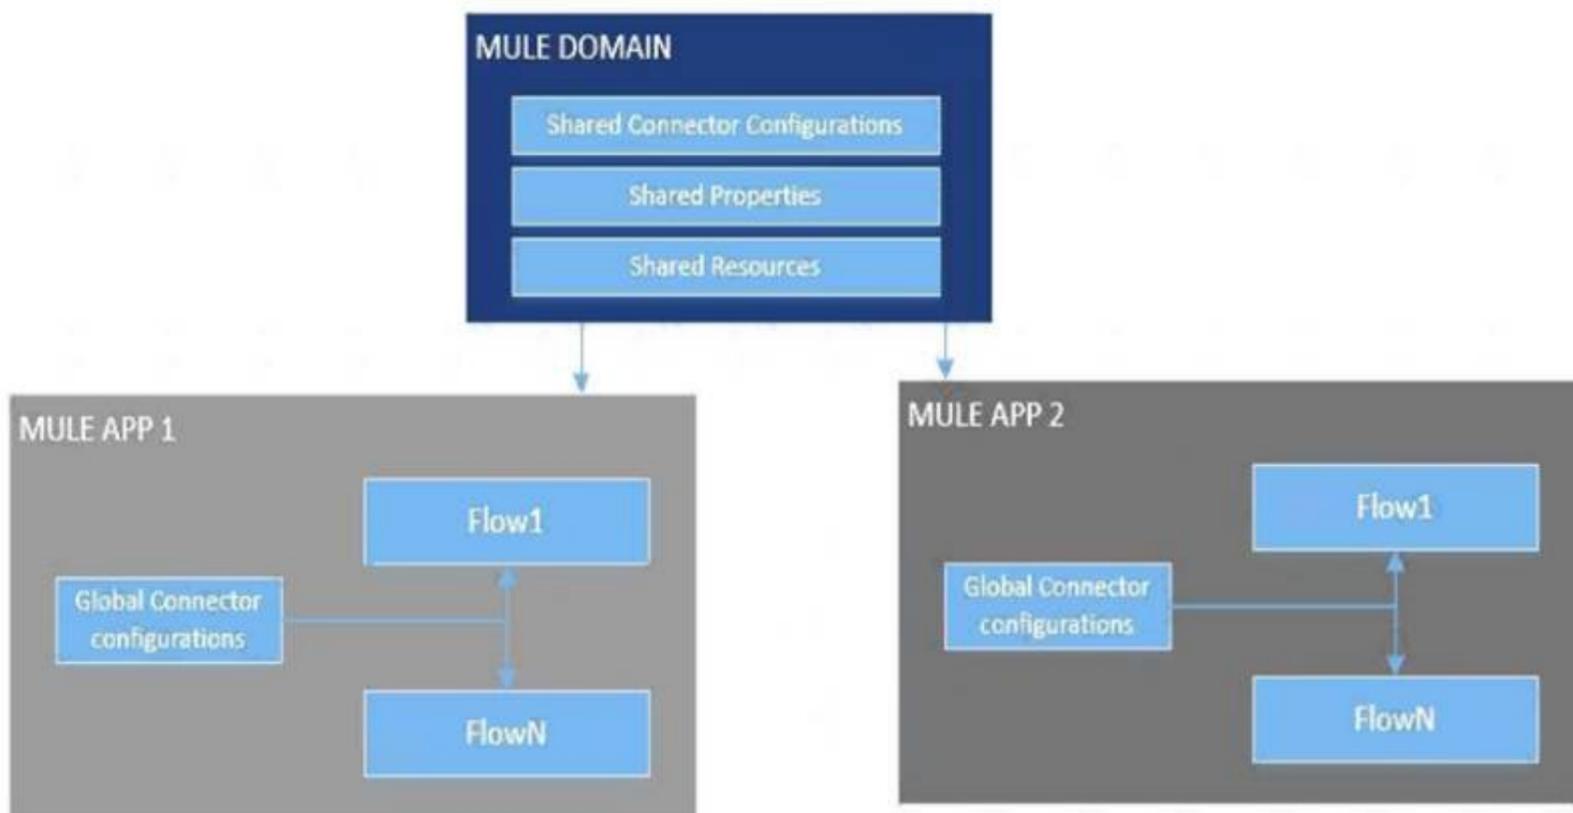
- Expose multiple services within the domain through the same port.
- Share the connection to persistent storage.
- Share services between apps through a well-defined interface.
- Ensure consistency between apps upon any changes because the configuration is only set in one place.

\* Use domains Project to share the same host and port among multiple projects. You can declare the http connector within a domain project and associate the domain project with other projects. Doing this also allows to control thread settings, keystore configurations, time outs for all the requests made within multiple applications. You may think that one can also achieve this by duplicating the http connector configuration across all the applications. But, doing this may pose a nightmare if you have to make a change and redeploy all the applications.

\* If you use connector configuration in the domain and let all the applications use the new domain instead of a default domain, you will maintain only one copy of the http connector configuration. Any changes will require only the domain to be redeployed instead of all the applications.

You can start using domains in only three steps:

- 1) Create a Mule Domain project
- 2) Create the global connector configurations which needs to be shared across the applications inside the Mule Domain project
- 3) Modify the value of domain in mule-deploy.properties file of the applications Graphical user interface Description automatically generated



Use a certificate defined in already deployed Mule domain Configure the certificate in the domain so that the API proxy HTTPS Listener references it, and then deploy the secure API proxy to the target Runtime Fabric, or on-premises target. (CloudHub is not supported with this approach because it does not support Mule domains.)

**NEW QUESTION 9**

A Mule application is built to support a local transaction for a series of operations on a single database. The mule application has a Scatter-Gather scope that participates in the local transaction.

What is the behavior of the Scatter-Gather when running within this local transaction?

- A. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
- B. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back
- C. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
- D. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back

Answer: A

**NEW QUESTION 10**

An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders. The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS).

