



# MCIA-LEVEL-1<sup>Q&As</sup>

MuleSoft Certified Integration Architect - Level 1

## Pass Mulesoft MCIA-LEVEL-1 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/mcia-level-1.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Mulesoft  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

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 was unavailable for four hours due to a major outage but is now fully operational again
- C. The credentials for accessing the database have been updated and the previous credentials are no longer valid
- D. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade

Correct Answer: B

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 -- The credentials for accessing the database have been updated and the previous credentials are no longer valid.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

---

### QUESTION 2

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.

Correct Answer: D

Explanation: ?xplanation

\*

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 Reference:  
<https://docs.mulesoft.com/object-store/osv2-apis>

Additional Info:

When to use Object Store and when to use VM

### QUESTION 3

Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improved throughput of the system. But it was observed over the period of time that few of the important exception log messages which were used to rollback transactions are not working as expected causing huge loss to the Organization. Organization wants to avoid these losses. Application also has constraints due to which they cant compromise on throughput much. What is the possible option in this case?

- A. Logging needs to be changed from asynchronous to synchronous
- B. External log appender needs to be used in this case
- C. Persistent memory storage should be used in such scenarios
- D. Mixed configuration of asynchronous or synchronous loggers should be used to log exceptions via synchronous way

Correct Answer: D

Explanation: Correct approach is to use Mixed configuration of asynchronous or synchronous loggers should be used to log exceptions via synchronous way Asynchronous logging poses a performance-reliability trade-off. You may lose



some messages if Mule crashes before the logging buffers flush to the disk. In this case, consider that you can have a mixed configuration of asynchronous or synchronous loggers in your app. Best practice is to use asynchronous logging over synchronous with a minimum logging level of WARN for a production application. In some cases, enable INFO logging level when you need to confirm events such as successful policy installation or to perform troubleshooting. Configure your logging strategy by editing your application's src/main/resources/log4j2.xml file

---

#### QUESTION 4

What is an advantage that Anypoint Platform offers by providing universal API management and Integration-Platform-as-a-Service (iPaaS) capabilities in a unified platform?

- A. Ability to use a single iPaaS to manage and integrate all API gateways
- B. Ability to use a single connector to manage and integrate all APIs
- C. Ability to use a single control plane for both full-lifecycle API management and integration
- D. Ability to use a single iPaaS to manage all API developer portals

Correct Answer: C

Explanation: Anypoint Platform offers universal API management and Integration- Platform-as-a-Service (iPaaS) capabilities in a unified platform, meaning that it provides a single control plane to manage both full-lifecycle API management and integration. This allows organizations to easily manage their APIs and integrations, as well as deploy APIs and integrations quickly and efficiently. According to the MuleSoft Certified Integration Architect - Level 1 Course Book, "Anypoint Platform provides a unified platform for managing, deploying, and monitoring both API and integration solutions, allowing organizations to quickly and easily build and manage their APIs and integrations."

---

#### QUESTION 5

An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders.

The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS).

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

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

Correct Answer: C



Reference: <https://www.caeliusconsulting.com/blogs/one-way-and-two-way-tls-and-their-implementation-in-mulesoft/>

---

#### QUESTION 6

A key CI/CD capability of any enterprise solution is a testing framework to write and run repeatable tests.

Which component of Anypoint Platform provides the test automation capabilities for customers to use in their pipelines?

- A. Anypoint CLI
- B. Mule Maven Plugin
- C. Exchange Mocking Service
- D. MUnit

Correct Answer: D

---

#### QUESTION 7

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

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

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

Correct Answer: A

---

#### QUESTION 8

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

Correct 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.

