

# Salesforce MuleSoft-Integration-Architect-I Exam Practice Test Questions Updated on a Regular Basis



2026 Latest Exams-boost MuleSoft-Integration-Architect-I PDF Dumps and MuleSoft-Integration-Architect-I Exam Engine Free Share: <https://drive.google.com/open?id=1-gCJ17CbICXYrqHVqrQpFxRz3puJZuAN>

Our MuleSoft-Integration-Architect-I study materials just need you to memorize all keypoints of the knowledge of the real exam. It is unnecessary to review all irrelevant knowledges. At present, our MuleSoft-Integration-Architect-I exam questions have helped thousands of people pass the exam and obtain the certificate. Also, the passing rate of our MuleSoft-Integration-Architect-I Training Materials is the highest according to our investigation. None of the other exam braindumps in the market has the pass rate high as 98% to 100% as our MuleSoft-Integration-Architect-I learning quiz.

## Salesforce MuleSoft-Integration-Architect-I Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>Designing Integration Solutions to Meet Security Requirements: This topic emphasizes securing access to the Anypoint Platform and APIs, using Anypoint Security, counteracting security vulnerabilities, and understanding audit logging capabilities.</li></ul>
Topic 2	<ul style="list-style-type: none"><li>Designing Integration Solutions to Meet Reliability Requirements: It includes selecting alternatives to traditional transactions, recognizing the purpose of various scopes and strategies, differentiating disaster recovery and high availability, and using local and XA transactions.</li></ul>

Topic 3	<ul style="list-style-type: none"> <li>Initiating Integration Solutions on Anypoint Platform: Summarizing MuleSoft Catalyst and Catalyst Knowledge Hub, differentiating between functional and non-functional requirements, selecting features for designing and managing APIs, and choosing deployment options are its sub-topics.</li> </ul>
Topic 4	<ul style="list-style-type: none"> <li>Designing for the Runtime Plane Technology Architecture: It includes analyzing Mule runtime clusters, designing solutions for CloudHub, choosing Mule runtime domains, leveraging Mule 4 class loader isolation, and understanding the reactive event processing model.</li> </ul>
Topic 5	<ul style="list-style-type: none"> <li>Applying DevOps Practices and Operating Integration Solutions: Its sub-topics are related to designing CI</li> <li>CD pipelines with MuleSoft plugins, automating interactions with Anypoint Platform, designing logging configurations, and identifying Anypoint Monitoring features.</li> </ul>

>> Reliable MuleSoft-Integration-Architect-I Braindumps Ppt <<

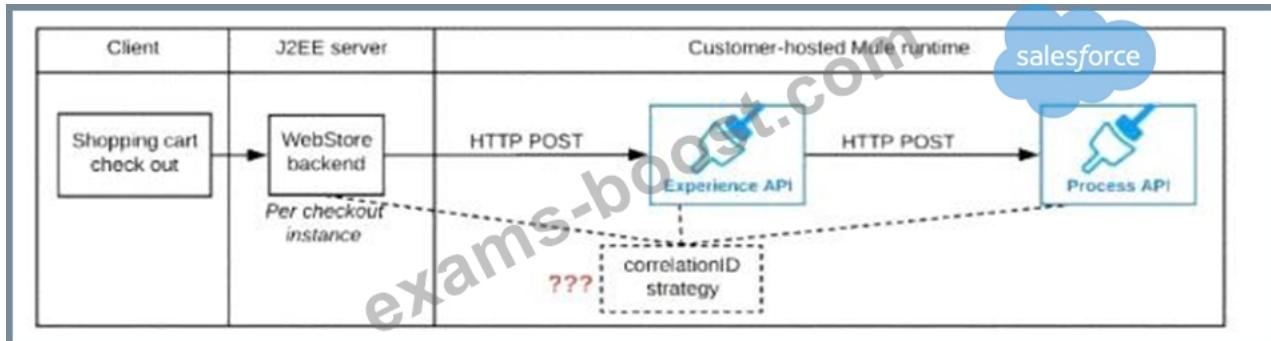
## MuleSoft-Integration-Architect-I Exam Certification, MuleSoft-Integration-Architect-I Review Guide

Being anxious for the MuleSoft-Integration-Architect-I exam ahead of you? Have a look of our MuleSoft-Integration-Architect-I training engine please. Presiding over the line of our practice materials over ten years, our experts are proficient as elites who made our MuleSoft-Integration-Architect-I learning questions, and it is their job to officiate the routines of offering help for you. All points are predominantly related with the exam ahead of you. You will find the exam is a piece of cake with the help of our MuleSoft-Integration-Architect-I Study Materials.