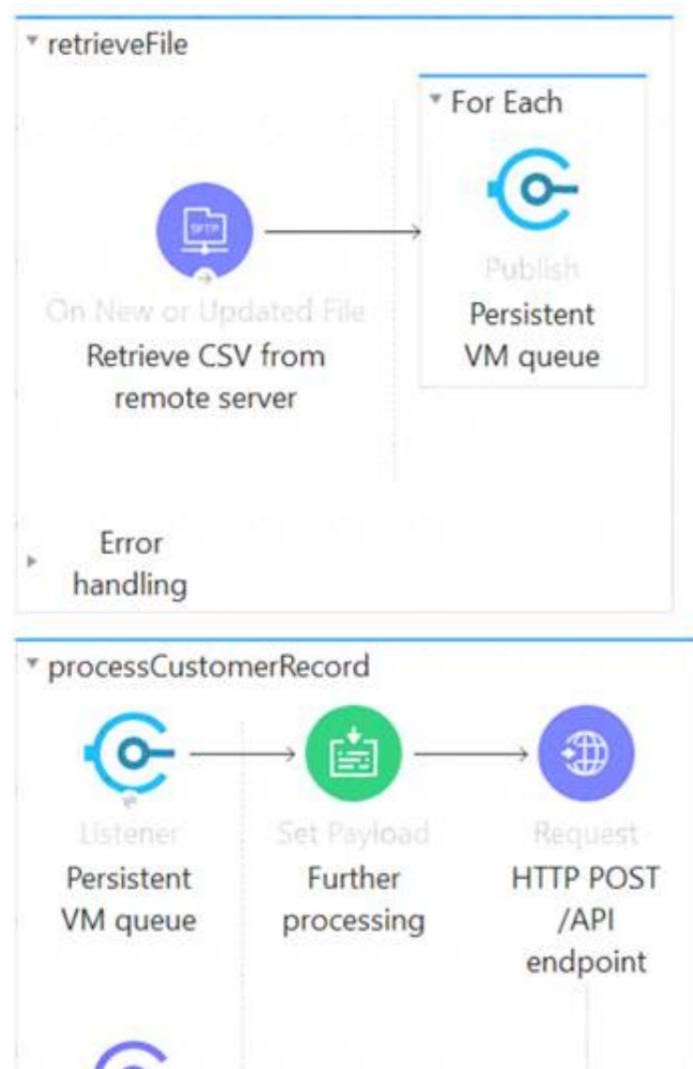
At a minimum, what must be stored in each truststore and keystore of these two Mule applications to properly support two-way TLS between the two Mule applications while properly protecting each Mule application's keys?

- A. Orders API truststore: The Orders API public key Process Orders keystore: The Process Orders private key and public key
- B. Orders API truststore: The Orders API private key and public key Process Orders keystore: The Process Orders private key public key
- C. Orders API truststore: The Process Orders public key Orders API keystore: The Orders API private key and public key Process Orders truststore: The Orders API public key Process Orders keystore: The Process Orders private key and public key
- D. Orders API truststore: The Process Orders public key Orders API keystore: The Orders API private key Process Orders truststore: The Orders API public key Process Orders keystore: The Process Orders private key

Answer: C

### NEW QUESTION 10

Refer to the exhibit.



This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow's VM Listener receives messages from the same VM queue and then processes each message separately.

How are messages routed to the cloudhub workers as messages are received by the VM Listener?

- A. Each message is routed to ONE of the Cloudhub workers in a DETERMINISTIC round robin fashion thereby EXACTLY BALANCING messages among the cloudhub workers
- B. Each messages routes to ONE of the available Clouhub workers in a NON- DETERMINISTIC non round-robin fashion thereby APPROXIMATELY BALANCING messages among the cloudhub workers
- C. Each message is routed to the SAME Cloudhub worker that retrieved the file, thereby BINDING ALLmessages to ONLY that ONE Cloudhub worker
- D. Each message is duplicated to ALL of the Cloudhub workers, thereby SHARING EACH message with ALL the Cloudhub workers.

Answer: B

### NEW QUESTION 12

An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments. What must the testing team do before they can start integration testing the API in the Staging environment?

- A. They must access the API portal and create an API notebook using the Client ID and Client Secret supplied by the API portal in the Staging environment
- B. They must request access to the API instance in the Staging environment and obtain a Client ID and Client Secret to be used for testing the API
- C. They must be assigned as an API version owner of the API in the Staging environment
- D. They must request access to the Staging environment and obtain the Client ID and Client Secret for that environment to be used for testing the API

Answer: B

#### Explanation:

- \* It's mentioned that the API is governed by a Client ID Enforcement policy in all environments.
- \* Client ID Enforcement policy allows only authorized applications to access the deployed API implementation.
- \* Each authorized application is configured with credentials: client\_id and client\_secret.
- \* At runtime, authorized applications provide the credentials with each request to the API implementation. MuleSoft Reference: <https://docs.mulesoft.com/api-manager/2.x/policy-mule3-client-id-based-policies>

### NEW QUESTION 17

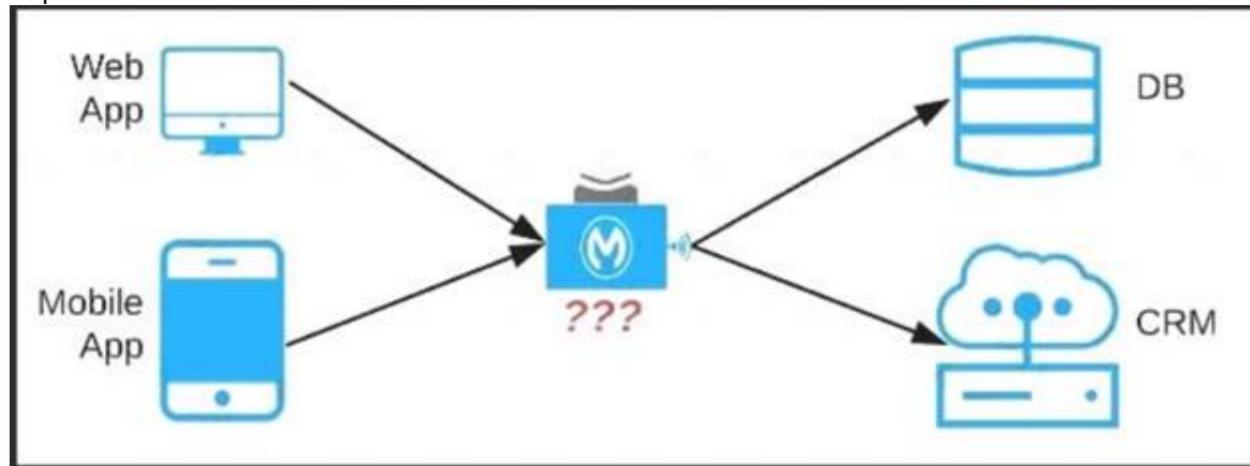
As a part of project requirement, Java Invoke static connector in a mule 4 application needs to invoke a static method in a dependency jar file. What are two ways to add the dependency to be visible by the connectors class loader?  
 (Choose two answers)

- A. In the Java Invoke static connector configuration, configure a path and name of the dependency jar file
- B. Add the dependency jar file to the java classpath by setting the JVM parameters
- C. Use Maven command to include the dependency jar file when packaging the application
- D. Configure the dependency as a shared library in the project POM
- E. Update mule-artefact.json to export the Java package

Answer: BD

**NEW QUESTION 18**

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.

