



## **MuleSoft**

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

MuleSoft Certified Integration Architect - Level 1

### NEW QUESTION 1

Additional nodes are being added to an existing customer-hosted Mule runtime cluster to improve performance. Mule applications deployed to this cluster are invoked by API clients through a load balancer.

What is also required to carry out this change?

- A. A new load balancer must be provisioned to allow traffic to the new nodes in a round-robin fashion
- B. External monitoring tools or log aggregators must be configured to recognize the new nodes
- C. API implementations using an object store must be adjusted to recognize the new nodes and persist to them
- D. New firewall rules must be configured to accommodate communication between API clients and the new nodes

**Answer: B**

#### Explanation:

- \* Clustering is a group of servers or mule runtime which acts as a single unit.
  - \* Mulesoft Enterprise Edition supports scalable clustering to provide high availability for the Mulesoft application.
  - \* In simple terms, virtual servers composed of multiple nodes and they communicate and share information through a distributed shared memory grid.
  - \* By default, Mulesoft ensures the High availability of applications if clustering implemented.
  - \* Let's consider the scenario one of the nodes in cluster crashed or goes down and under maintenance. In such cases, Mulesoft will ensure that requests are processed by other nodes in the cluster. Mulesoft clustering also ensures that the request is load balanced between all the nodes in a cluster.
  - \* Clustering is only supported by on-premise Mule runtime and it is not supported in Cloudhub.
- Correct answer is External monitoring tools or log aggregators must be configured to recognize the new nodes
- \* Rest of the options are automatically taken care of when a new node is added in cluster.

### NEW QUESTION 2

In Anypoint Platform, a company wants to configure multiple identity providers (Idps) for various lines of business (LOBs) Multiple business groups and environments have been defined for these LOBs. What Anypoint Platform feature can use multiple Idps access the company's business groups and environment?

- A. User management
- B. Roles and permissions
- C. Dedicated load balancers
- D. Client Management

**Answer: D**

#### Explanation:

Correct answer is Client Management

- \* Anypoint Platform acts as a client provider by default, but you can also configure external client providers to authorize client applications.
- \* As an API owner, you can apply an OAuth 2.0 policy to authorize client applications that try to access your API. You need an OAuth 2.0 provider to use an OAuth 2.0 policy.
- \* You can configure more than one client provider and associate the client providers with different environments. If you configure multiple client providers after you have already created environments, you can associate the new client providers with the environment.
- \* You should review the existing client configuration before reassigning client providers to avoid any downtime with existing assets or APIs.
- \* When you delete a client provider from your master organization, the client provider is no longer available in environments that used it.
- \* Also, assets or APIs that used the client provider can no longer authorize users who want to access them.

-----MuleSoft

Reference: <https://docs.mulesoft.com/access-management/managing-api-clients>  
<https://www.folkstalk.com/2019/11/mulesoft-integration-and-platform.html>

### NEW QUESTION 3

An organization has implemented the cluster with two customer hosted Mule runtimes is hosting an application.

This application has a flow with a JMS listener configured to consume messages from a queue destination. As an integration architect can you advise which JMS listener configuration must be used to receive messages in all the nodes of the cluster?

- A. Use the parameter primaryNodeOnly= "false" on the JMS listener
- B. Use the parameter primaryNodeOnly= "false" on the JMS listener with a shared subscription
- C. Use the parameter primaryNodeOnly= "true" on the JMS listener with a non-shared subscription
- D. Use the parameter primaryNodeOnly= "true" on the JMS listener

**Answer: A**

### NEW QUESTION 4

A company is using Mulesoft to develop API's and deploy them to Cloudhub and on premises targets. Recently it has decided to enable Runtime Fabric deployment option as well and infrastructure is set up for this option.

What can be used to deploy Runtime Fabric?

- A. AnypointCLI
- B. Anypoint platform REST API's
- C. Directly uploading ajar file from the Runtime manager
- D. Mule maven plug-in

**Answer: D**

### NEW QUESTION 5

A mule application is being designed to perform product orchestration. The Mule application needs to join together the responses from an inventory API and a Product Sales History API with the least latency.

To minimize the overall latency. What is the most idiomatic (used for its intended purpose) design to call each API request in the Mule application?

- A. Call each API request in a separate lookup call from Dataweave reduce operator
- B. Call each API request in a separate route of a Scatter-Gather
- C. Call each API request in a separate route of a Parallel For Each scope
- D. Call each API request in a separate Async scope

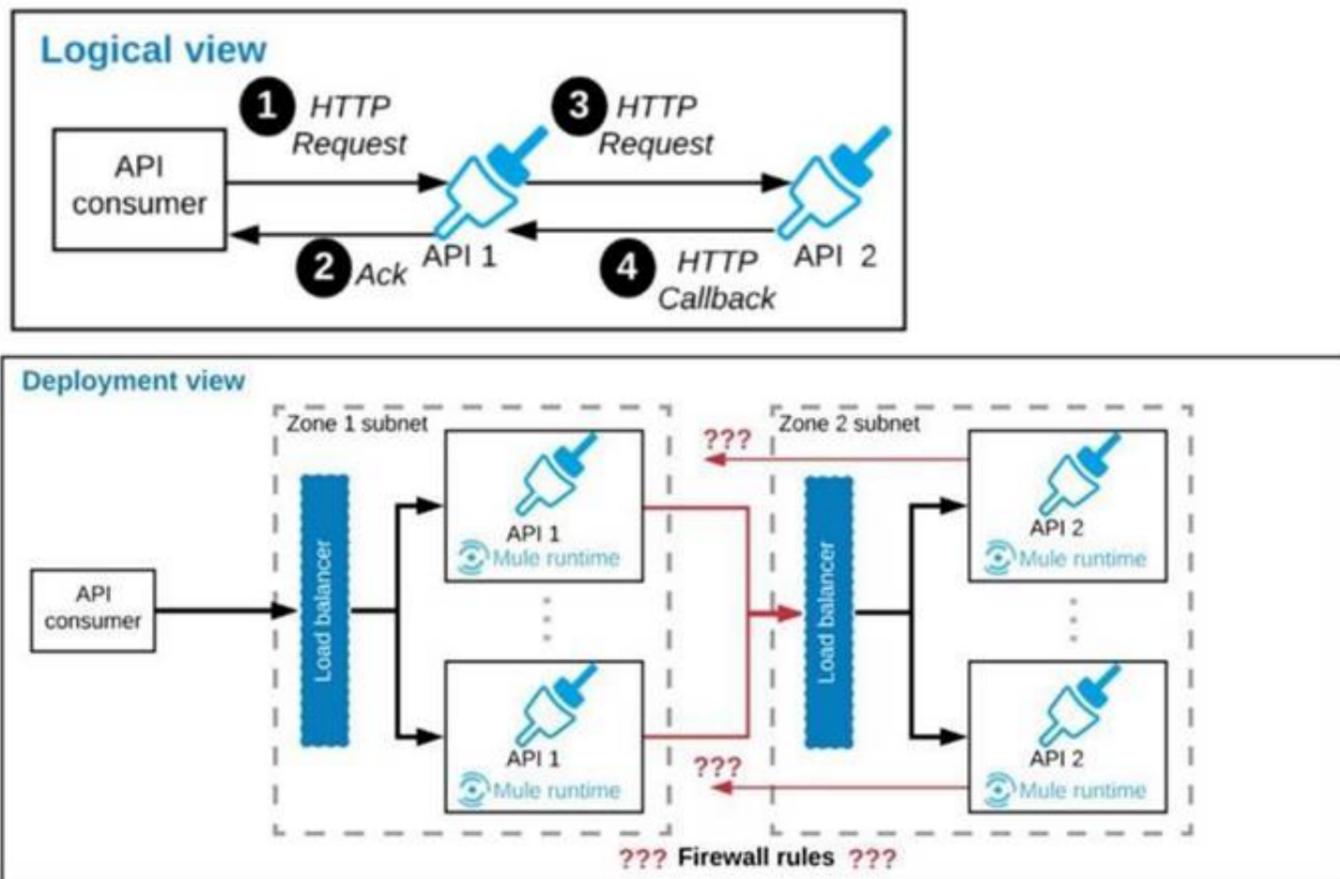
Answer: B

**Explanation:**

Scatter-Gather sends a request message to multiple targets concurrently. It collects the responses from all routes, and aggregates them into a single message.

**NEW QUESTION 6**

Refer to the exhibit.



A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API 1 to notify about completion of the asynchronous process.

Each API is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.