\*

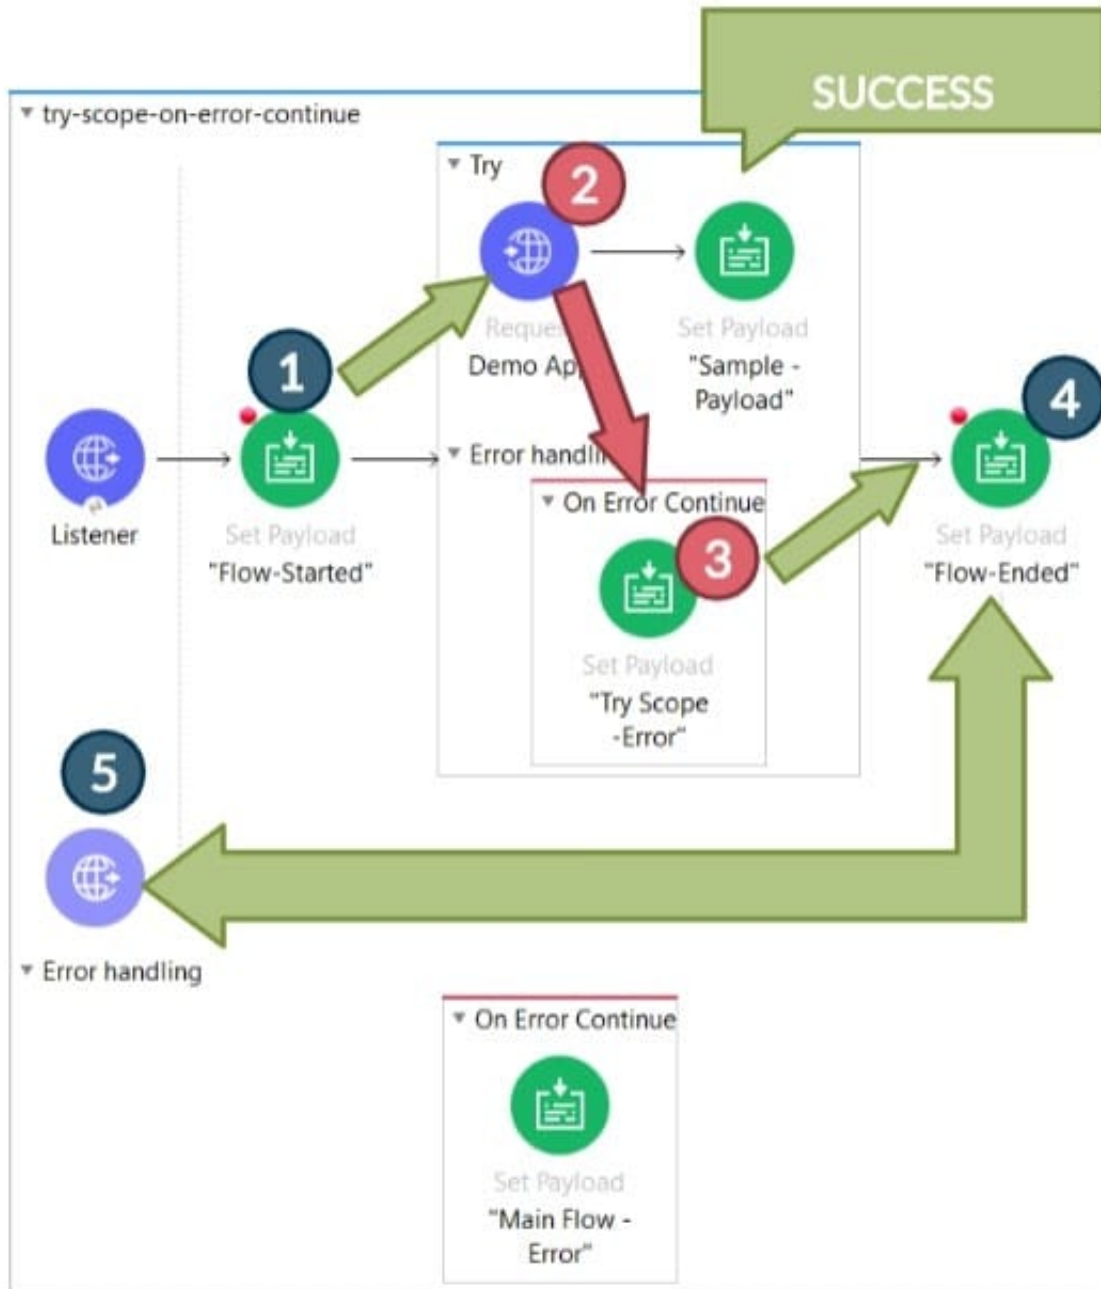
Scatter gather is used to club multiple responses together before processing. In this scenario, we need sequential processing. So option A is out of choice.