Answer: C

**Explanation:**

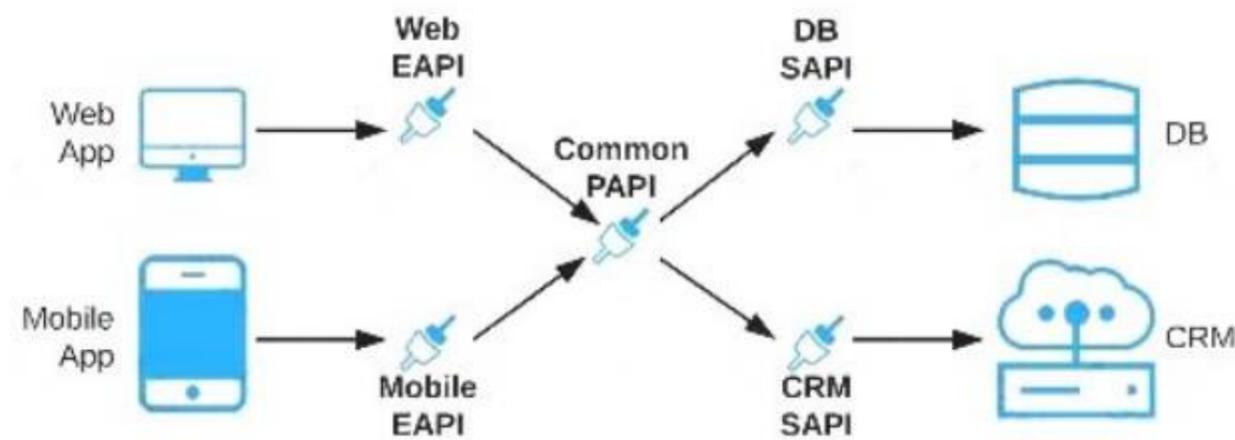
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements 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

A screenshot of a computer Description automatically generated with low confidence



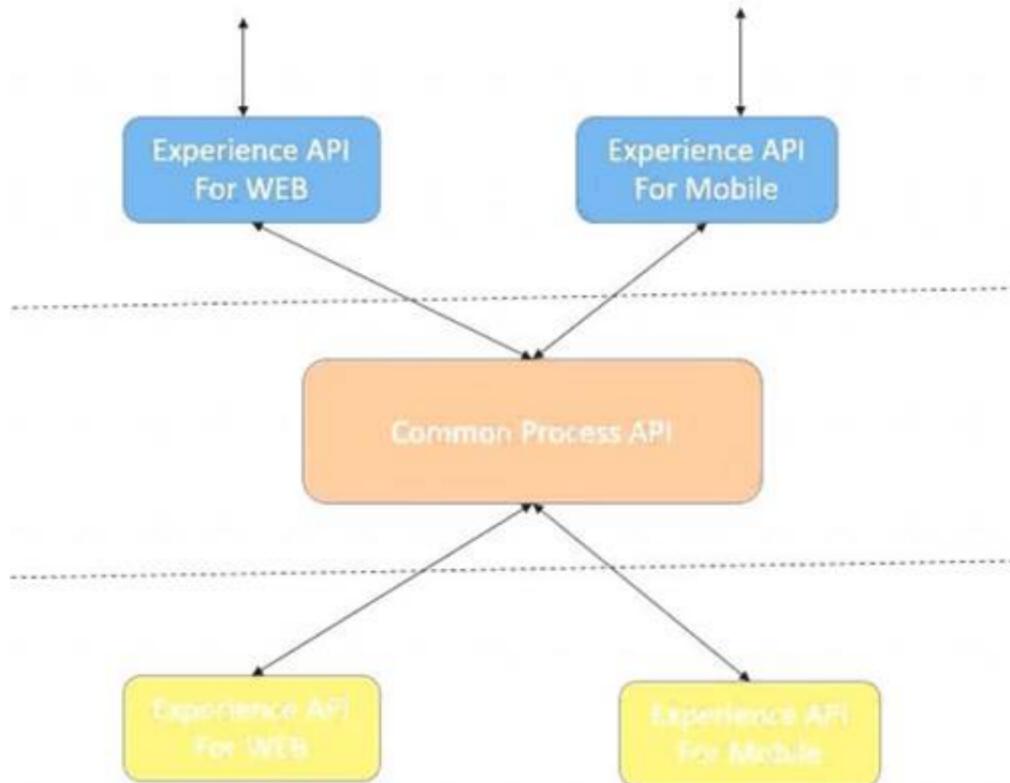
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements 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

Diagram Description automatically generated

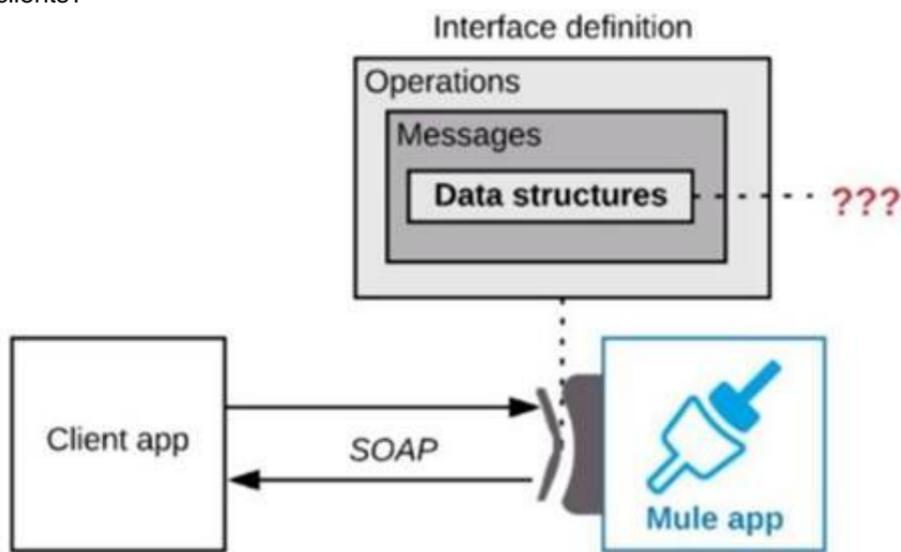


**NEW QUESTION 19**

Refer to the exhibit.

A Mule application is being designed to expose a SOAP web service to its clients.

What language is typically used inside the web service's interface definition to define the data structures that the web service is expected to exchange with its clients?



- A. WSDL
- B. XSD
- C. JSON Schema
- D. RAML

**Answer: B**

**Explanation:**

Correct Answer XSD In this approach to developing a web service, you begin with an XML schema (XSD file) that defines XML data structures to be used as parameters and return types in the web service operations.

Reference:

[https://www.w3schools.com/xml/schema\\_intro.asp](https://www.w3schools.com/xml/schema_intro.asp)

**NEW QUESTION 21**

What best describes the Fully Qualified Domain Names (FQDNs), also known as DNS entries, created when a Mule application is deployed to the CloudHub Shared Worker Cloud?