In the network architecture, how must the firewall rules be configured to enable the above Interaction between API 1 and API 2?

- A. To authorize the certificate to be used both APIs
- B. To enable communication from each API's Mule Runtimes and Network zone to the load balancer of the other API
- C. To open direct two-way communication between the Mule Runtimes of both API's
- D. To allow communication between load balancers used by each API

Answer: B

**Explanation:**

\* If your API implementation involves putting a load balancer in front of your APIkit application, configure the load balancer to redirect URLs that reference the baseUrl of the application directly. If the load balancer does not redirect URLs, any calls that reach the load balancer looking for the application do not reach their destination.

\* When you receive incoming traffic through the load balancer, the responses will go out the same way. However, traffic that is originating from your instance will not pass through the load balancer. Instead, it is sent directly from the public IP address of your instance out to the Internet. The ELB is not involved in that scenario.

\* The question says "each API is deployed to multiple redundant Mule runtimes", that seems to be a hint for self hosted Mule runtime cluster. Set Inbound allowed for the LB, outbound allowed for runtime to request out.

\* Hence correct way is to enable communication from each API's Mule Runtimes and Network zone to the load balancer of the other API. Because communication is asynchronous one

**NEW QUESTION 7**

What operation can be performed through a JMX agent enabled in a Mule application?

- A. View object store entries
- B. Replay an unsuccessful message
- C. Set a particular tog4J2 log level to TRACE
- D. Deploy a Mule application

Answer: C

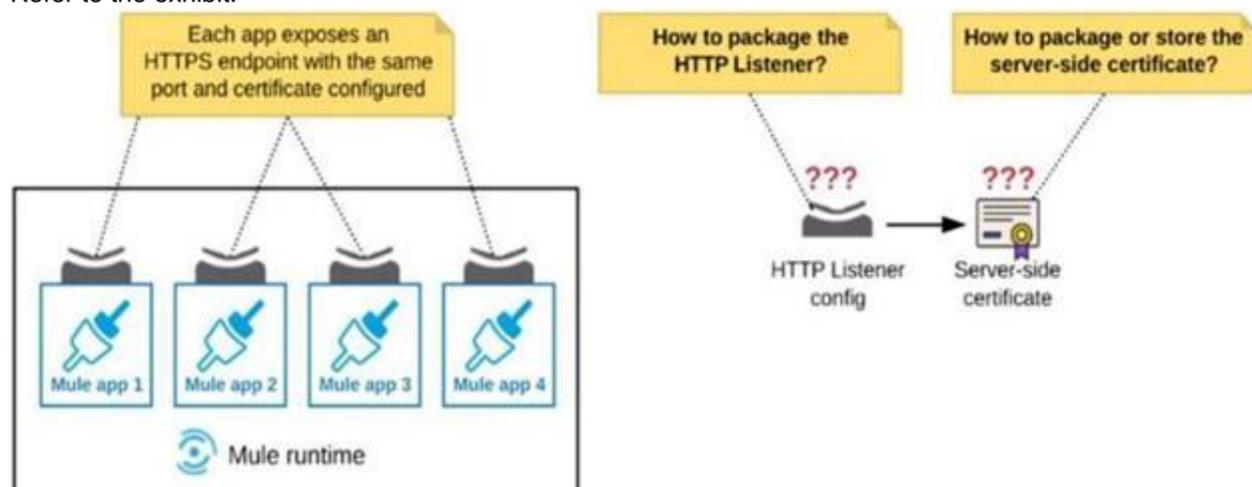
**Explanation:**

JMX Management Java Management Extensions (JMX) is a simple and standard way to manage applications, devices, services, and other resources. JMX is dynamic, so you can use it to monitor and manage resources as they are created, installed, and implemented. You can also use JMX to monitor and manage the Java Virtual Machine (JVM). Each resource is instrumented by one or more Managed Beans, or MBeans. All MBeans are registered in an MBean Server. The JMX server agent consists of an MBean Server and a set of services for handling Mbeans. There are several agents provided with Mule for JMX support. The easiest way to configure JMX is to use the default JMX support agent. Log4J Agent The log4j agent exposes the configuration of the Log4J instance used by Mule for JMX

management. You enable the Log4J agent using the <jmx-log4j> element. It does not take any additional properties MuleSoft Reference: <https://docs.mulesoft.com/mule-runtime/3.9/jmx-management>

**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 repackaging 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:

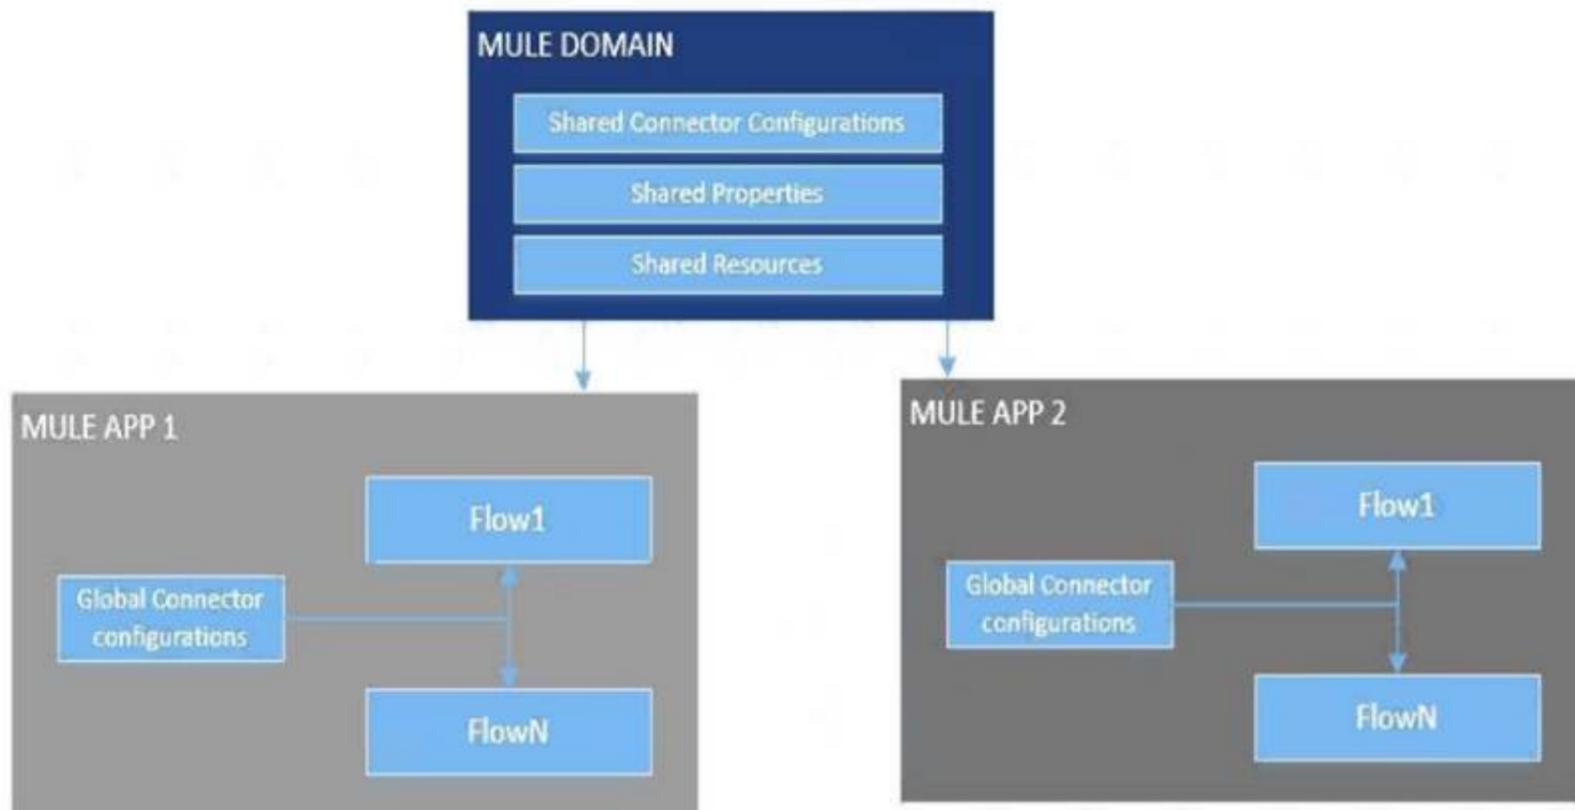
- Expose multiple services within the domain through the same port.
- Share the connection to persistent storage.
- Share services between apps through a well-defined interface.
- Ensure consistency between apps upon any changes because the configuration is only set in one place.

