

## Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1

<https://www.2passeasy.com/dumps/MCIA-Level-1/>



### NEW QUESTION 1

Mule applications need to be deployed to CloudHub so they can access on-premises database systems. These systems store sensitive and hence tightly protected data, so are not accessible over the internet.

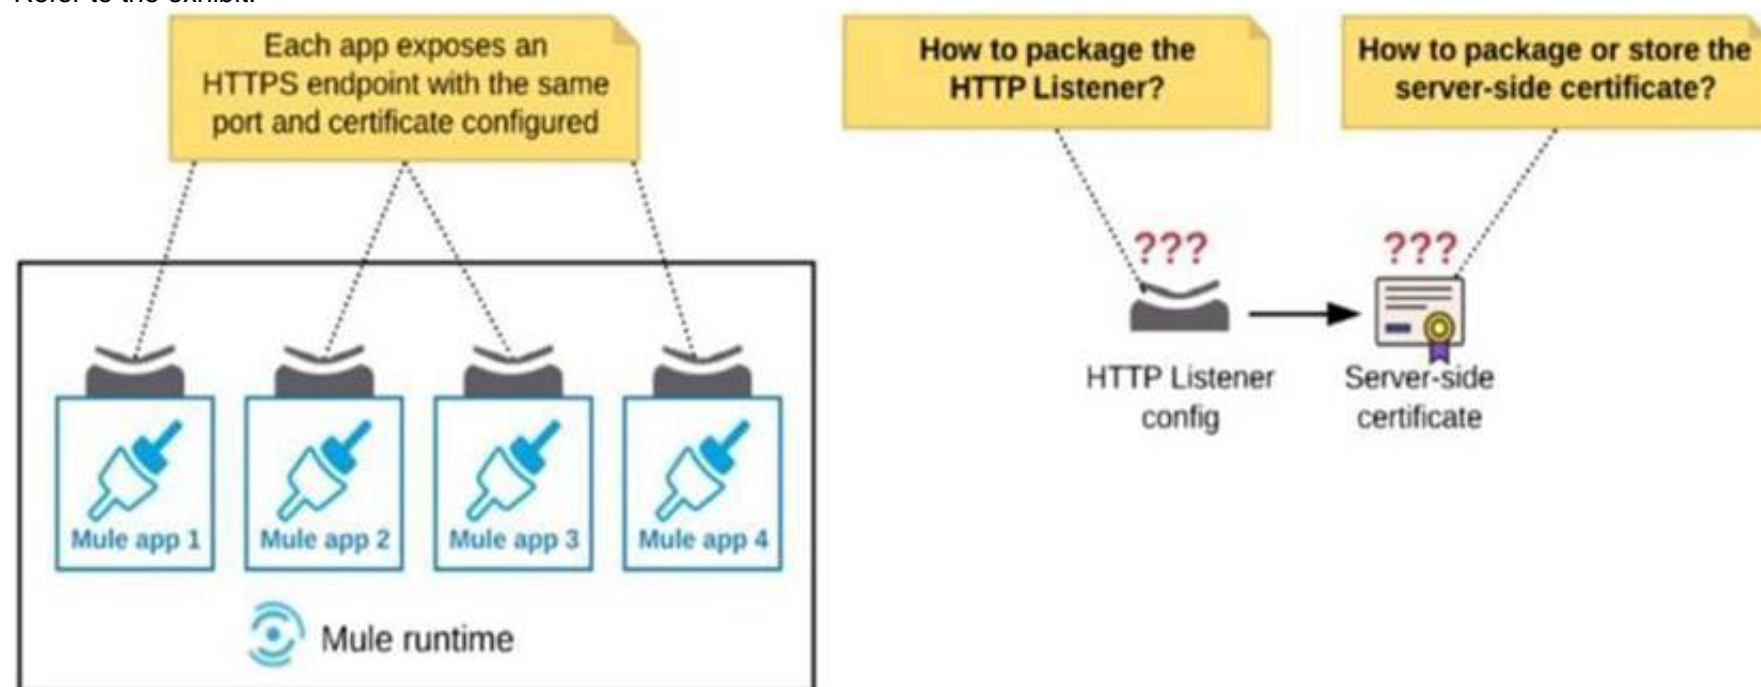
What network architecture supports this requirement?