- A. A fixed number of FQDNs are created, IRRESPECTIVE of the environment and VPC design
- B. The FQDNs are determined by the application name chosen, IRRESPECTIVE of the region
- C. The FQDNs are determined by the application name, but can be modified by an administrator after deployment
- D. The FQDNs are determined by both the application name and the region

**Answer: D**

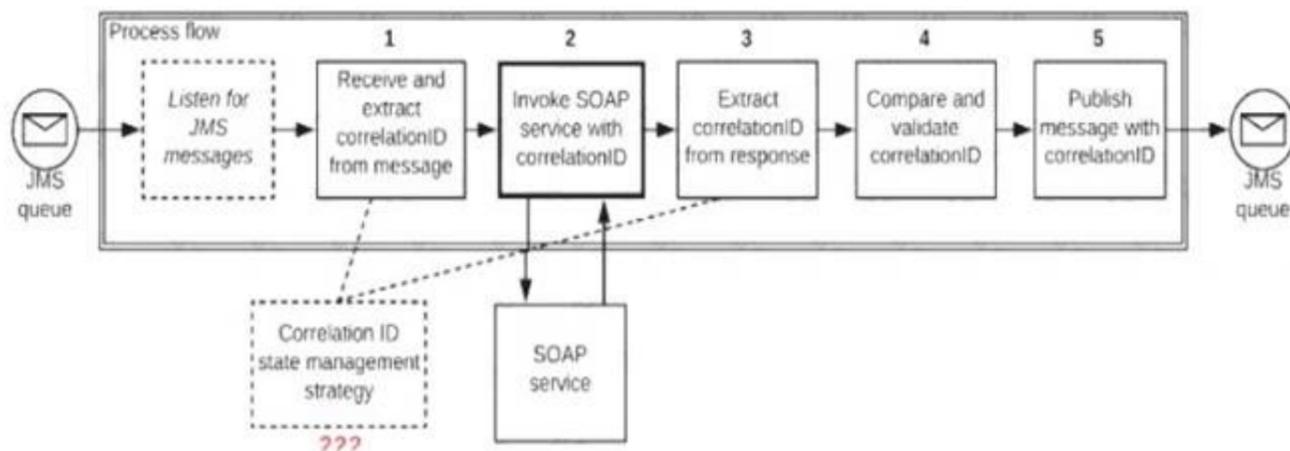
**Explanation:**

Every Mule application deployed to CloudHub receives a DNS entry pointing to the CloudHub. The DNS entry is a CNAME for the CloudHub Shared Load Balancer in the region to which the Mule application is deployed. When we deploy the application on CloudHub, we get a generic url to access the endpoints. Generic URL looks as below:  
 <application-name>.<region>.cloudhub.io <application-name> is the deployed application name which is unique across all the MuleSoft clients. <region> is the region name in which an application is deployed.

The public CloudHub (shared) load balancer already redirects these requests, where myApp is the name of the Mule application deployment to CloudHub: HTTP requests to `http://myApp.<region>.cloudhub.io` redirects to `http://mule-worker-myApp.<region>.cloudhub.io:8081`  
 HTTPS traffic to `https://myApp.<region>.cloudhub.io` redirects to `https://mule-worker-myApp.<region>.cloudhub.io:8082`

**NEW QUESTION 22**

Refer to the exhibit.



A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to receive JMS messages from a JMS queue. To process each received JMS message, the following steps are performed in a flow:

- Step 1: The JMS Correlation ID header is read from the received JMS message.
- Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.
- Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.
- Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.
- Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.

Where should the Mule application store the JMS Correlation ID values received in Step 1 and Step 3 so that the validation in Step 4 can be performed, while also making the overall Mule application highly available, fault-tolerant, performant, and maintainable?

- A. Both Correlation ID values should be stored in a persistent object store
- B. Both Correlation ID values should be stored in a non-persistent object store
- C. The Correlation ID value in Step 1 should be stored in a persistent object storeThe Correlation ID value in step 3 should be stored as a Mule event variable/attribute
- D. Both Correlation ID values should be stored as Mule event variable/attribute

**Answer: C**

**Explanation:**

- \* If we store Correlation id value in step 1 as Mule event variables/attributes, the values will be cleared after server restart and we want system to be fault tolerant.
- \* The Correlation ID value in Step 1 should be stored in a persistent object store.
- \* We don't need to store Correlation ID value in Step 3 to persistent object store. We can store it but as we also need to make application performant. We can avoid this step of accessing persistent object store.
- \* Accessing persistent object stores slow down the performance as persistent object stores are by default stored in shared file systems.
- \* As the SOAP service is idempotent in nature. In case of any failures , using this Correlation ID saved in first step we can make call to SOAP service and validate the Correlation ID.

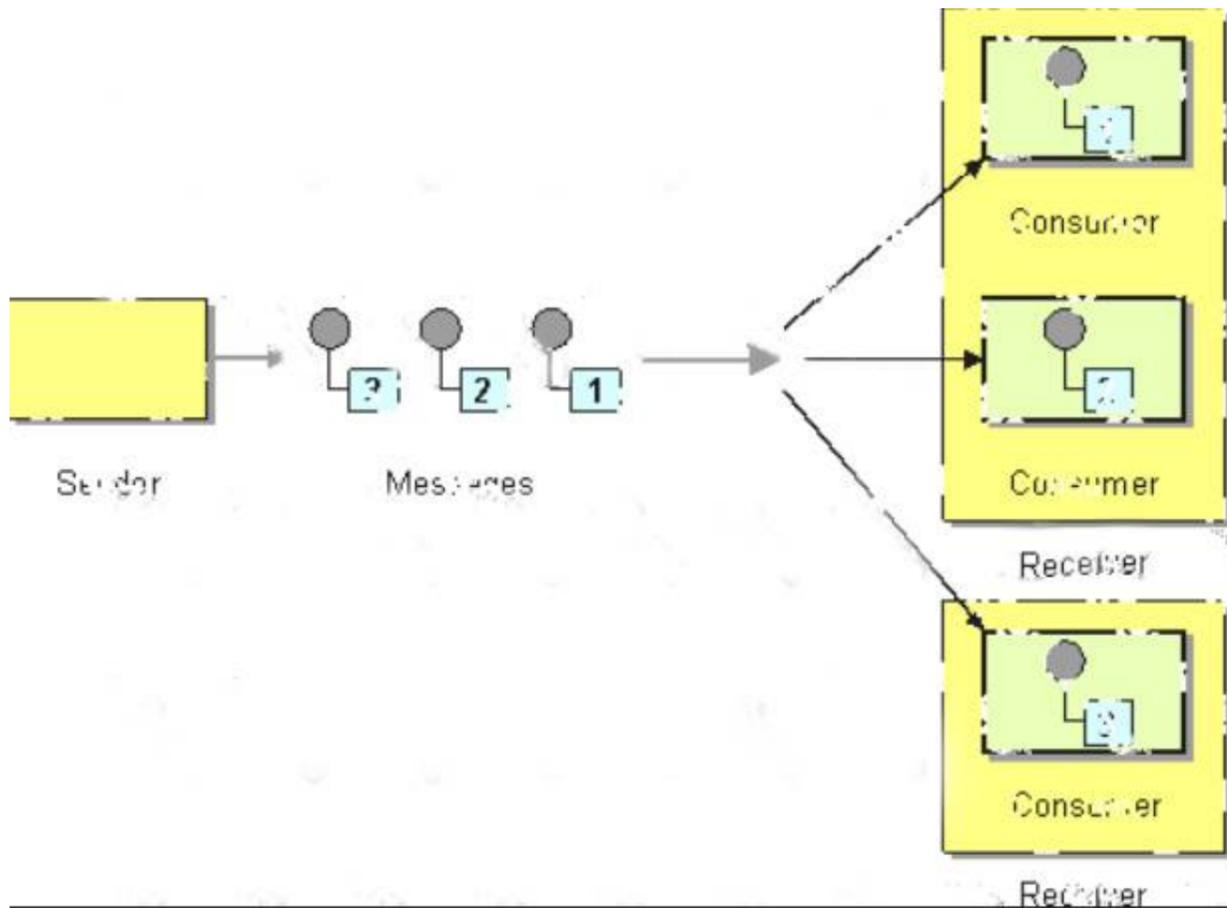
Top of Form

Additional Information:

- \* Competing Consumers

are multiple consumers that are all created to receive messages from a single Point-to-Point Channel. When the channel delivers a message, any of the consumers could potentially receive it. The messaging system's implementation determines which consumer actually receives the message, but in effect the consumers compete with each other to be the receiver. Once a consumer receives a message, it can delegate to the rest of its application to help process the message.