### Salesforce Certified MuleSoft Integration Architect I Sample Questions (Q238-Q243):

#### NEW QUESTION # 238

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?

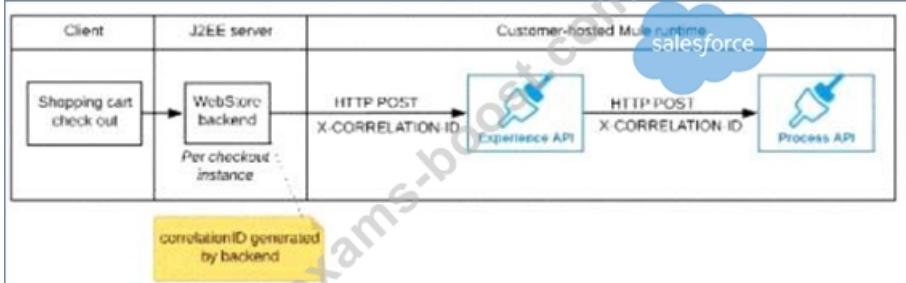
A)

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)

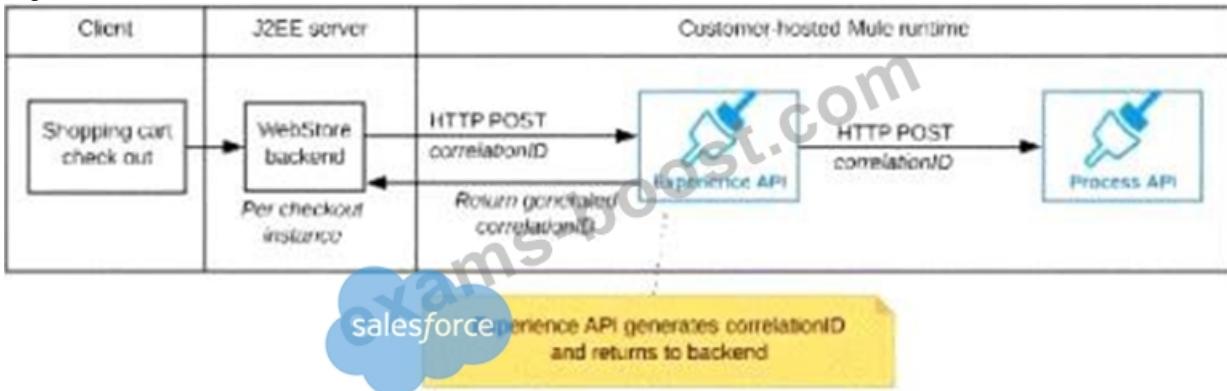
The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-ID HTTP request header. In each API invocation belonging to that checkout, no special code or configuration is included in the Experience API and Process API implementations to generate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



C)

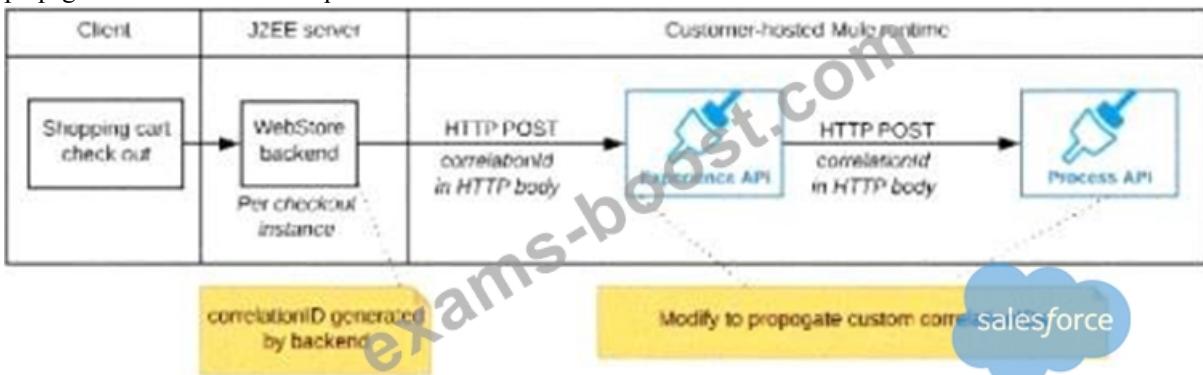
The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

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.