- A. An Anypoint VPC connected to the on-premises network using an IPsec tunnel or AWSDirectConnect, plus matching firewall rules in the VPC and on-premises network
- B. Static IP addresses for the Mule applications deployed to the CloudHub Shared Worker Cloud, plus matching firewall rules and IP whitelisting in the on-premises network
- C. An Anypoint VPC with one Dedicated Load Balancer fronting each on-premises database system, plus matching IP whitelisting in the load balancer and firewall rules in the VPC and on-premises network
- D. Relocation of the database systems to a DMZ in the on-premises network, with Mule applications deployed to the CloudHub Shared Worker Cloud connecting only to the DMZ

Answer: A

### NEW QUESTION 2

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 Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION
- C. 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
- D. 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 the SAME Mule DOMAIN project Go to Set

Answer: B

### NEW QUESTION 3

An API client is implemented as a Mule application that includes an HTTP Request operation using a default configuration. The HTTP Request operation invokes an external API that follows standard HTTP status code conventions, which causes the HTTP Request operation to return a 4xx status code. What is a possible cause of this status code response?

- A. An error occurred inside the external API implementation when processing the HTTP request that was received from the outbound HTTP Request operation of the Mule application
- B. The external API reported that the API implementation has moved to a different external endpoint
- C. The HTTP response cannot be interpreted by the HTTP Request operation of the Mule application after it was received from the external API
- D. The external API reported an error with the HTTP request that was received from the outbound HTTP Request operation of the Mule application

Answer: D

### NEW QUESTION 4

An organization's governance process requires project teams to get formal approval from all key stakeholders for all new Integration design specifications. An integration Mule application is being designed that interacts with various backend systems. The Mule application will be created using Anypoint Design Center or Anypoint Studio and will then be deployed to a customer-hosted runtime.

What key elements should be included in the integration design specification when requesting approval for this Mule application?

- A. SLAs and non-functional requirements to access the backend systems

- B. Snapshots of the Mule application's flows, including their error handling
- C. A list of current and future consumers of the Mule application and their contact details
- D. The credentials to access the backend systems and contact details for the administrator of each system

**Answer:** A

#### NEW QUESTION 5

An integration Mute 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 falls, then a log entry must be written for that row, but processing of other rows must not be affected.

What combination of Mute 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

#### NEW QUESTION 6

Anypoint Exchange is required to maintain the source code of some of the assets committed to it, such as Connectors, Templates, and API specifications.

What is the best way to use an organization's source-code management (SCM) system in this context?

- A. Organizations should continue to use an SCM system of their choice, in addition to keeping source codefor these asset types in Anypoint Exchange, thereby enabling parallel development, branching, and merging
- B. Organizations need to use Anypoint Exchange as the main SCM system to centralize versioning and avoid code duplication
- C. Organizations can continue to use an SCM system of their choice for branching and merging, as long as they follow the branching and merging strategy enforced by Anypoint Exchange
- D. Organizations need to point Anypoint Exchange to their SCM system so Anypoint Exchange can pull source code when requested by developers and provide it to Anypoint Studio

**Answer:** A

#### NEW QUESTION 7

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.