\*

Correct answer is For Each scope and 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.



**QUESTION 9**

An insurance company is implementing a MuleSoft API to get inventory details from the two vendors. Due to network issues, the invocations to vendor applications are getting timed- out intermittently. But the transactions are successful upon reprocessing

What is the most performant way of implementing this requirement?

A. Implement a scatter-gather scope to invoke the two vendor applications on two different route Use the Until-Successful scope to implement the retry mechanism for timeout errors on each route



- B. Implement a Choice scope to invoke the two vendor applications on two different route Use the try-catch scope to implement the retry mechanism for timeout errors on each route
- C. Implement a For-Each scope to invoke the two vendor applications Use until successful scope to implement the retry mechanism for the timeout errors
- D. Implement Round-Robin scope to invoke the two vendor applications on two different routes Use the Try-Catch scope to implement retry mechanism for timeout errors on each route

Correct Answer: A

---

### QUESTION 10

A stock broking company makes use of CloudHub VPC to deploy Mule applications. Mule application needs to connect to a database application in the customers on-premises corporate data center and also to a Kafka cluster running in AWS VPC.

How is access enabled for the API to connect to the database application and Kafka cluster securely?

- A. Set up a transit gateway to the customers on-premises corporate datacenter to AWS VPC
- B. Setup AnyPoint VPN to the customer's on-premise corporate data center and VPC peering with AWS VPC
- C. Setup VPC peering with AWS VPC and the customers devices corporate data center
- D. Setup VPC peering with the customers onto my service corporate data center and Anypoint VPN to AWS VPC

Correct Answer: B

---

### QUESTION 11

A company is designing a mule application to consume batch data from a partner's ftps server The data files have been compressed and then digitally signed using PGP.

What inputs are required for the application to securely consumed these files?

- A. ATLS context Key Store requiring the private key and certificate for the company PGP public key of partner PGP private key for the company
- B. ATLS context first store containing a public certificate for partner ftps server and the PGP public key of the partner TLS contact Key Store containing the FTP credentials
- C. TLS context trust or containing a public certificate for the ftps server The FTP username and password The PGP public key of the partner
- D. The PGP public key of the partner The PGP private key for the company The FTP username and password

Correct Answer: D

---





## QUESTION 12

What requires configuration of both a key store and a trust store for an HTTP Listener?

- A. Support for TLS mutual (two-way) authentication with HTTP clients
- B. Encryption of requests to both subdomains and API resource endpoints `fhttps://aDi.customer.com/` and `https://customer.com/api`
- C. Encryption of both HTTP request and HTTP response bodies for all HTTP clients
- D. Encryption of both HTTP request header and HTTP request body for all HTTP clients

Correct Answer: A

Explanation:

1 way SSL : The server presents its certificate to the client and the client adds it to its list of trusted certificate. And so, the client can talk to the server. 2-way SSL: The same principle but both ways. i.e. both the client and the server has to establish trust between themselves using a trusted certificate. In this way of a digital handshake, the server needs to present a certificate to authenticate itself to client and client has to present its certificate to server.

\*

TLS is a cryptographic protocol that provides communications security for your Mule app.

\*

TLS offers many different ways of exchanging keys for authentication, encrypting data, and guaranteeing message integrity. Keystores and Truststores. Truststore and keystore contents differ depending on whether they are used for clients or