\* 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**

Insurance organization is planning to deploy Mule application in MuleSoft Hosted runtime plane. As a part of requirement , application should be scalable . highly available. It also has regulatory requirement which demands logs to be retained for at least 2 years. As an Integration Architect what step you will recommend in order to achieve this?

- A. It is not possible to store logs for 2 years in CloudHub deployment
- B. External log management system is required.
- C. When deploying an application to CloudHub , logs retention period should be selected as 2 years
- D. When deploying an application to CloudHub, worker size should be sufficient to store 2 years data
- E. Logging strategy should be configured accordingly in log4j file deployed with the application.

**Answer: A**

**Explanation:**

Correct answer is It is not possible to store logs for 2 years in CloudHub deployment. External log management system is required. CloudHub has a specific log retention policy, as described in the documentation: the platform stores logs of up to 100 MB per app & per worker or for up to 30 days, whichever limit is hit first. Once this limit has been reached, the oldest log information is deleted in chunks and is irretrievably lost. The recommended approach is to persist your logs to an external logging system of your choice (such as Splunk, for instance) using a log appender. Please note that this solution results in the logs no longer being stored on our platform, so any support cases you lodge will require for you to provide the appropriate logs for review and case resolution

**NEW QUESTION 10**

What is maximum vCores can be allocated to application deployed to CloudHub?

- A. 1 vCores
- B. 2 vCores
- C. 4 vCores
- D. 16 vCores

**Answer: D**

**NEW QUESTION 10**

An organization is designing multiple new applications to run on CloudHub in a single Anypoint VPC and that must share data using a common persistent Anypoint object store V2 (OSv2).

Which design gives these mule applications access to the same object store instance?

- A. AVM connector configured to directly access the persistence queue of the persistent object store
- B. An Anypoint MQ connector configured to directly access the persistent object store
- C. Object store V2 can be shared across cloudhub applications with the configured osv2 connector
- D. The object store V2 rest API configured to access the persistent object store

**Answer: D**

**NEW QUESTION 13**

Which Salesforce API is invoked to deploy, retrieve, create or delete customization information such as custom object definitions using a Mule Salesforce connector in a Mule application?

- A. Metadata API
- B. REST API

- C. SOAP API
- D. Bulk API

Answer: B

**NEW QUESTION 16**

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

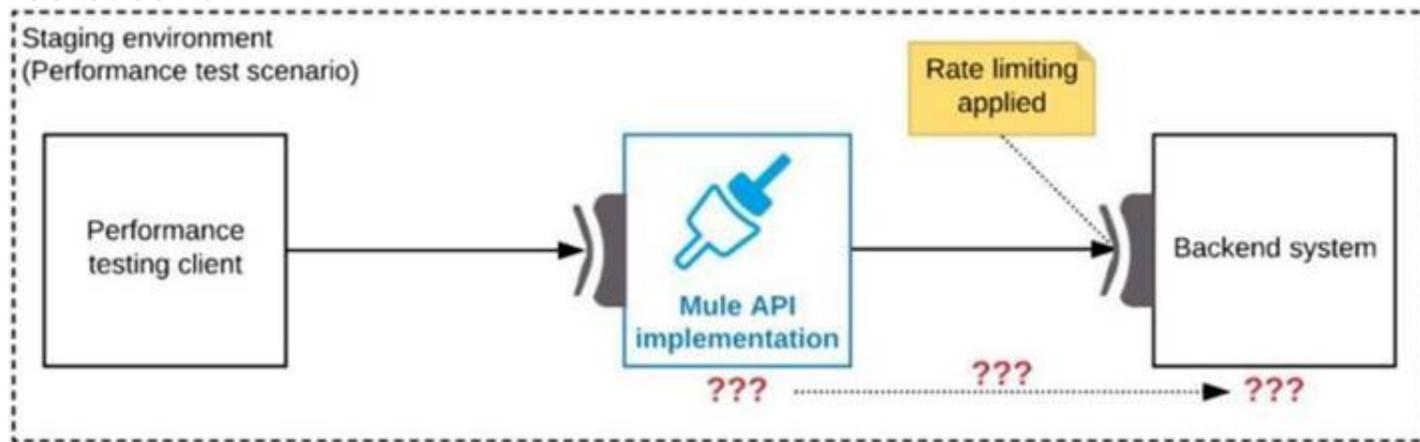
Answer: C

**Explanation:**

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

**NEW QUESTION 18**

Refer to the exhibit.



One of the backend systems invoked by an API implementation enforces rate limits on the number of requests a particular client can make. Both the backend system and the API implementation are deployed to several non-production environments in addition to production. Rate limiting of the backend system applies to all non-production environments. The production environment, however, does NOT have any rate limiting. What is the most effective approach to conduct performance tests of the API implementation in a staging (non-production) environment?

- A. Create a mocking service that replicates the backend system's production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests
- B. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
- C. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situatio
- D. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
- E. Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale

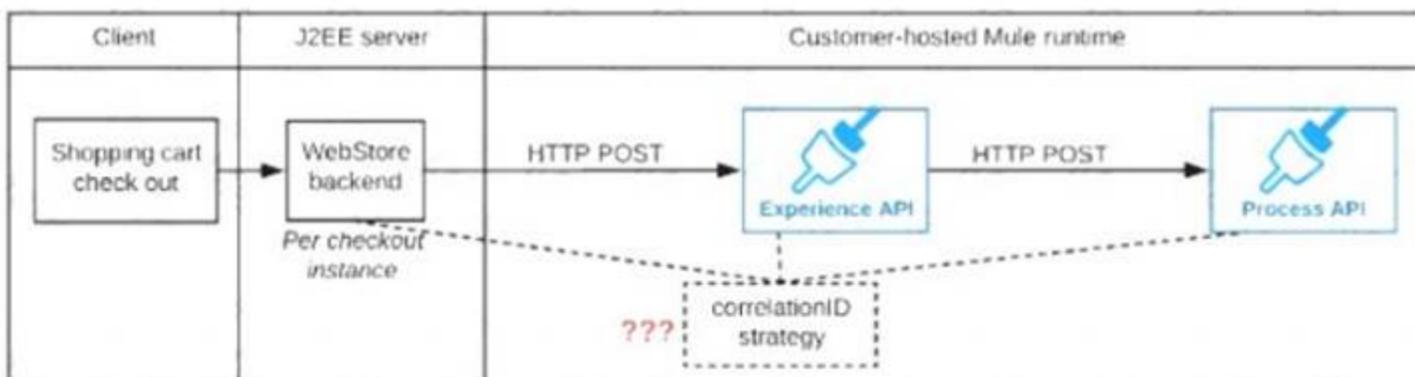
Answer: A

**Explanation:**

Correct answer is Create a mocking service that replicates the backend system's production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests  
 \* MUnit is for only Unit and integration testing for APIs and Mule apps. Not for performance Testing, even if it has the ability to Mock the backend.  
 \* Bypassing the backend invocation defeats the whole purpose of performance testing. Hence it is not a valid answer.  
 \* Scaled down performance tests cant be relied upon as performance of API's is not linear against load.

**NEW QUESTION 19**

Refer to the exhibit.



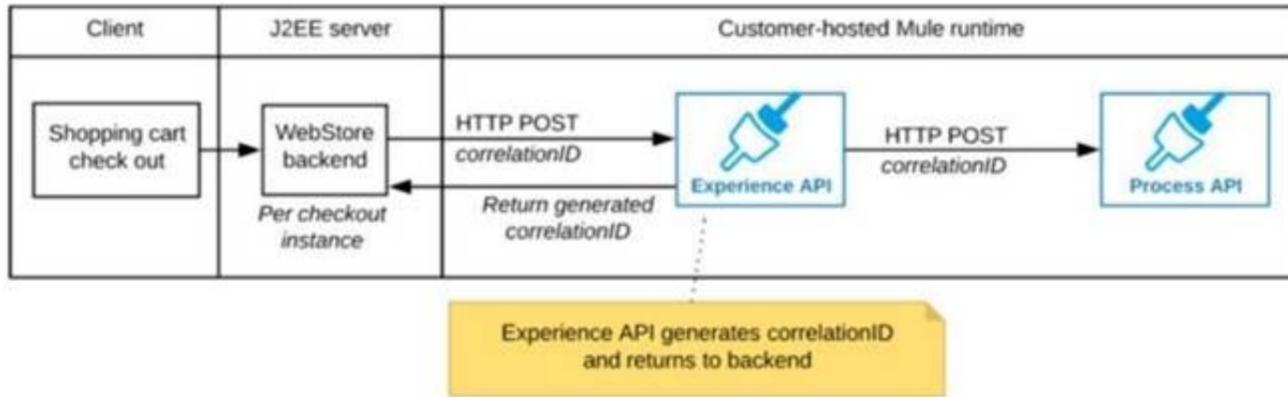
A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations are Mule applications executing in a customer -hosted Mule runtime. End-to-end correlation of all HTTP requests and responses belonging to each individual checkout Instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the same correlation ID for all requests and responses belonging to the same checkout instance.