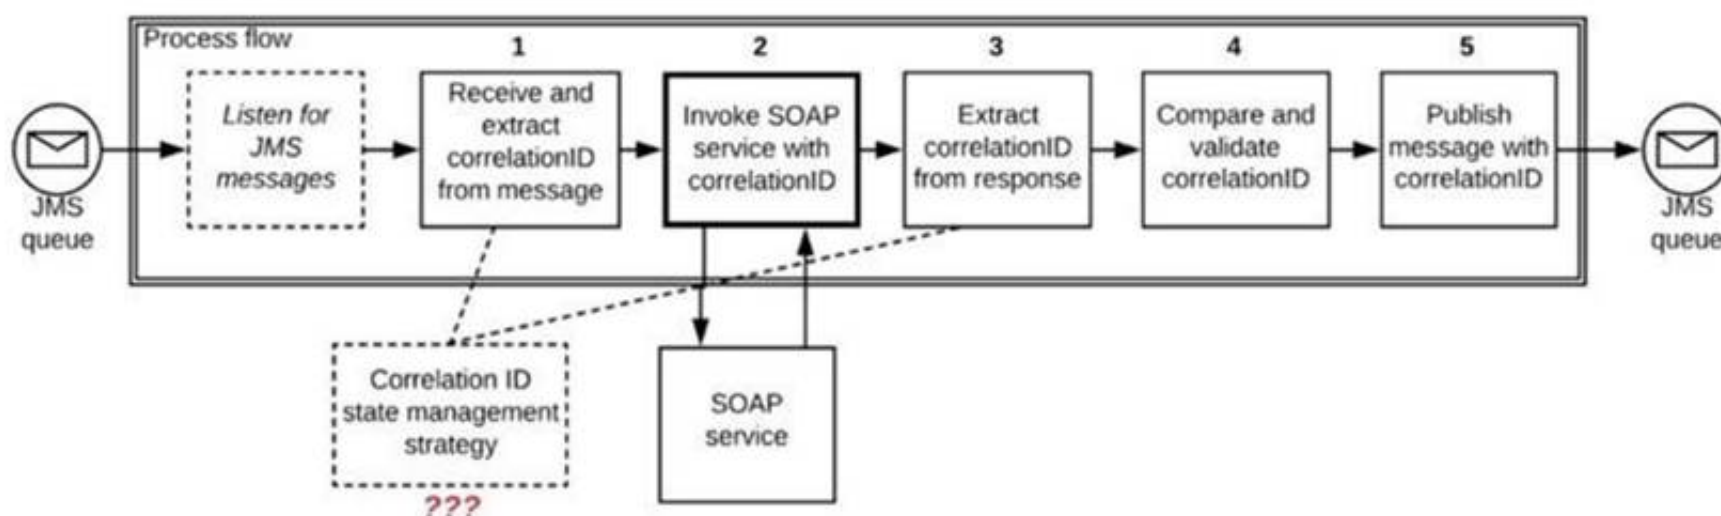
What combination of deployment options of the Anypoint Platform control plane and runtime plane(s) best serves this organization at the start of this strategic journey?