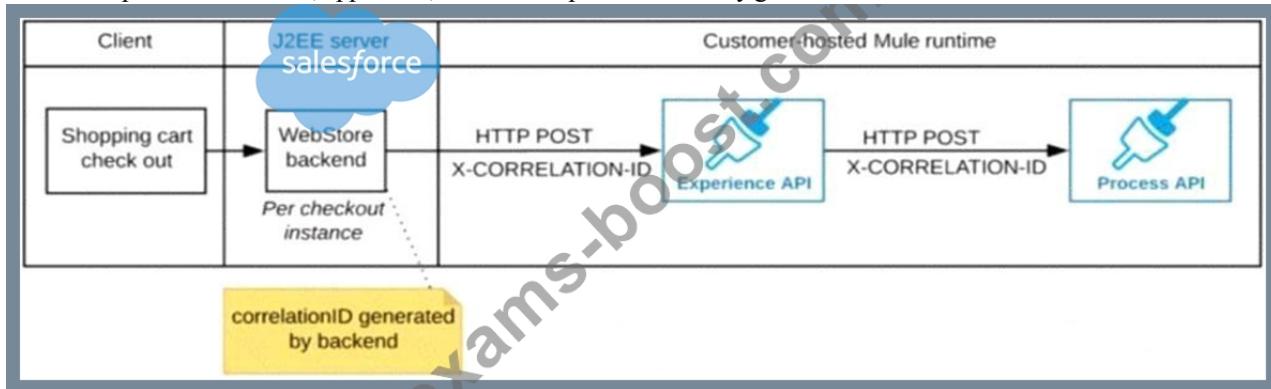


- A. Option A
- B. Option C
- C. Option B
- D. Option D

Answer: C

#### Explanation:

Correct answer is "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." 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-4>



#### NEW QUESTION # 239

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. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade
- B. The credentials for accessing the database have been updated and the previous credentials are no longer valid
- C. One of the stored procedures being called by the Mule application has been renamed
- D. The database server was unavailable for four hours due to a major outage but is now fully operational again**

#### Answer: D

#### Explanation:

\* Any change in the application will require a restart except when the issue outside the app. For below situations , you would need to redeploy the code after doing necessary changes

-- One of the stored procedures being called by the Mule application has been renamed. In this case, in the Mule application you will have to do changes to accommodate the new stored procedure name.

-- Required redesign of Mule applications to follow microservice architecture principles. As code is changed, deployment is must

-- If the credentials changed and you need to update the connector or the properties.

-- The credentials for accessing the database have been updated and the previous credentials are no longer valid. In this situation you need to restart or redeploy depending on how credentials are configured in Mule application.

\* So Correct answer is The database server was unavailable for four hours due to a major outage but is now fully operational again as this is the only external issue to application.

#### NEW QUESTION # 240

One of the backend systems involved by the API implementation enforces rate limits on the number of requests a particular client can make.

Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments.

The production environment however does not have any rate limiting.

What is the cost-effective approach to conduct performance test of the API implementation in the non- production staging environment?

- A. Conduct scaled-down performance tests in the staging environment against rate-limiting back-end system. Then upscale performance results to full production scale
- B. Use MUnit to simulate standard responses from the back-end system. Then conduct performance test to identify other bottlenecks in the system
- C. Include logic within the API implementation that bypasses locations of the back-end system in the staging environment and invoke a Mocking service that replicates typical back-end system responses. Then conduct performance test using this API implementation
- D. Create a Mocking service that replicates the back-end system's production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance test

**Answer: D**

Explanation:

To conduct performance testing in a non-production environment where rate limits are enforced, the most cost-effective approach is: C). Create a Mocking service that replicates the back-end system's production performance characteristics.

Then configure the API implementation to use the mocking service and conduct the performance test.

\* Mocking Service: Develop a mock service that emulates the performance characteristics of the production back-end system. This service should mimic the response times, data formats, and any relevant behavior of the actual back-end system without imposing rate limits.

\* Configuration: Modify the API implementation to route requests to the mocking service instead of the actual back-end system. This ensures that the performance tests are not impacted by the rate limits imposed in the non-production environment.

\* Performance Testing: Conduct the performance tests using the API implementation configured with the mocking service. This approach allows you to assess the performance under expected production load conditions without being constrained by non-production rate limits.