What is the most efficient way (using the least amount of custom coding or configuration) for the web store backend and the implementations of the Experience API and Process API to participate in end-to-end correlation of the API invocations for each checkout instance?

A)

The web store backend, being a Java EE application, automatically makes use of the thread-local correlation ID generated by the Java EE application server and automatically transmits that to the Experience API using HTTP-standard headers

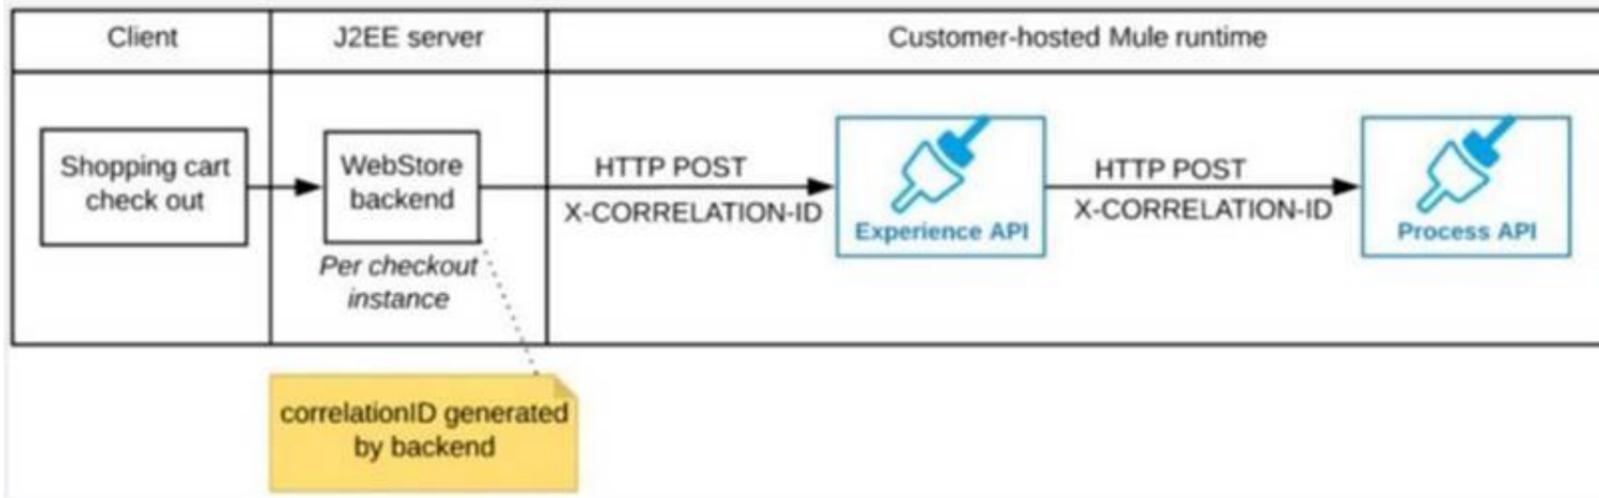
No special code or configuration is included in the web store backend, Experience API, and Process API implementations to generate and manage the correlation ID



B)

The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header In each API invocation belonging to that checkout

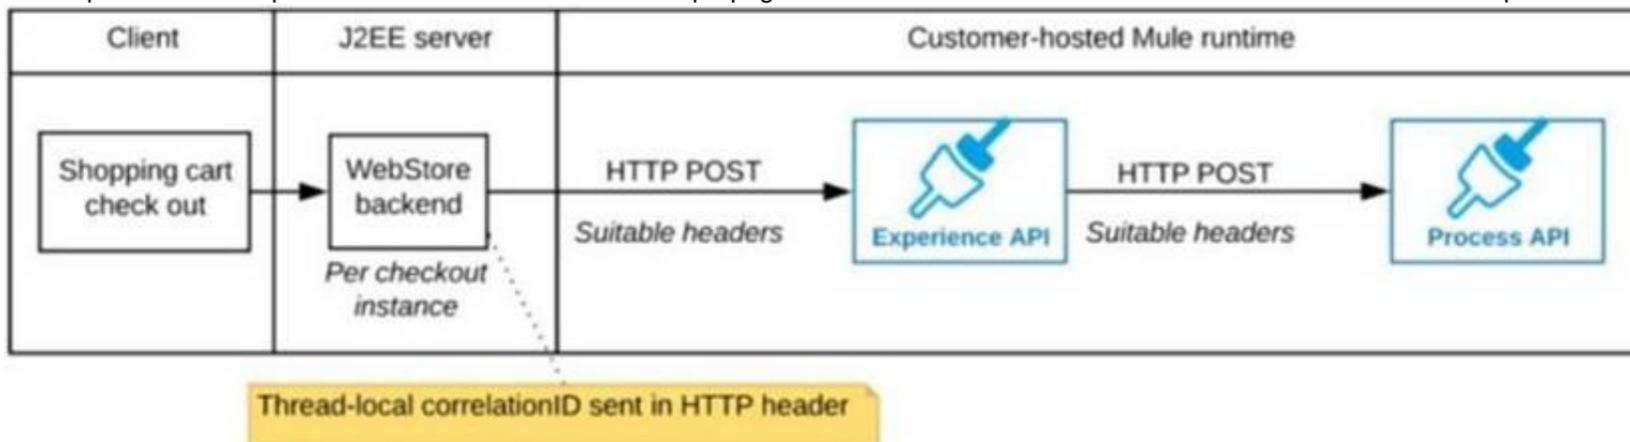
No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID



C)

The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

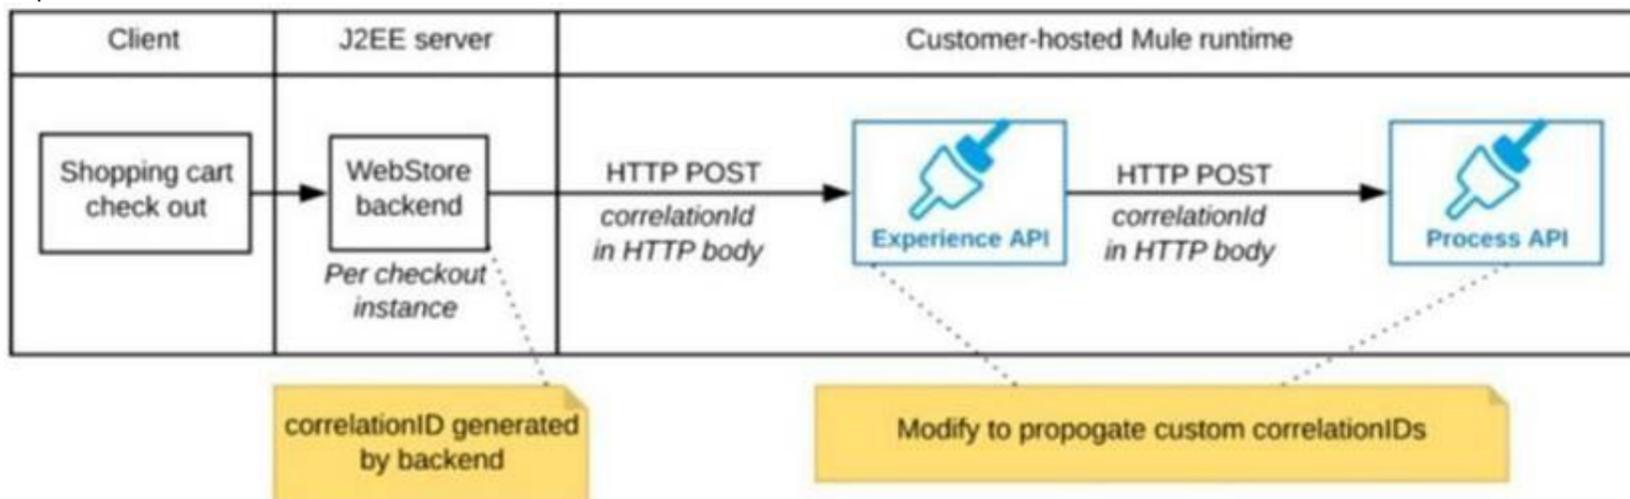
The Experience API implementation must be coded to also propagate the correlation ID to the Process API in a suitable HTTP request header