- 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:** B

#### NEW QUESTION 8

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 variables/attributes

Answer: C

#### NEW QUESTION 9

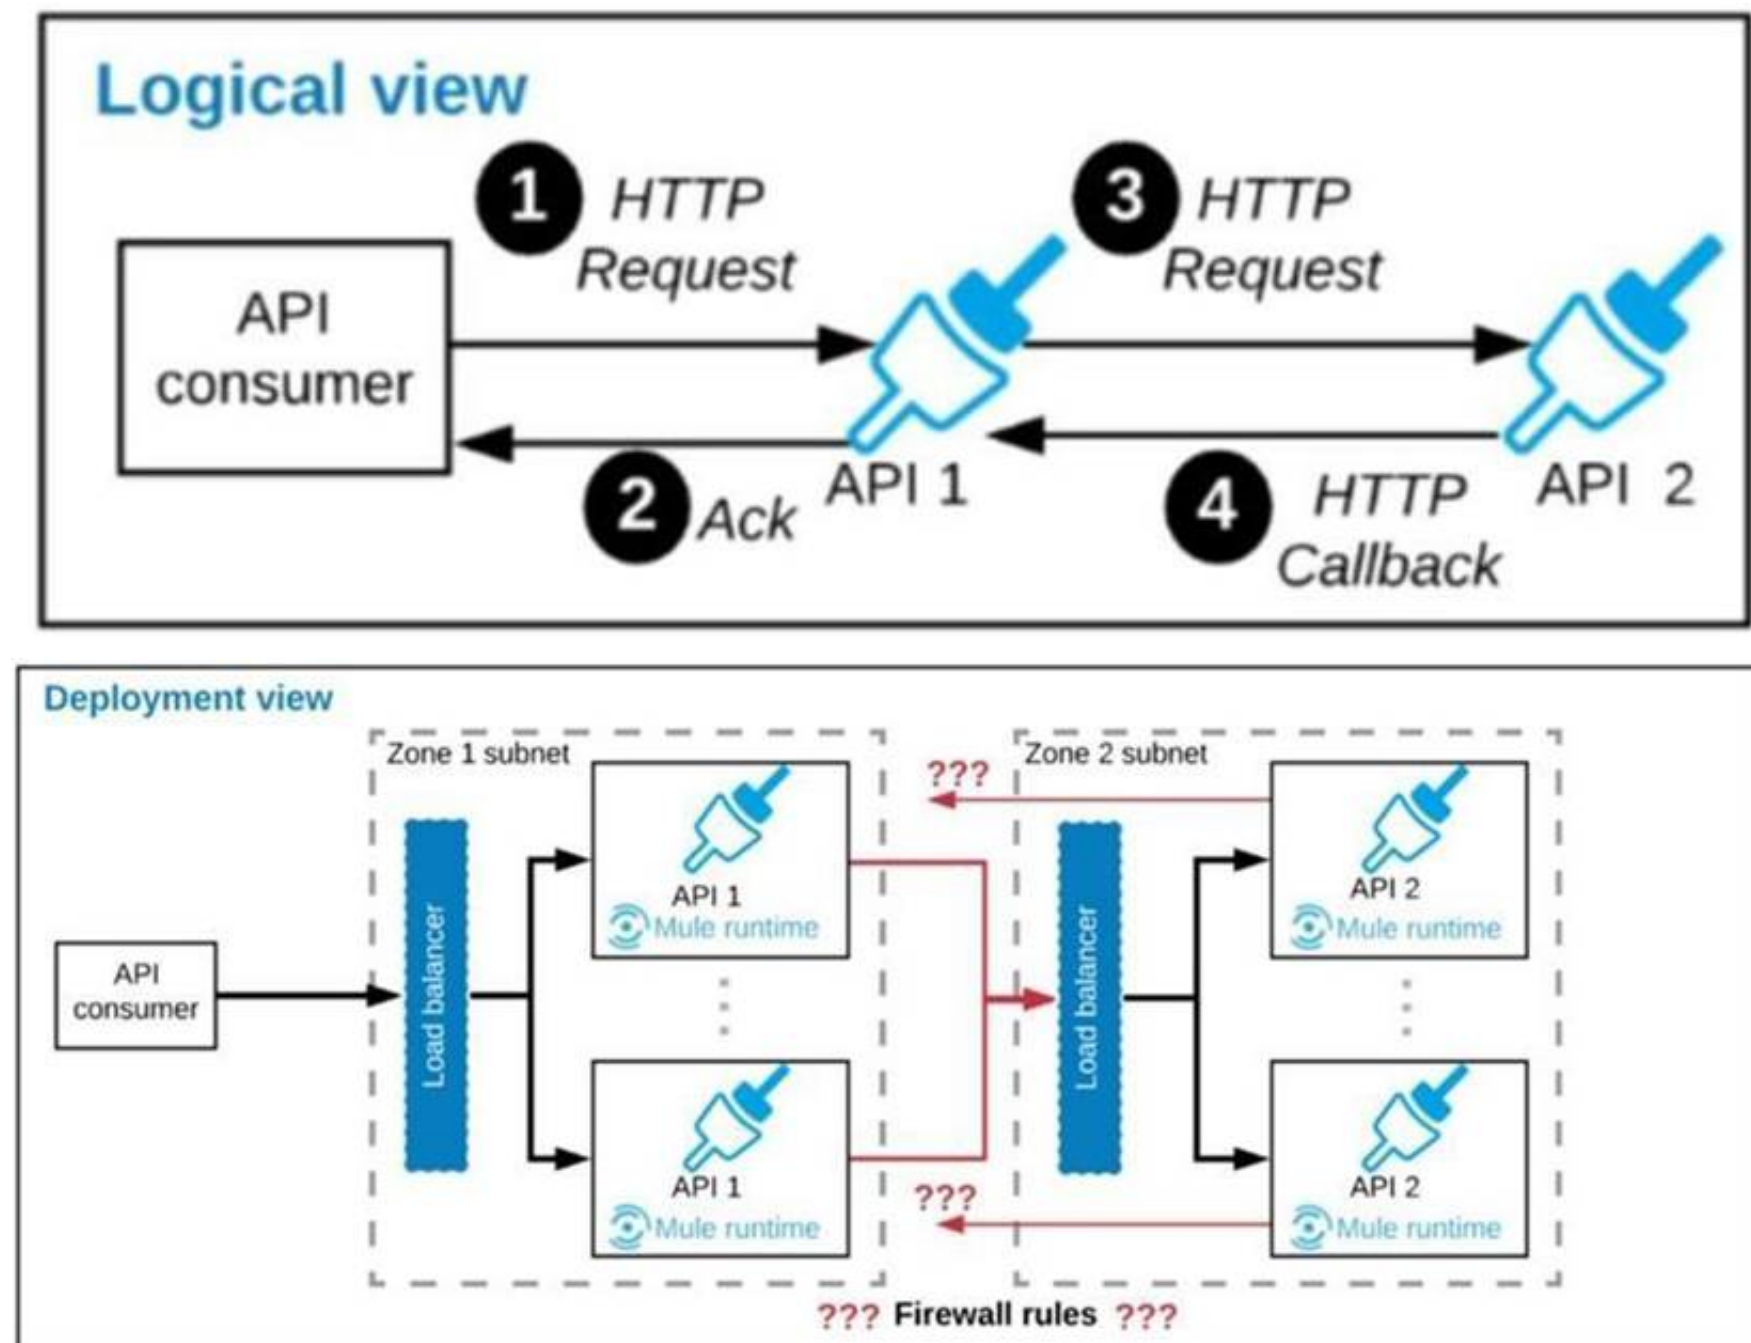
An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform. What feature of Anypoint Platform should be used to meet this requirement?