Diagram Description automatically generated



\* In case you are unaware about term idempotent re is more info:  
 Idempotent operations means their result will always same no matter how many times these operations are invoked.  
 Table Description automatically generated

# IDEMPOTENCE

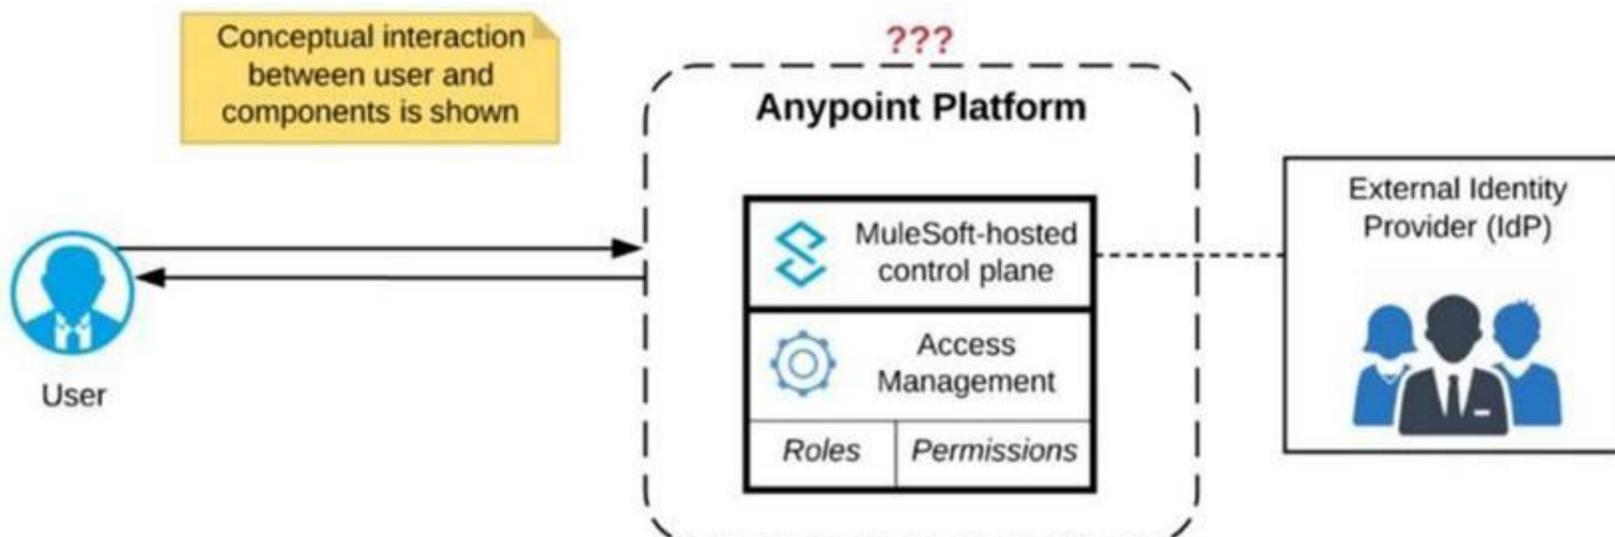
WHEN PERFORMING AN OPERATION AGAIN GIVES THE SAME RESULT

HTTP METHOD	IDEMPOTENCE	SAFETY
GET	YES	YES
HEAD	YES	YES
PUT	YES	NO
DELETE	YES	NO
POST	NO	NO
PATCH	NO	NO

Bottom of Form

**NEW QUESTION 26**

Refer to the exhibit.



Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity

management with Anypoint Platform.

What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled via the external Identity Provider?

- A. Controlling the business group within Anypoint Platform to which the user belongs
- B. Assigning Anypoint Platform permissions to a role
- C. Assigning Anypoint Platform role(s) to a user
- D. Removing a user's access to Anypoint Platform when they no longer work for the organization

**Answer: B**

**Explanation:**

\* By default, Anypoint Platform performs its own user management

– For user management, one external IdP can be integrated with the Anypoint Platform organization (note: not at business group level)

– Permissions and access control are still enforced inside Anypoint Platform and CANNOT be controlled via the external Identity Provider \* As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). \* You can map users in a federated organization's group to a role which also gives the flexibility of controlling the business group within Anypoint Platform to which the user belongs to.

Also user can nbe removed from external identity

management system when they no longer work for the organization. So they wont be able to authenticate using SSO to login to Anypoint Platform. \* Using external identity we can no change permissions of a particular role in Mulesoft Anypoint platform.

\* So Correct answer is Assigning Anypoint Platform permissions to a role

**NEW QUESTION 31**

When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

- A. Environment setup
- B. Identity management setup
- C. Role and permission setup
- D. Dedicated Load Balancer setup

**Answer: B**

**Explanation:**

\* Roles are business group specific. Configure identity management in the Anypoint Platform master organization. As the Anypoint Platform organization

administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). \* Roles and permissions can be set up at business group and organization level also. But Identity Management setup is only done at Organization level \* Business groups are self-contained resource groups that contain Anypoint Platform resources such as applications and APIs. Business groups provide a way to separate and control access to Anypoint Platform resources because users have access only to the busine

**NEW QUESTION 36**

An organization is struggling frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects.

How can these best practices be achieved with the applications having the least amount of responsibility?

- A. Create a Mule plugin project with all the dependencies and add it as a dependency in each application's POM.xml file
- B. Create a mule domain project with all the dependencies define in its POM.xml file and add each application to the domain Project
- C. Add all dependencies in each application's POM.xml file
- D. Create a parent POM of all the required dependencies and reference each in each application's POM.xml file

**Answer: D**

**NEW QUESTION 38**

An organization has deployed runtime fabric on an eight node 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

**Answer: D**

**NEW QUESTION 40**

An organization has decided on a cloud migration strategy to minimize the organization's own IT resources. Currently the organization has all of its new applications running on its own premises and uses an on-premises load balancer that exposes all APIs under the base URL (<https://api.rutujar.com>).

As part of migration strategy, the organization is planning to migrate all of its new applications and load balancer CloudHub.

What is the most straightforward and cost-effective approach to Mule application deployment and load balancing that preserves the public URL's?