D)

The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API

The Experience API and Process API implementations must be coded to receive the custom correlation ID In the HTTP requests and propagate It in suitable HTTP request headers



- A. Option A
- B. Option B
- C. Option C

D. Option D

**Answer: B**

**Explanation:**

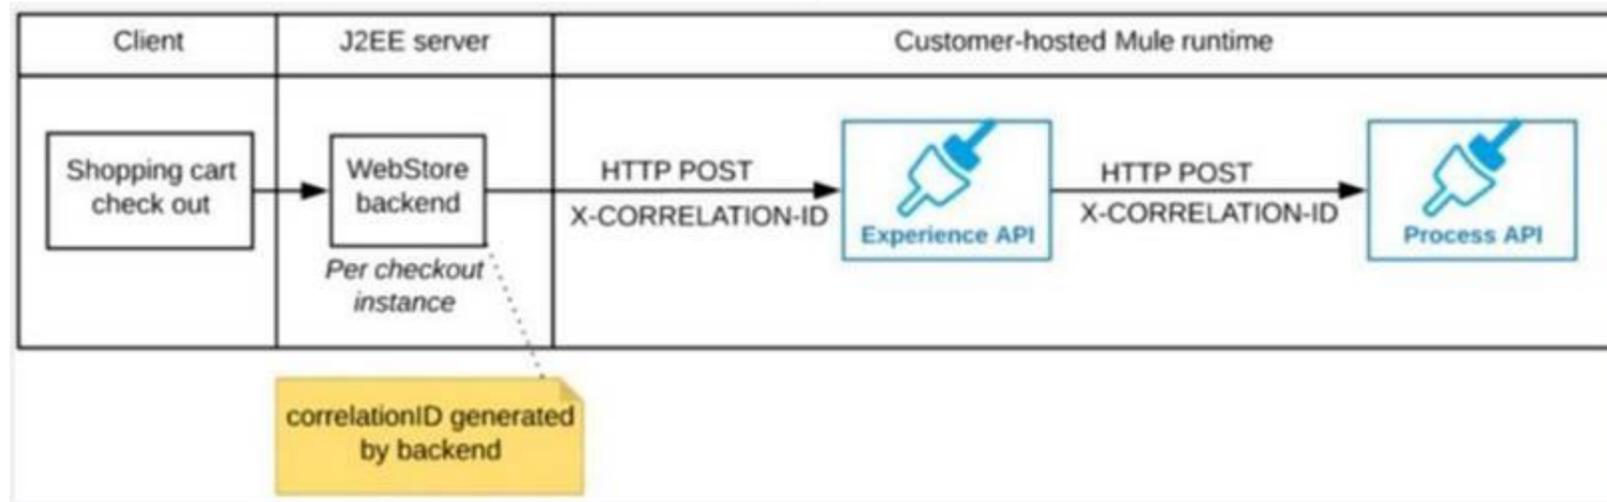
: By design, Correlation Ids cannot be changed within a flow in Mule 4 applications and can be set only at source. This ID is part of the Event Context and is generated as soon as the message is received by the application. When a HTTP Request is received, the request is inspected for "X-Correlation-Id" header. If "X-Correlation-Id" header is present, HTTP connector uses this as the Correlation Id. If "X-Correlation-Id" header is NOT present, a Correlation Id is randomly generated. For

Incoming HTTP Requests: In order to set a custom Correlation Id, the client invoking the HTTP request must set "X-Correlation-Id" header. This will ensure that the Mule Flow uses this Correlation Id. For Outgoing HTTP Requests: You can also propagate the existing Correlation Id to downstream APIs. By default, all outgoing HTTP Requests send "X-Correlation-Id" header. However, you can choose to set a different value to "X-Correlation-Id" header or set "Send Correlation Id" to NEVER.

Mulesoft Reference:

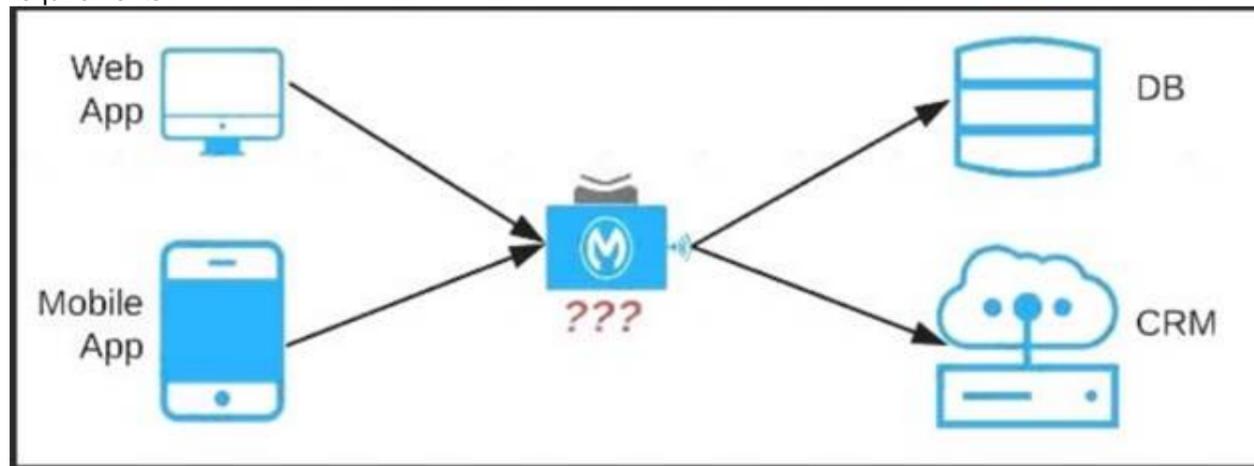
<https://help.mulesoft.com/s/article/How-to-Set-Custom-Correlation-Id-for-Flows-with-HTTP-Endpoint-in-Mule>

Graphical user interface, application, Word Description automatically generated



**NEW QUESTION 20**

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:**

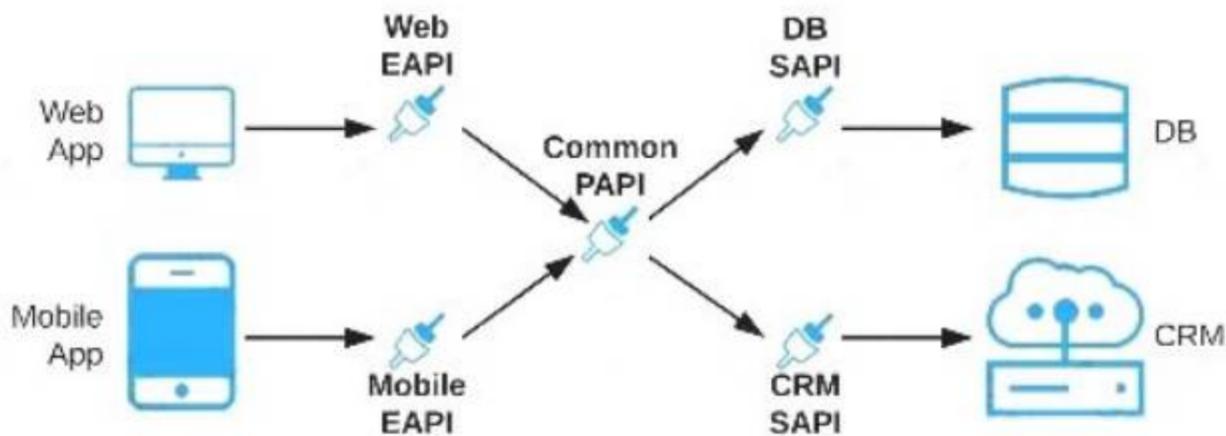
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements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