- A. Enterprise Security Module  
 B. Client ID Secret  
 C. Federated Identity Management  
 D. Federated Client Management

Answer: C

#### NEW QUESTION 10

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

Answer: C

#### NEW QUESTION 10

A Mule application uses the Database connector. What condition can the Mule application automatically adjust to or recover from without needing to restart or redeploy the Mule application?

- A. One of the stored procedures being called by the Mule application has been renamed  
 B. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade  
 C. The database server was unavailable for four hours due to a major outage but is now fully operational again

D. The credentials for accessing the database have been updated and the previous credentials are no longer valid

Answer: D

#### NEW QUESTION 15

A team would like to create a project skeleton that developers can use as a starting point when creating API implementations with Anypoint Studio. This skeleton should help drive consistent use of best practices within the team.

What type of Anypoint Exchange artifact(s) should be added to Anypoint Exchange to publish the project skeleton?

- A. A RAML archetype and reusable trait definitions to be reused across API implementations
- B. A custom asset with the default API implementation
- C. An example of an API implementation following best practices
- D. A Mule application template with the key components and minimal integration logic

Answer: D

#### NEW QUESTION 20

What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

- A. To send Mule application log entries to Splunk
- B. To change tog4j2 tog levels in Anypoint Runtime Manager without having to restart the Mule application
- C. To log certain messages to a custom log category
- D. To directly reference one shared and customized log4j2.xml file from multiple Mule applications

Answer: D

#### NEW QUESTION 21

An Integration Mule application is being designed to synchronize customer data between two systems. One system is an IBM Mainframe and the other system is a Salesforce Marketing Cloud (CRM) instance. Both systems have been deployed in their typical configurations, and are to be invoked using the native protocols provided by Salesforce and IBM.