- A. Deploy the Mule application to CloudhubCreate a CNAME record for base URL( <https://api.rutujar.com>) in the Cloudhub shared load balancer that points to the A record of the on-premises load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications
- B. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL ( <https://api.rutujar.com>) in the organization's DNS server to point to the A record of the Cloudhub dedicated load balancerApply mapping rules in DLB to map URLto their corresponding Mule applications
- C. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL ( <https://api.rutujar.com>) in the organization's DNS server to point to the A record of the CloudHub shared load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications
- D. For each migrated Mule application, deploy an API proxy application to Cloudhub with all traffic to themule applications routed through a Cloud Hub Dedicated load balancer (DLB)Update a CNAME record for base URL ( <https://api.rutujar.com>) in the organization's DNS server to point to the A record of the CloudHub dedicated load balancerApply mapping rules in DLB to map each API proxy application who is responding new application

**Answer: B**

#### NEW QUESTION 43

An organization will deploy Mule applications to Cloudhub, Business requirements mandate that all application logs be stored ONLY in an external splunk consolidated logging service and NOT in Cloudhub.

In order to most easily store Mule application logs ONLY in Splunk, how must Mule application logging be configured in Runtime Manager, and where should the log4j2 splunk appender be defined?

- A. Keep the default logging configuration in RuntimeManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support at Mule application deployments.
- B. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in EACH Mule application's log4j2.xml file
- C. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manger tosupport at Mule application deployments.
- D. Keep the default logging configuration in Runtime ManagerDefine the Splunk appender in EACH Mule application log4j2.xml file

**Answer: B**

#### Explanation:

By default, CloudHub replaces a Mule application's log4j2.xml file with a CloudHub log4j2.xml file. In CloudHub, you can disable the CloudHub provided Mule application log4j2 file. This allows integrating Mule application logs with custom or third-party log management systems

#### NEW QUESTION 48

Which Mulesoft feature helps users to delegate their access without sharing sensitive credentials or giving full control of accounts to 3rd parties?

- A. Secure Scheme
- B. client id enforcement policy
- C. Connected apps
- D. Certificates

**Answer: C**

#### Explanation:

Connected Apps

The Connected Apps feature provides a framework that enables an external application to integrate with Anypoint Platform using APIs through OAuth 2.0 and OpenID Connect. Connected apps help users delegate their access without sharing sensitive credentials or giving full control of their accounts to third parties. Actions taken by connected apps are audited, and users can also revoke access at any time. Note that some products do not currently include client IDs in this release of the Connected Apps feature. The Connected Apps feature enables you to use secure authentication protocols and control an app's access to user data. Additionally, end users can authorize the app to access their Anypoint Platform data.

Mule Ref Doc : <https://docs.mulesoft.com/access-management/connected-apps-overview>

#### NEW QUESTION 53

As a part of business requirement , old CRM system needs to be integrated using Mule application. CRM system is capable of exchanging data only via SOAP/HTTP protocol. As an integration architect who follows API led approach , what is the the below step you will perform so that you can share document with CRM team?

- A. Create RAML specification using Design Center
- B. Create SOAP API specification using Design Center
- C. Create WSDL specification using text editor
- D. Create WSDL specification using Design Center

**Answer: C**

#### Explanation:

Correct answer is Create WSDL specification using text editor SOAP services are specified using WSDL. A client program connecting to a web service can read the WSDL to determine what functions are available on the server. We can not create WSDL specification in Design Center. We need to use external text editor to create WSDL.

#### NEW QUESTION 54

An organization is creating a set of new services that are critical for their business. The project team prefers using REST for all services but is willing to use SOAP with common WS-" standards if a particular service requires it.

What requirement would drive the team to use SOAP/WS-\* for a particular service?

- A. Must use XML payloads for the service and ensure that it adheres to a specific schema
- B. Must publish and share the service specification (including data formats) with the consumers of the service
- C. Must support message acknowledgement and retry as part of the protocol
- D. Must secure the service, requiring all consumers to submit a valid SAML token

**Answer: D**

#### Explanation:

Security Assertion Markup Language (SAML) is an open standard that allows identity providers (IdP) to pass authorization credentials to service providers (SP). SAML transactions use Extensible Markup Language (XML) for standardized communications between the identity provider and service providers.

SAML is the link between the authentication of a user's identity and the authorization to use a service. WS-Security is the key extension that supports many authentication models including: basic username/password credentials, SAML, OAuth and more.

A common way that SOAP API's are authenticated is via SAML Single Sign On (SSO). SAML works by facilitating the exchange of authentication and authorization credentials across applications. However, there is no specification that describes how to add SAML to REST web services.

Reference: <https://www.oasis-open.org/committees/download.php/16768/wss-v1.1-spec-os-SAMLTOKENProfile.pdf>

#### NEW QUESTION 55

A corporation has deployed multiple mule applications implementing various public and private API's to different cloudhub workers. These API's are Critical applications that must be highly available and in line with the reliability SLA as defined by stakeholders.

How can API availability (liveliness or readiness) be monitored so that Ops team receives outage notifications?

- A. Enable monitoring of individual applications from Anypoint monitoring
- B. Configure alerts with failure conditions in runtime manager
- C. Configure alerts failure conditions in API manager
- D. Use any point functional monitoring test API's functional behavior

**Answer: A**

#### NEW QUESTION 59

The AnyAirline organization's passenger reservations center is designing an integration solution that combines invocations of three different System APIs (bookFlight, bookHotel, and bookCar) in a business transaction. Each System API makes calls to a single database.

The entire business transaction must be rolled back when at least one of the APIs fails.

What is the most idiomatic (used for its intended purpose) way to integrate these APIs in near real-time that provides the best balance of consistency, performance, and reliability?

- A. Implement eXtended Architecture (XA) transactions between the API implementations Coordinate between the API implementations using a Saga patternImplement caching in each API implementation to improve performance
- B. Implement local transactions within each API implementationConfigure each API implementation to also participate in the same eXtended Architecture (XA) transactionImplement caching in each API implementation to improve performance
- C. Implement local transactions in each API implementation Coordinate between the API implementations using a Saga patternApply various compensating actions depending on where a failure occurs
- D. Implement an eXtended Architecture (XA) transaction manager in a Mule application using a Saga patternConnect each API implementation with the Mule application using XA transactions Apply various compensating actions depending on where a failure occurs

**Answer: C**

#### NEW QUESTION 61

An external web UI application currently accepts occasional HTTP requests from client web browsers to change (insert, update, or delete) inventory pricing information in an inventory system's database. Each inventory pricing change must be transformed and then synchronized with multiple customer experience systems in near real-time (in under 10 seconds). New customer experience systems are expected to be added in the future.

The database is used heavily and limits the number of SELECT queries that can be made to the database to 10 requests per hour per user.

What is the most scalable, idiomatic (used for its intended purpose), decoupled, reusable, and maintainable integration mechanism available to synchronize each inventory pricing change with the various customer experience systems in near real-time?