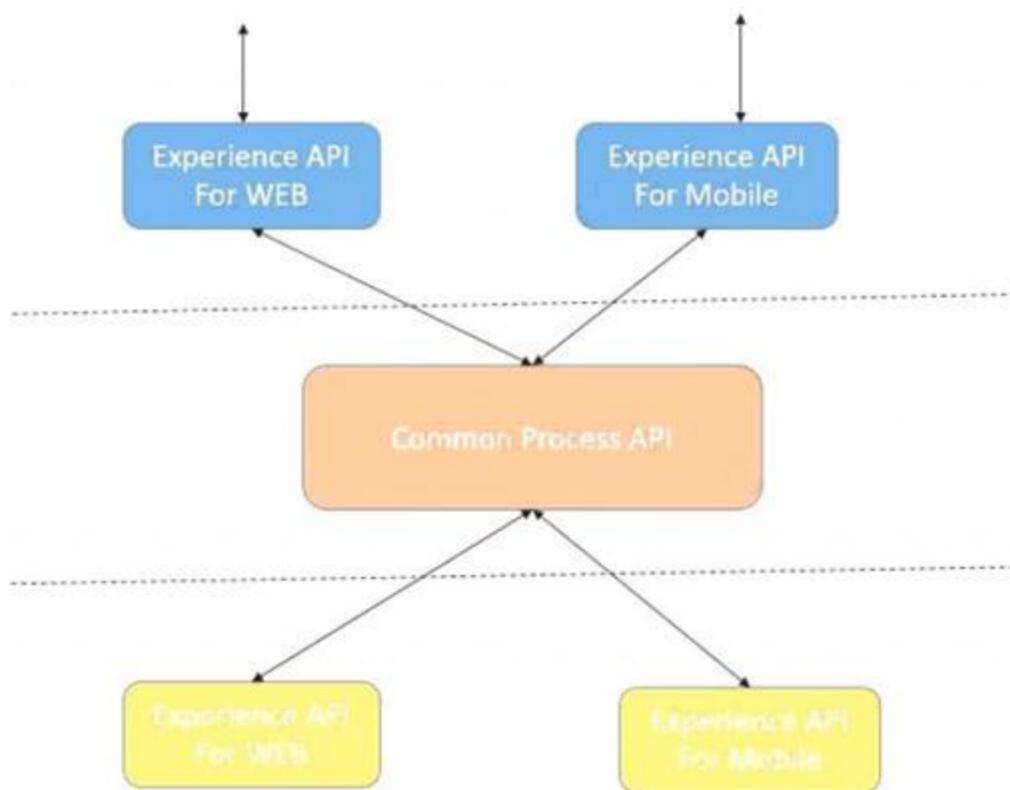
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements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

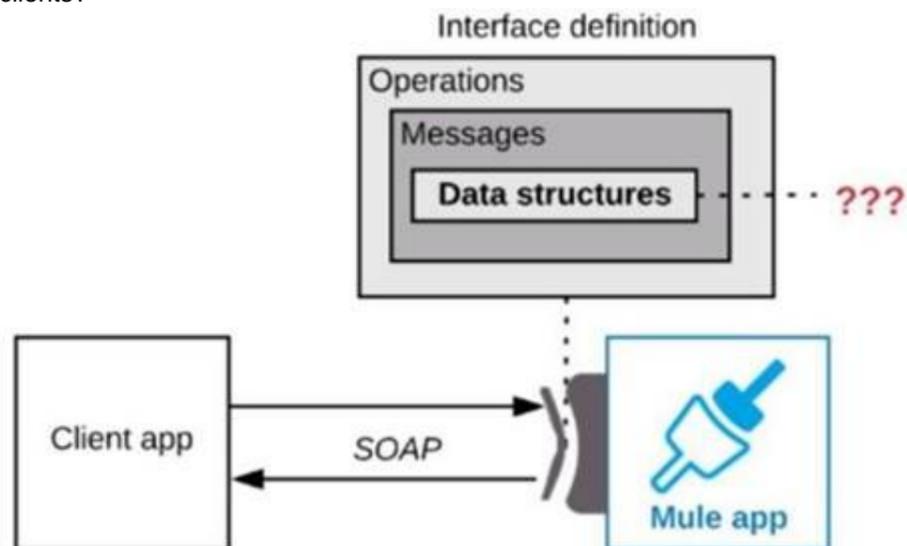


**NEW QUESTION 22**

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 27**

When designing an upstream API and its implementation, the development team has been advised to not set timeouts when invoking downstream API. Because the downstream API has no SLA that can be relied upon. This is the only downstream API dependency of that upstream API. Assume the downstream API runs uninterrupted without crashing. What is the impact of this advice?

- A. The invocation of the downstream API will run to completion without timing out.
- B. An SLA for the upstream API CANNOT be provided.
- C. A default timeout of 500 ms will automatically be applied by the Mule runtime in which the upstream API implementation executes.
- D. A load-dependent timeout of less than 1000 ms will be applied by the Mule runtime in which the downstream API implementation executes.

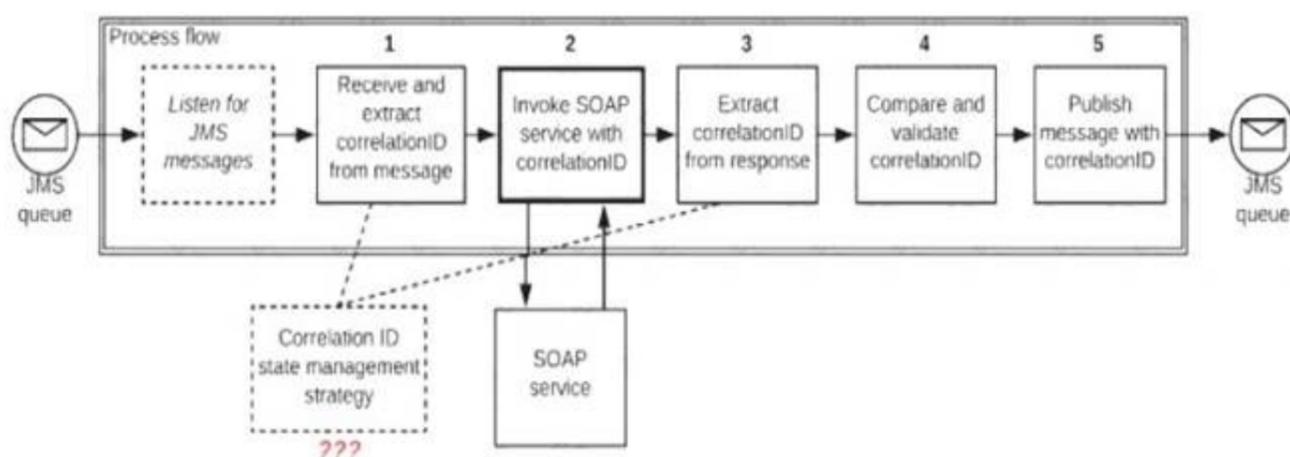
**Answer: B**

**Explanation:**

An SLA for the upstream API CANNOT be provided.