What interface technologies are the most straightforward and appropriate to use in this Mule application to interact with these systems, assuming that Anypoint Connectors exist that implement these interface technologies?

- A. IBM: DB access CRM:gRPC
- B. IBM: REST CRM:REST
- C. IBM: ActiveMQ CRM: REST
- D. IBM:QCS CRM: SOAP

Answer: A

#### NEW QUESTION 26

An integration Mule application is deployed to a customer-hosted multi-node Mule 4 runtime duster. The Mule application uses a Listener operation of a JMS connector to receive incoming messages from a JMS queue.

How are the messages consumed by the Mule application?

- A. Depending on the JMS provider's configuration, either all messages are consumed by ONLY the primary cluster node or else ALL messages are consumed by ALL cluster nodes
- B. Regardless of the Listener operation configuration, all messages are consumed by ALL cluster nodes
- C. Depending on the Listener operation configuration, either all messages are consumed by ONLY the primary cluster node or else EACH message is consumed by ANY ONE cluster node
- D. Regardless of the Listener operation configuration, all messages are consumed by ONLY the primary cluster node

Answer: C

#### NEW QUESTION 27

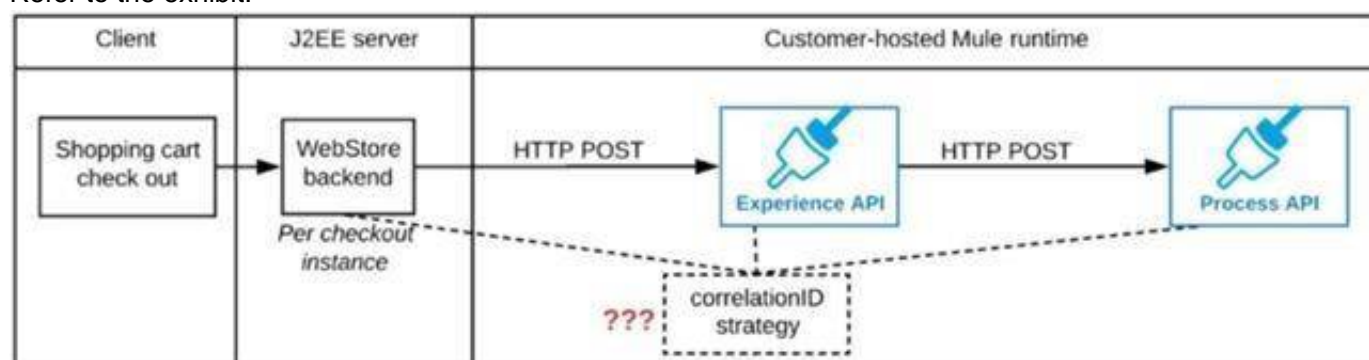
What metrics about API invocations are available for visualization in custom charts using Anypoint Analytics?

- A. Request size, request HTTP verbs, response time
- B. Request size, number of requests, JDBC Select operation result set size
- C. Request size, number of requests, JDBC Select operation response time
- D. Request size, number of requests, response size, response time

Answer: D

#### NEW QUESTION 30

Refer to the exhibit.