This method ensures that performance testing is accurate and reflective of the production environment without additional costs or constraints due to rate limiting in staging environments.

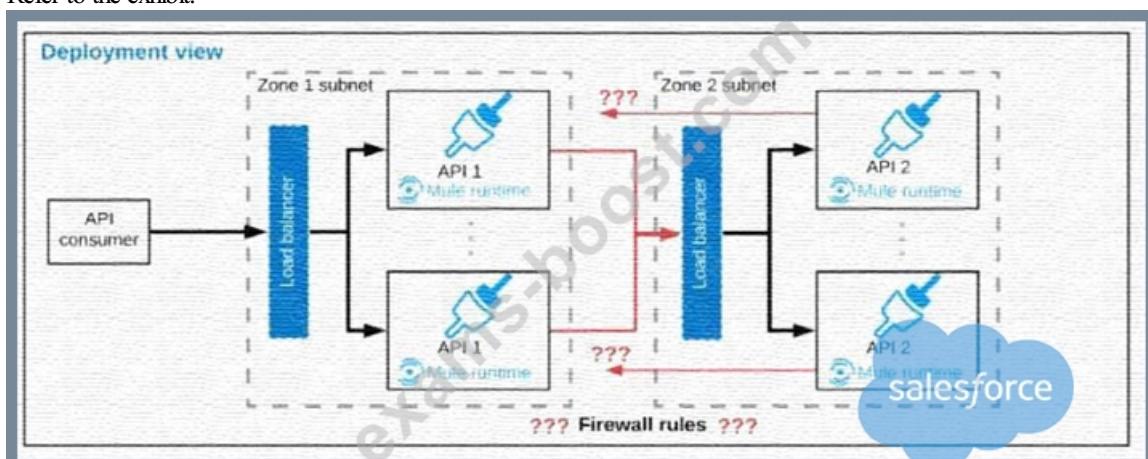
References:

MuleSoft Documentation: Mocking Services

MuleSoft Documentation: Performance Testing

## NEW QUESTION # 241

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

## Answer: A

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 baseUri 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 Reference: <https://docs.mulesoft.com/apikit/4.x/configure-load-balancer-task>

## NEW QUESTION # 242

What is a key difference between synchronous and asynchronous logging from Mule applications?

- A. Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event
- B. Synchronous logging within an ongoing transaction writes log messages in the same thread that processes the current Mule event
- C. Synchronous logging writes log messages in a single logging thread but does not block the Mule event being processed by the next event processor
- D. Asynchronous logging produces more reliable audit trails with more accurate timestamps

## Answer: A

Explanation:

Types of logging:

A) Synchronous: The execution of thread that is processing messages is interrupted to wait for the log message to be fully handled before it can continue.

# The execution of the thread that is processing your message is interrupted to wait for the log message to be fully output before it can continue

# Performance degrades because of synchronous logging

# Used when the log is used as an audit trail or when logging ERROR/CRITICAL messages

# If the logger fails to write to disk, the exception would raise on the same thread that's currently processing the Mule event. If logging is critical for you, then you can rollback the transaction.

B) Asynchronous:

# The logging operation occurs in a separate thread, so the actual processing of your message won't be delayed to wait for the logging to complete

# Substantial improvement in throughput and latency of message processing

# Mule runtime engine (Mule) 4 uses Log4j 2 asynchronous logging by default

# The disadvantage of asynchronous logging is error handling.

# If the logger fails to write to disk, the thread doing the processing won't be aware of any issues writing to the disk, so you won't be able to rollback anything. Because the actual writing of the log gets deferred, there's a chance that log messages might never make it to disk and get lost, if Mule were to crash before the buffers are flushed.

----- So Correct answer is: Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event

## NEW QUESTION # 243

.....

For most people who have no much time to prepare the Salesforce real exam, latest MuleSoft-Integration-Architect-I exam questions will be your excellent partner to help you get high passing score in the valid test. Once you receive our MuleSoft-Integration-Architect-I Dumps Torrent, it will just need one or two days to practice test questions and answers. If you finished it well, clearing exam will be easy.

**MuleSoft-Integration-Architect-I Exam Certification:** <https://www.exams-boost.com/MuleSoft-Integration-Architect-I-valid-materials.html>

P.S. Free & New MuleSoft-Integration-Architect-I dumps are available on Google Drive shared by Exams-boost: <https://drive.google.com/open?id=1-gCJ17CbfxYRqHVqrQpFxRz3puJZuAN>