**NEW QUESTION 30**

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

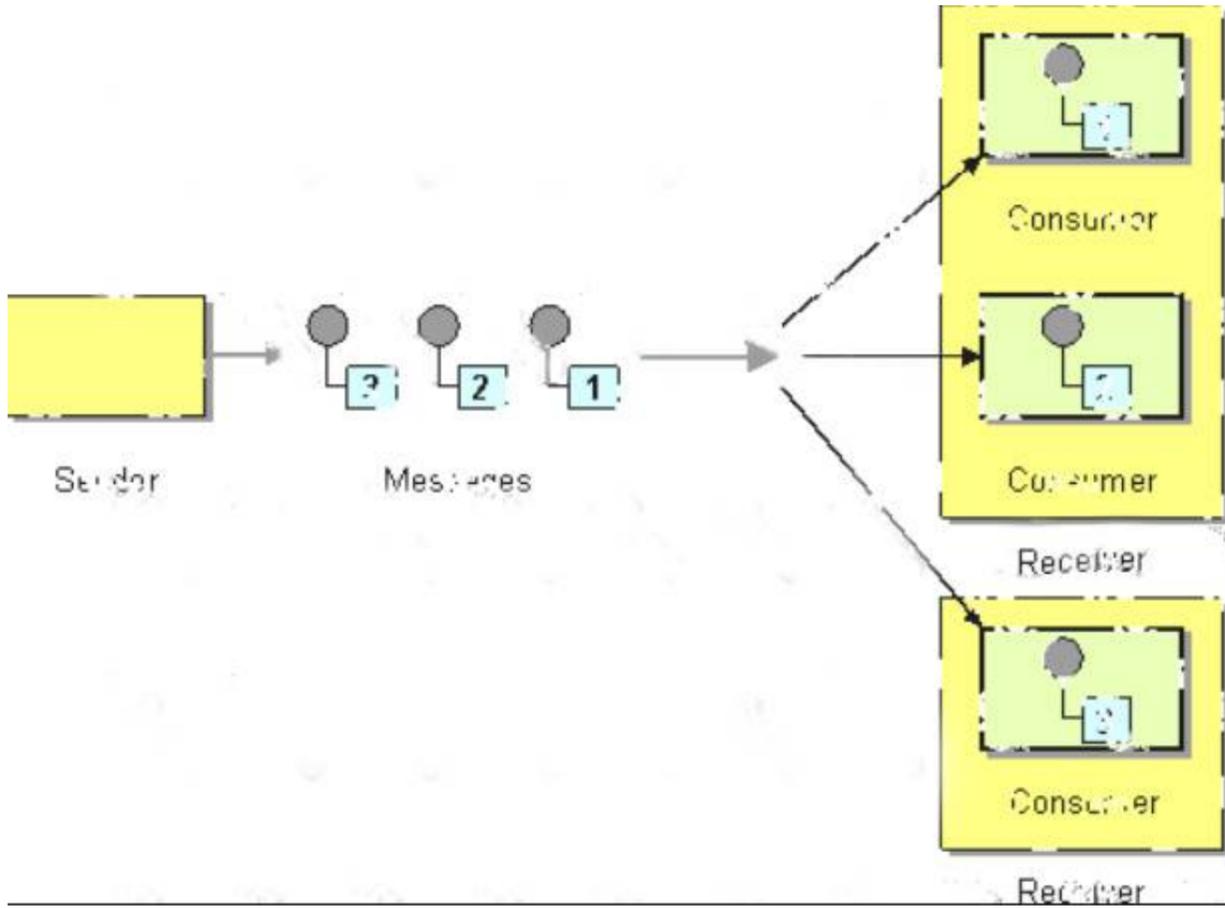
Additional Information:

\* 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 34**

A Mule application name Pub uses a persistence object store. The Pub Mule application is deployed to Cloudhub and it configured to use Object Store v2. Another Mule application name sub is being developed to retrieve values from the Pub Mule application persistence object Store and will also be deployed to cloudhub.

What is the most direct way for the Sub Mule application to retrieve values from the Pub Mule application persistence object store with the least latency?

- A. Use an object store connector configured to access the Pub Mule application persistence object store
- B. Use a VM connector configured to directly access the persistence queue of the Pub Mule application persistence object store.
- C. Use an Anypoint MQ connector configured to directly access the Pub Mule application persistence object store
- D. Use the Object store v2 REST API configured to access the Pub Mule application persistence object store.

**Answer: D**

**Explanation:**

- \* The Object Store V2 API enables API access to Anypoint Platform Object Store v2.
- \* You can configure a Mule app to use the Object Store REST API to store and retrieve values from an object store in another Mule app. However, Object Store v2 is not designed for app-to-app communication. To share data between two Mule4 apps, use a queue in Anypoint MQ.
- \* The Object Store v2 APIs enable you to use REST to perform the following:
  - Retrieve a list of object stores and keys associated with an application.
  - Store and retrieve key-value pairs in an object store.

- Delete key-value pairs from an object store.
- Retrieve Object Store usage statistics for your organization.
- Object Store provides these APIs: Object Store API  
Object Store Stats API

**NEW QUESTION 39**

An integration Mule application consumes and processes a list of rows from a CSV file. Each row must be read from the CSV file, validated, and the row data sent to a JMS queue, in the exact order as in the CSV file.

If any processing step for a row fails, then a log entry must be written for that row, but processing of other rows must not be affected.

What combination of Mule components is most idiomatic (used according to their intended purpose) when Implementing the above requirements?

- A. Scatter-Gather component On Error Continue scope
- B. VM connector first Successful scope On Error Propagate scope
- C. For Each scope On Error Continue scope
- D. Async scope On Error Propagate scope

**Answer: C**

**Explanation:**

\* On Error Propagate halts execution and sends error to the client. In this scenario it's mentioned that "processing of other rows must not be affected" so Option B and C are ruled out.

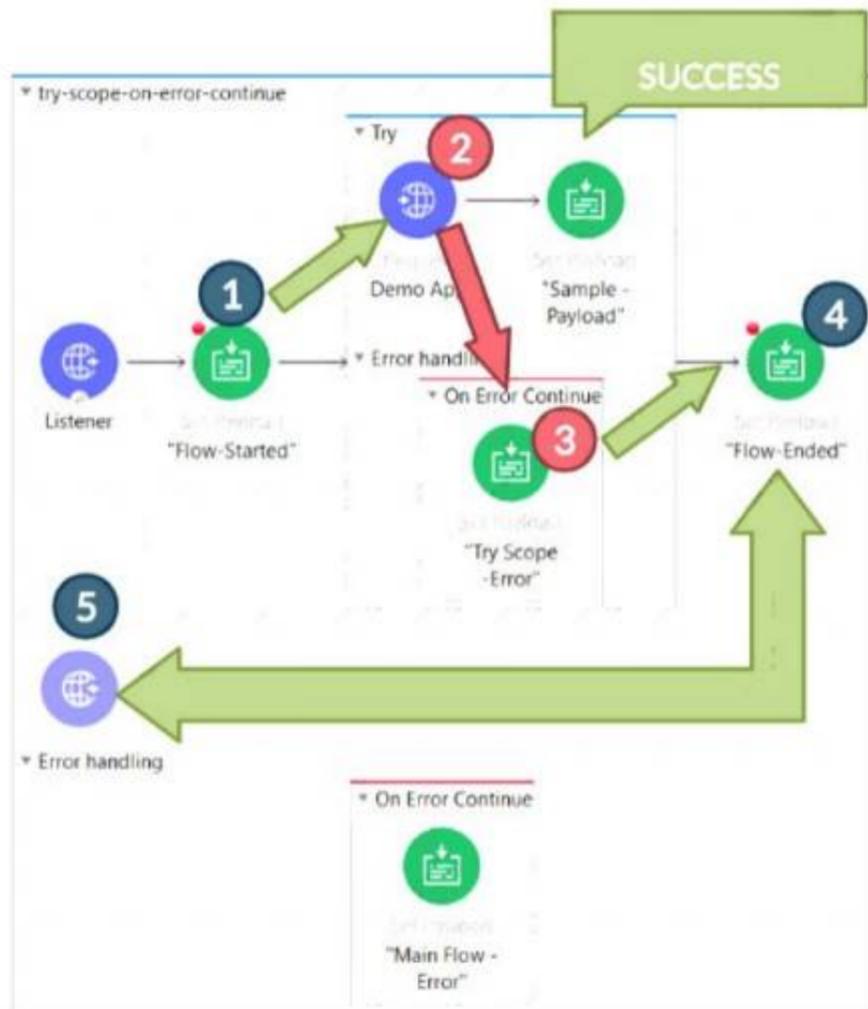
\* Scatter gather is used to club multiple responses together before processing. In this scenario, we need sequential processing. So option A is out of choice.

\* Correct answer is For Each scope & On Error Continue scope Below requirement can be fulfilled in the below way

1) Using For Each scope , which will send each row from csv file sequentially. each row needs to be sent sequentially as requirement is to send the message in exactly the same way as it is mentioned in the csv file

2) Also other part of requirement is if any processing step for a row fails then it should log an error but should not affect other record processing . This can be achieved using On error Continue scope on these set of activities. so that error will not halt the processing. Also logger needs to be added in error handling section so that it can be logged.

\* Attaching diagram for reference. Here it's try scope, but similar would be the case with For Each loop. Diagram Description automatically generated



**NEW QUESTION 44**

A mule application is deployed to a Single Cloudhub worker and the public URL appears in Runtime Manager as the APP URL.

Requests are sent by external web clients over the public internet to the mule application App url. Each of these requests routed to the HTTPS Listener event source of the running Mule application.

Later, the DevOps team edits some properties of this running Mule application in Runtime Manager. Immediately after the new property values are applied in runtime manager, how is the current Mule application deployment affected and how will future web client requests to the Mule application be handled?

- A. Cloudhub will redeploy the Mule application to the OLD Cloudhub workerNew web client requests will RETURN AN ERROR until the Mule application is redeployed to the OLD Cloudhub worker
- B. CloudHub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests will RETURN AN ERROR until the NEW Cloudhub worker is available
- C. Cloudhub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker until the NEW Cloudhub worker is available.
- D. Cloudhub will redeploy the mule application to the OLD Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.