- A. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the watermark attribute set to an appropriate database columnIn the same flow, use a Scatter-Gather to call each customer experience system's REST API with transformed inventory-pricing records
- B. Add a trigger to the inventory-pricing database table so that for each change to the inventory pricing database, a stored procedure is called that makes a REST call to a Mule applicationWrite the Mule application to publish each Mule event as a message to an Anypoint MQ exchange Write other Mule applications to subscribe to the Anypoint MQ exchange, transform each received message, and then update the Mule application's corresponding customer experience system(s)
- C. Replace the external web UI application with a Mule application to accept HTTP requests from client web browsersIn the same Mule application, use a Batch Job scope to test if the database request will succeed, aggregate pricing changes within a short time window, and then update both the inventory pricing database and each customer experience system using a Parallel For Each scope
- D. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the ID attribute set to an appropriate database columnIn the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queueWrite other Mule applications to subscribe to the Anypoint MQ queue, transform each received message, and then update the Mule application's corresponding customer experience system(s)

**Answer: B**

#### NEW QUESTION 66

An organization has just developed a Mule application that implements a REST API. The mule application will be deployed to a cluster of customer hosted Mule runtimes.

What additional infrastructure component must the customer provide in order to distribute inbound API requests across the Mule runtimes of the cluster?

- A. A message broker
- B. An HTTP Load Balancer
- C. A database
- D. An Object Store

**Answer: B**

#### Explanation:

Correct answer is An HTTP Load Balancer.

Key thing to note here is that we are deploying application to customer hosted Mule runtime. This means we will need load balancer to route the requests to different instances of the cluster.

#### NEW QUESTION 71

A travel company wants to publish a well-defined booking service API to be shared with its business partners. These business partners have agreed to ONLY consume SOAP services and they want to get the service contracts in an easily consumable way before they start any development. The travel company will publish the initial design documents to Anypoint Exchange, then share those documents with the business partners. When using an API-led approach, what is the first design document the travel company should deliver to its business partners?

- A. Create a WSDL specification using any XML editor
- B. Create a RAML API specification using any text editor
- C. Create an OAS API specification in Design Center
- D. Create a SOAP API specification in Design Center

**Answer:** A

**Explanation:**

SOAP API specifications are provided as WSDL. Design center doesn't provide the functionality to create WSDL file. Hence WSDL needs to be created using XML editor

**NEW QUESTION 72**

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 75**

An insurance provider is implementing Anypoint platform to manage its application infrastructure and is using the customer hosted runtime for its business due to certain financial requirements it must meet. It has built a number of synchronous API's and is currently hosting these on a mule runtime on one server. These applications make use of a number of components including heavy use of object stores and VM queues. Business has grown rapidly in the last year and the insurance provider is starting to receive reports of reliability issues from its applications.

The DevOps team indicates that the API's are currently handling too many requests and this is over loading the server. The team has also mentioned that there is a significant downtime when the server is down for maintenance.

As an integration architect, which option would you suggest to mitigate these issues?

- A. Add a load balancer and add additional servers in a server group configuration
- B. Add a load balancer and add additional servers in a cluster configuration
- C. Increase physical specifications of server CPU memory and network
- D. Change applications by use an event-driven model

**Answer:** B

**NEW QUESTION 78**

An organization is designing a mule application to support an all or nothing transaction between several 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

**Answer:** A

**Explanation:**

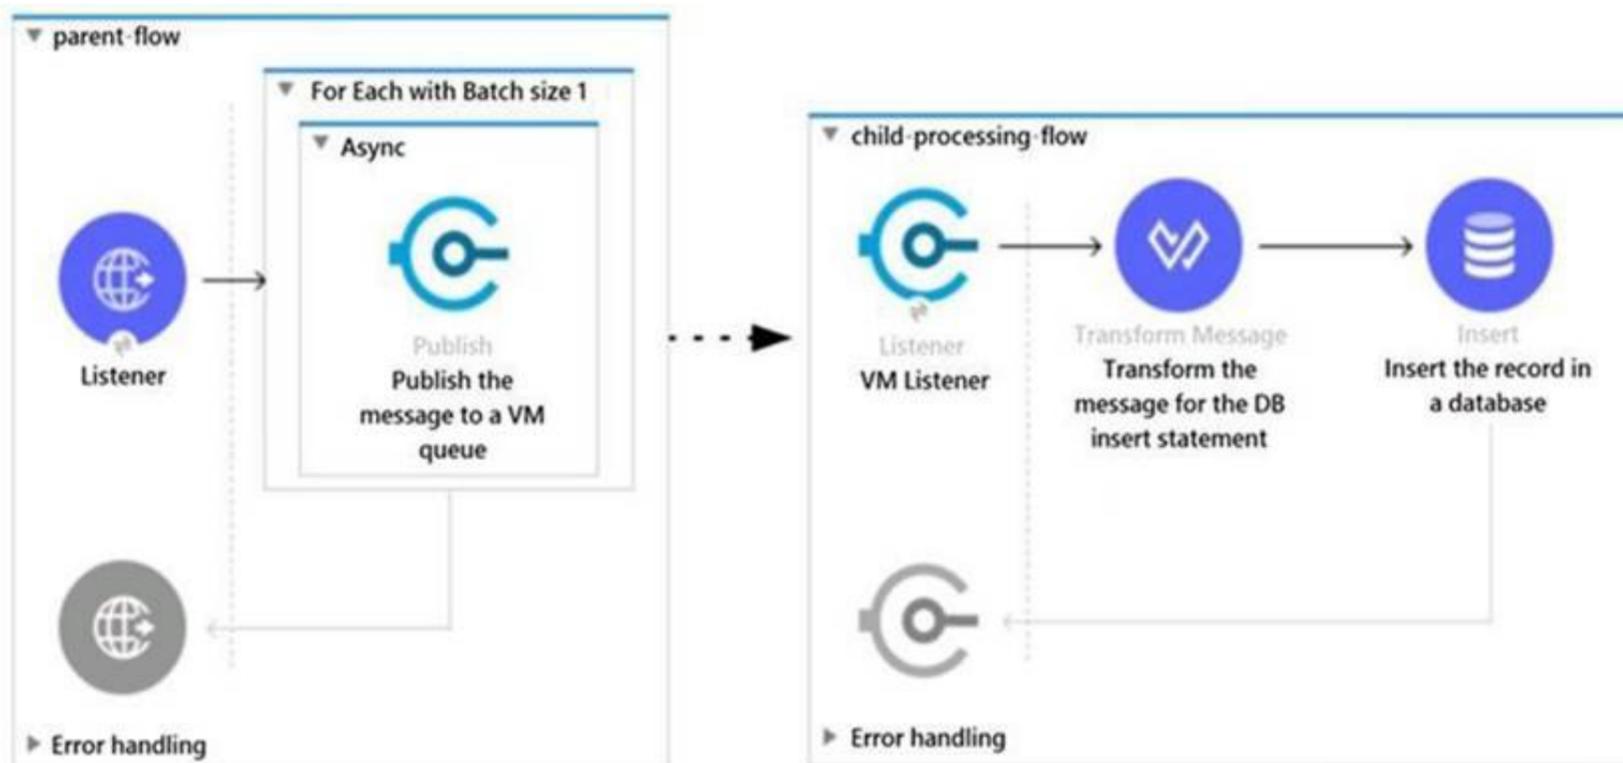
Correct answer is VM. VM support Transactional Type. When an exception occurs, the transaction rolls back to its original state for reprocessing. This feature is not supported by other connectors.