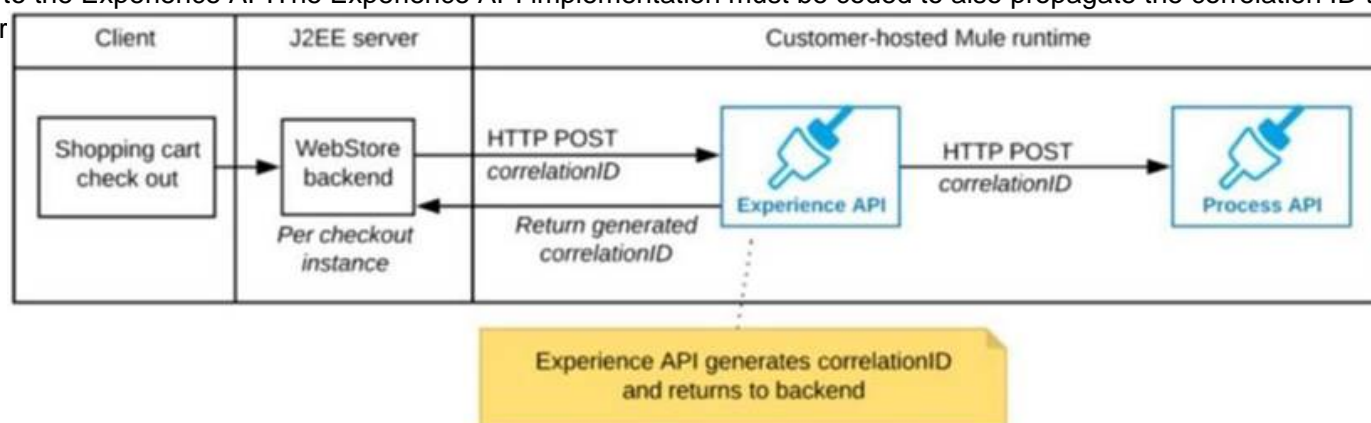


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations are Mule applications executing in a customer-hosted Mule runtime.

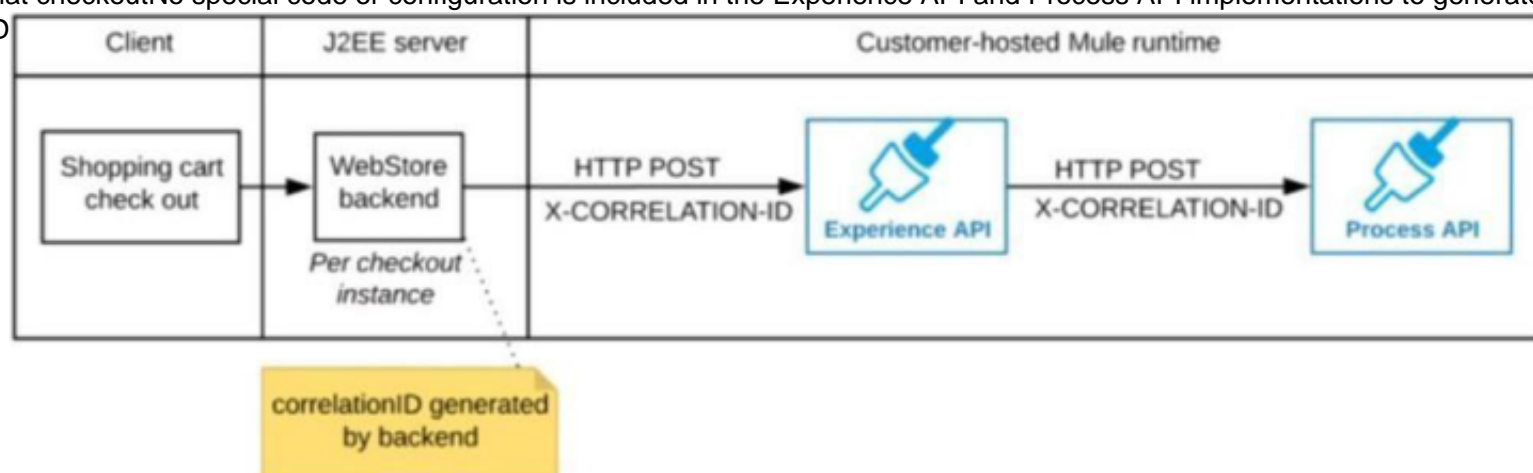
End-to-end correlation of all HTTP requests and responses belonging to each individual checkout instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the 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?