**Answer:** C

**Explanation:**

CloudHub supports updating your applications at runtime so end users of your HTTP APIs experience zero downtime. While your application update is deploying, CloudHub keeps the old version of your application running. Your domain points to the old version of your application until the newly uploaded version is fully started. This allows you to keep servicing requests from your old application while the new version of your application is starting.

**NEW QUESTION 45**

A retailer is designing a data exchange interface to be used by its suppliers. The interface must support secure communication over the public internet. The interface must also work with a wide variety of programming languages and IT systems used by suppliers.

What are suitable interface technologies for this data exchange that are secure, cross-platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

- A. EDIFACT XML over SFTP JSON/REST over HTTPS
- B. SOAP over HTTPS HOP over TLS gRPC over HTTPS
- C. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
- D. CSV over FTP YAML over TLS JSON over HTTPS

**Answer:** C

**Explanation:**

As per definition of API by Mulesoft, it is Application Programming Interface using HTTP-based protocols. Non-HTTP-based programmatic interfaces are not APIs.

\* HTTP-based programmatic interfaces are APIs even if they don't use REST or JSON. Hence implementation based on Java RMI, CORBA/IIOP, raw TCP/IP interfaces are not API's as they are not using HTTP.

\* One more thing to note is FTP was not built to be secure. It is generally considered to be an insecure protocol because it relies on clear-text usernames and passwords for authentication and does not use encryption.

\* Data sent via FTP is vulnerable to sniffing, spoofing, and brute force attacks, among other basic attack methods.

Considering the above points only correct option is

- XML over ActiveMQ
- XML over SFTP
- XML/REST over HTTPS

**NEW QUESTION 49**

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 52**

How does timeout attribute help inform design decisions while using JMS connector listening for incoming messages in an extended architecture (XA) transaction?

- A. After the timeout is exceeded, stale JMS consumer threads are destroyed and new threads are created
- B. The timeout specifies the time allowed to pass between receiving JMS messages on the same JMS connection and then after the timeout new JMS connection is established
- C. The time allowed to pass between committing the transaction and the completion of the mule flow and then after the timeout flow processing triggers an error
- D. The timeout defines the time that is allowed to pass without the transaction ending explicitly and after the timeout expires, the transaction rolls back

**Answer:** D

**NEW QUESTION 54**

An organization has deployed both Mule and non-Mule API implementations to integrate its customer and order management systems. All the APIs are available to REST clients on the public internet.

The organization wants to monitor these APIs by running health checks: for example, to determine if an API can properly accept and process requests. The organization does not have subscriptions to any external monitoring tools and also does not want to extend its IT footprint.

What Anypoint Platform feature provides the most idiomatic (used for its intended purpose) way to monitor the availability of both the Mule and the non-Mule API implementations?

- A. API Functional Monitoring
- B. Runtime Manager
- C. API Manager
- D. Anypoint Visualizer

**Answer:** D

**NEW QUESTION 56**

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 60**

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.

**Answer: C**

**Explanation:**

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.

**NEW QUESTION 62**

A project uses Jenkins to implement CI/CD process. It was observed that each Mule package contains some of the Jenkins files and folders for configurations of CI/CD jobs.

As these files and folders are not part of the actual package, expectation is that these should not be part of deployed archive. Which file can be used to exclude these files and folders from the deployed archive?

- A. muleignore
- B. \_unTrackMule
- C. muleInclude
- D. \_muleExclude

**Answer: D**

**NEW QUESTION 66**

A mule application must periodically process a large dataset which varies from 6 GB lo 8 GB from a back-end database and write transform data lo an FTPS server using a properly configured bad job scope.

The performance requirements of an application are approved to run in the cloud hub 0.2 vCore with 8 GB storage capacity and currency requirements are met. How can the high rate of records be effectively managed in this application?

- A. Use streaming with a file storage repeatable strategy for reading records from the database and batch aggregator with streaming to write to FTPS
- B. Use streaming with an in-memory reputable store strategy for reading records from the database and batch aggregator with streaming to write to FTPS
- C. Use streaming with a file store repeatable strategy for reading records from the database and batch aggregator with an optimal size
- D. Use streaming with a file store repeatable strategy reading records from the database and batch aggregator without any required configuration

**Answer: A**

**NEW QUESTION 68**

An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/CD pipeline. Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?

- A. Use runtime manager rest apis for API management and mavenforAPI deployment
- B. Use Maven with a custom configuration required for the API lifecycle
- C. Use Anypoint CLI or Anypoint Platform REST apis with scripting language such as groovy
- D. Use Exchange rest api's for API management and MavenforAPI deployment

**Answer: D**

**NEW QUESTION 72**

An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub.

For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub's default zero-downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers. This is expected to result in several Mule application deployments per hour.

- A. 10.0.0.0/21(2048 IPs)
- B. 10.0.0.0/22(1024IPs)
- C. 10.0.0.0/23(512 IPs)
- D. 10.0.0.0/24(256 IPs)

**Answer:** A

**Explanation:**

- \* When you create an Anypoint VPC, the range of IP addresses for the network must be specified in the form of a Classless Inter-Domain Routing (CIDR) block, using CIDR notation.
- \* This address space is reserved for Mule workers, so it cannot overlap with any address space used in your data center if you want to peer it with your VPC.
- \* To calculate the proper sizing for your Anypoint VPC, you first need to understand that the number of dedicated IP addresses is not the same as the number of workers you have deployed.
- \* For each worker deployed to CloudHub, the following IP assignment takes place: For better fault tolerance, the VPC subnet may be divided into up to four Availability Zones.
- \* A few IP addresses are reserved for infrastructure. At least two IP addresses per worker to perform at zero-downtime.
- \* Hence in this scenario 2048 IP's are required to support the requirement.

**NEW QUESTION 73**

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 74**

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 76**

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL <https://api.acme.com>

As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub

What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

- A. Deploy the Mule applications to CloudhubUpdate the CNAME record for an api.acme.com in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map URLs to their corresponding Mule applications
- B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB)Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications
- C. Deploy the Mule applications to CloudhubCreate CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancerApply mapping rules in the SLB to map URLs to their corresponding Mule applications
- D. Deploy the Mule applications to CloudhubUpdate the CNAME record for api.acme.com in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB)Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

**Answer:** A

**Explanation:**

<https://help.mulesoft.com/s/feed/0D52T000055pzgsSAA>.

**NEW QUESTION 80**

Mule application muleA deployed in cloudhub uses Object Store v2 to share data across instances. As a part of new requirement , application muleB which is

deployed in same region wants to access this Object Store.

Which of the following option you would suggest which will have minimum latency in this scenario?

- A. Object Store REST API
- B. Object Store connector
- C. Both of the above option will have same latency
- D. Object Store of one mule application cannot be accessed by other mule application.

**Answer:** A

**Explanation:**

V2 Rest API is recommended for on premise applications to access Object Store. It also comes with overhead of encryption and security of using rest api. With Object Store v2, the API call is localized to the same data center as the Runtime Manager app.

But in this case requirement is to access the OS of other mule application and not the same mule application. You can configure a Mule app to use the Object Store REST API to store and retrieve values from an object store in another Mule app.

However, Object Store v2 is not designed for app-to-app communication.

**NEW QUESTION 81**

What limits if a particular Anypoint Platform user can discover an asset in Anypoint Exchange?

- A. Design Center and RAML were both used to create the asset
- B. The existence of a public Anypoint Exchange portal to which the asset has been published
- C. The type of the asset in Anypoint Exchange
- D. The business groups to which the user belongs

**Answer:** D

**Explanation:**

\* "The existence of a public Anypoint Exchange portal to which the asset has been published" - question does not mention anything about the public portal. Beside the public portal is open to the internet, to anyone. \* If you cannot find an asset in the current business group scopes, search in other scopes. In the left navigation bar click All assets (assets provided by MuleSoft and your own master organization), Provided by MuleSoft, or a business group scope. User belonging to one Business Group can see assets related to his group only Reference: <https://docs.mulesoft.com/exchange/to-find-info> <https://docs.mulesoft.com/exchange/asset-details> Correct answer is The business groups to which the user belongs

**NEW QUESTION 83**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### MCIA-Level-1 Practice Exam Features:

- \* MCIA-Level-1 Questions and Answers Updated Frequently
- \* MCIA-Level-1 Practice Questions Verified by Expert Senior Certified Staff
- \* MCIA-Level-1 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* MCIA-Level-1 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The MCIA-Level-1 Practice Test Here](#)