Here is additional information about Transaction management: Table Description automatically generated

	Shared Load Balancer	Dedicated Load Balancer
VPC	Shared VPC (Mulesoft)	VPC (Customer)
Default Load Balancer	Cloudhub provides Default Shared Load Balancer available in All Environment	Need to Purchase
Organization Use	Multiple Organization	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)
Algorithm	Round Robin	Round Robin
Supports HTTPS Protocol	Yes	Yes
Worker Assignment	No	Yes
IP Blacklisting/Whitelisting	No <small><a href="https://docs.mulesoft.com/runtime-manager/ip-whitelists">https://docs.mulesoft.com/runtime-manager/ip-whitelists</a></small>	Yes
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

**NEW QUESTION 82**

Refer to the exhibit.



A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items, then sends each item one at a time inside an Async scope to a VM queue.

A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing the item to a database. This Mule application is deployed to four CloudHub workers with persistent queues enabled.

What message processing guarantees are provided by the VM queue and the CloudHub workers, and how are VM messages routed among the CloudHub workers for each invocation of the parent flow under normal operating conditions where all the CloudHub workers remain online?

- A. EACH item VM message is processed AT MOST ONCE by ONE CloudHub worker, with workers chosen in a deterministic round-robin fashion Each of the four CloudHub workers can be expected to process 1/4 of the Item VM messages (about 50 items)
- B. EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker Each of the four CloudHub workers can be expected to process some item VM messages

- C. ALL Item VM messages are processed AT LEAST ONCE by the SAME CloudHub worker where the parent flow was invoked This one CloudHub worker processes ALL 200 item VM messages
- D. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub worker This one CloudHub worker processes ALL 200 item VM messages

**Answer: B**

**Explanation:**

Correct answer is EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker. Each of the four CloudHub workers can be expected to process some item VM messages In Cloudhub, each persistent VM queue is listened on by every CloudHub worker - But each message is read and processed at least once by only one CloudHub worker and the duplicate processing is possible - If the CloudHub worker fails , the message can be read by another worker to prevent loss of messages and this can lead to duplicate processing - By default , every CloudHub worker's VM Listener receives different messages from VM Queue Referenece: <https://dzone.com/articles/deploying-mulesoft-application-on-1-worker-vs-mult>

**NEW QUESTION 85**

What is required before an API implemented using the components of Anypoint Platform can be managed and governed (by applying API policies) on Anypoint Platform?

- A. The API must be published to Anypoint Exchange and a corresponding API instance ID must be obtained from API Manager to be used in the API implementation
- B. The API implementation source code must be committed to a source control management system (such as GitHub)
- C. A RAML definition of the API must be created in API designer so it can then be published to Anypoint Exchange
- D. The API must be shared with the potential developers through an API portal so API consumers can interact with the API

**Answer: A**

**Explanation:**

Context of the question is about managing and governing mule applications deployed on Anypoint platform.

Anypoint API Manager (API Manager) is a component of Anypoint Platform that enables you to manage, govern, and secure APIs. It leverages the runtime capabilities of API Gateway and Anypoint Service Mesh, both of which enforce policies, collect and track analytics data, manage proxies, provide encryption and authentication, and manage applications.

Mule Ref Doc : <https://docs.mulesoft.com/api-manager/2.x/getting-started-proxy>

**NEW QUESTION 89**

A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.

The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.

What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

- A. Retrieve the OpenAPI Specification file(s) from API Manager
- B. Retrieve the WSDL file(s) from the deployed Mule application
- C. Retrieve the RAML file(s) from the deployed Mule application
- D. Retrieve the XML file(s) from Runtime Manager

**Answer: D**

**NEW QUESTION 92**

An API implementation is being developed to expose data from a production database via HTTP requests. The API implementation executes a database SELECT statement that is dynamically created based upon data received from each incoming HTTP request. The developers are planning to use various types of testing to make sure the Mule application works as expected, can handle specific workloads, and behaves correctly from an API consumer perspective. What type of testing would typically mock the results from each SELECT statement rather than actually execute it in the production database?

- A. Unit testing (white box)
- B. Integration testing
- C. Functional testing (black box)
- D. Performance testing

**Answer: A**

**Explanation:**

In Unit testing instead of using actual backends, stubs are used for the backend services. This ensures that developers are not blocked and have no dependency on other systems.

In Unit testing instead of using actual backends, stubs are used for the backend services. This ensures that developers are not blocked and have no dependency on other systems.

Below are the typical characteristics of unit testing.

- Unit tests do not require deployment into any special environment, such as a staging environment
- Unit tests can be run from within an embedded Mule runtime
- Unit tests can/should be implemented using MUnit
- For read-only interactions to any dependencies (such as other APIs): allowed to invoke production endpoints
- For write interactions: developers must implement mocks using MUnit
- Require knowledge of the implementation details of the API implementation under test

**NEW QUESTION 96**

A new upstream API is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity. The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms. If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API's desired SLA?

- A. Set a timeout of 100 ms; that leaves 400 ms for the other two downstream APIs to complete
- B. Do not set a timeout; the Invocation of this API is mandatory and so we must wait until it responds
- C. Set a timeout of 50 ms; this times out more invocations of that API but gives additional room for retries
- D. No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

**Answer:** D

**Explanation:**

Before we answer this question, we need to understand what median (50th percentile) and 80th percentile means. If the 50th percentile (median) of a response time is 500ms that means that 50% of my transactions are either as fast or faster than 500ms.

If the 90th percentile of the same transaction is at 1000ms it means that 90% are as fast or faster and only 10% are slower. Now as per upstream SLA, 99th percentile is 800 ms which means 99% of the incoming requests should have response time less than or equal to 800 ms. But as per one of the backend API, their 95th percentile is 1000 ms which means that backend API will take 1000 ms or less than that for 95% of requests. As there are three API invocation from upstream API, we can not conclude a timeout that can be set to meet the desired SLA as backend SLA's do not support it.

Let see why other answers are not correct.

1) Do not set a timeout --> This can potentially violate SLA's of upstream API

2) Set a timeout of 100 ms; ---> This will not work as backend API has 100 ms as median meaning only 50% requests will be answered in this time and we will get timeout for 50% of the requests. Important thing to note here is, All APIs need to be executed sequentially, so if you get timeout in first API, there is no use of going to second and third API. As a service provider you wouldn't want to keep 50% of your consumers dissatisfied. So not the best option to go with.

\*To quote an example: Let's assume you have built an API to update customer contact details.

- First API is fetching customer number based on login credentials

- Second API is fetching Info in 1 table and returning unique key

- Third API, using unique key provided in second API as primary key, updating remaining details

\* Now consider, if API times out in first API and can't fetch customer number, in this case, it's useless to call API 2 and 3 and that is why question mentions specifically that all APIs need to be executed sequentially.

3) Set a timeout of 50 ms --> Again not possible due to the same reason as above Hence correct answer is No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

**NEW QUESTION 98**

.....

## Relate Links

**100% Pass Your MCIA-Level-1 Exam with Exam Bible Prep Materials**

<https://www.exambible.com/MCIA-Level-1-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>