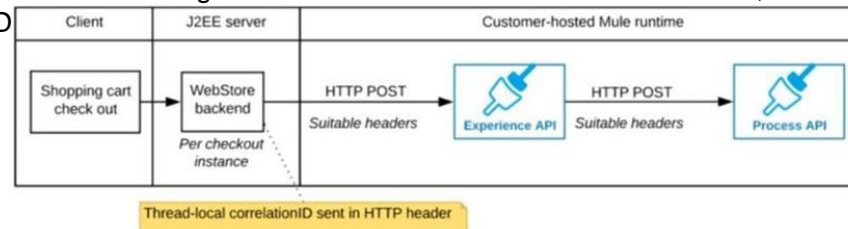
A. The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API. The Experience API implementation must be coded to also propagate the correlation ID to the Process API in a suitable HTTP request header.



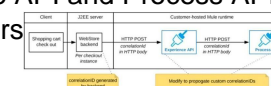
B. The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-ID HTTP request header in each API invocation belonging to that checkout. No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID.



C. The 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.



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.



Answer: B

### NEW QUESTION 34

A Mule application contains a Batch Job with two Batch Steps (Batch\_Step\_1 and Batch\_Step\_2). A payload with 1000 records is received by the Batch Job. How many threads are used by the Batch Job to process records, and how does each Batch Step process records within the Batch Job?

- A. Each Batch Job uses SEVERAL THREADS for the Batch Steps. Each Batch Step instance receives ONE record at a time as the payload, and BATCH STEP INSTANCES execute IN PARALLEL to process records and Batch Steps in ANY order as fast as possible.
- B. Each Batch Job uses SEVERAL THREADS for the Batch Steps. Each Batch Step instance receives ONE record at a time as the payload, and RECORDS are processed IN PARALLEL within and between the two Batch Steps.
- C. Each Batch Job uses a SINGLE THREAD for all Batch Steps. Each Batch Step instance receives ONE record at a time as the payload, and RECORDS are processed IN ORDER, first through Batch\_Step\_1 and then through Batch\_Step\_2.
- D. Each Batch Job uses a SINGLE THREAD to process a configured block size of record. Each Batch Step instance receives A BLOCK OF records as the payload, and BLOCKS of records are processed IN ORDER.

Answer: A

### NEW QUESTION 39

An Order microservice and a Fulfillment microservice are being designed to communicate with their clients through message-based integration (and NOT through API invocations).

The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mule application, the Fulfillment microservice.

The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message (a kind of event message). Each OrderFulfilled message can be consumed by any interested Mule application, and the Order microservice is one such Mule application.

What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

- A. Order messages are sent to an Anypoint MQ exchange. OrderFulfilled messages are sent to an Anypoint MQ queue. Both microservices interact with Anypoint

MQ as the message broker, which must therefore scale to support the load of both microservices

B. Order messages are sent to a JMS queueOrderFulfilled messages are sent to a JMS topicBoth microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices

C. Order messages are sent directly to the Fulfillment microservicesOrderFulfilled messages are sent directly to the Order microserviceThe Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load of each microservice

D. Order messages are sent to a JMS queueOrderFulfilled messages are sent to a JMS topicThe Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice

**Answer:** D

#### NEW QUESTION 44

An organization uses Mule runtimes which are managed by Anypoint Platform - Private Cloud Edition.

What MuleSoft component is responsible for feeding analytics data to non-MuleSoft analytics platforms?

- A. Anypoint Runtime Manager
- B. Anypoint Exchange
- C. Anypoint API Manager
- D. The Mule runtimes

**Answer:** A

#### NEW QUESTION 48

.....



## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual MCIA-Level-1 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the MCIA-Level-1 Product From:

<https://www.2passeasy.com/dumps/MCIA-Level-1/>

## Money Back Guarantee

### 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