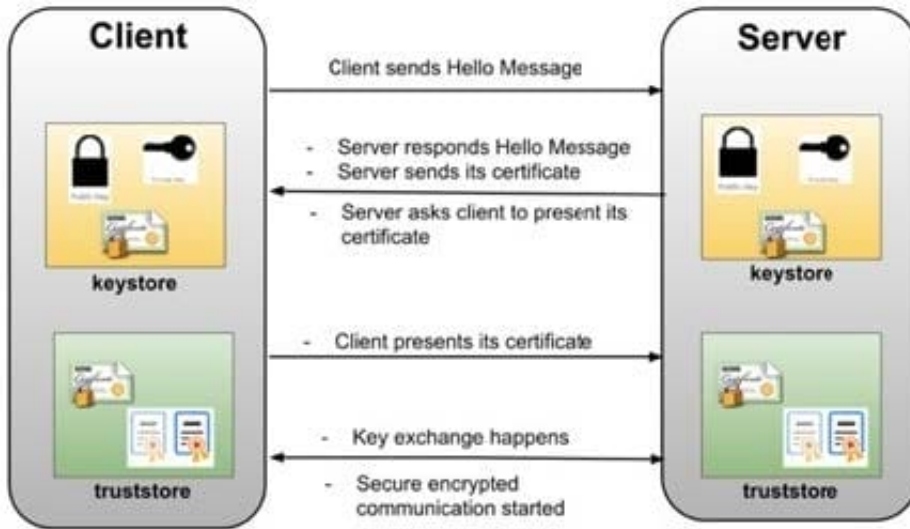
servers:

For servers: the truststore contains certificates of the trusted clients, the keystore contains the private and public key of the server. For clients: the truststore contains certificates of the trusted servers, the keystore contains the private and

public key of the client. Adding both a keystore and a truststore to the configuration implements two-way TLS authentication also known as mutual authentication.

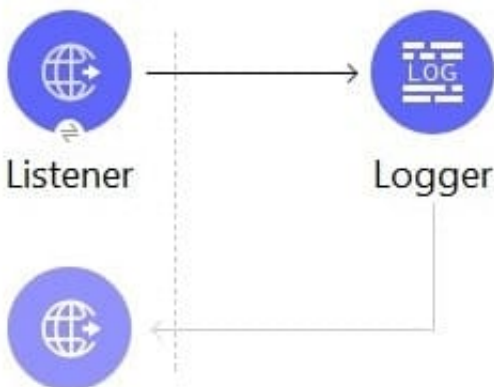
\*

in this case, correct answer is Support for TLS mutual (two-way) authentication with HTTP clients.



**QUESTION 13**

Refer to the exhibit.



**Error handling**

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

- A. CPU\_INTENSIVE and Dedicated Selector pool
- B. UBER and NONBLOCKING
- C. Shared Selector Pool and CPU LITE
- D. BLOCKING\_IO and UBER

Correct Answer: C

**QUESTION 14**



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

Correct 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

## QUESTION 15

A new Mule application under development must implement extensive data transformation logic. Some of the data transformation functionality is already available as external transformation services that are mature and widely used across the organization; the rest is highly specific to the new Mule application.

The organization follows a rigorous testing approach, where every service and application must be extensively acceptance tested before it is allowed to go into production.

What is the best way to implement the data transformation logic for this new Mule application while minimizing the overall testing effort?

- A. Implement and expose all transformation logic as microservices using DataWeave, so it can be reused by any application component that needs it, including the new Mule application
- B. Implement transformation logic in the new Mule application using DataWeave, replicating the transformation logic of existing transformation services
- C. Extend the existing transformation services with new transformation logic and Invoke them from the new Mule application
- D. Implement transformation logic in the new Mule application using DataWeave, invoking existing transformation services when possible

Correct Answer: D

Explanation: Correct answer is Implement transformation logic in the new Mule application using DataWeave, invoking existing transformation services when possible. \* The key here minimal testing effort, "Extend existing transformation logic" is not a feasible option because additional functionality is highly specific to the new Mule application so it should not be a part of commonly used functionality. So this option is ruled out. \* "Implement transformation logic in the new Mule application using DataWeave, replicating the transformation logic of existing transformation services" Replicating the transformation logic of existing transformation services will cause duplicity of code. So this option is ruled out. \* "Implement and expose all transformation logic as microservices using DataWeave, so it can be reused by any



application component that needs it, including the new Mule application" as question specifies that the transformation is app specific and wont be used outside

[Latest MCIA-LEVEL-1 Dumps](#)

[MCIA-LEVEL-1 VCE Dumps](#)

[MCIA-LEVEL-1 Practice Test](#)