

## Exam Questions MCI A-Level-1

MuleSoft Certified Integration Architect - Level 1

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



**NEW QUESTION 1**


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

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

Answer: B

**NEW QUESTION 2**

Refer to the exhibit.



An organization is designing a Mule application to receive data from one external business partner. The two companies currently have no shared IT infrastructure and do not want to establish one. Instead, all communication should be over the public internet (with no VPN).


What Anypoint Connector can be used in the organization's Mule application to securely receive data from this external business partner?

- A. File connector
- B. VM connector
- C. SFTP connector
- D. Object Store connector

Answer: C

**NEW QUESTION 3**

Refer to the exhibit.



An organization uses a 2-node Mule runtime cluster to host one stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution.

Two additional nodes have been added to the cluster and the load balancer has been configured to recognize the new nodes with no other change to the load balancer.

What average performance change is guaranteed to happen, assuming all cluster nodes are fully operational?

- A. 50% reduction in the response time of the API
- B. 100% increase in the throughput of the API
- C. 50% reduction in the JVM heap memory consumed by each node
- D. 50% reduction in the number of requests being received by each node

Answer: D

#### NEW QUESTION 4

What is true about the network connections when a Mule application uses a JMS connector to interact with a JMS provider (message broker)?

- A. The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider
- B. The AMQP protocol can be used by the JMS connector to portably establish connections to various types of JMS providers
- C. To receive messages into the Mule application, the JMS provider initiates a network connection to the JMS connector and pushes messages along this connection
- D. To complete sending a JMS message, the JMS connector must establish a network connection with the JMS message recipient

Answer: D

#### NEW QUESTION 5

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 log 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 6

An integration Mule application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mule application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system. Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system). The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed. The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages. What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

- A. An On Error scope Non-persistent VM ActiveMQ Dead Letter Queue for manual processing
- B. An On Error scope MuleSoft Object Store ActiveMQ Dead Letter Queue for manual processing
- C. Until Successful component MuleSoft Object Store ActiveMQ is NOT needed or used
- D. Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing

Answer: A

#### NEW QUESTION 7

An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to customhosted Mule runtimes (on-premises) while others execute in the MuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems. How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access the backend systems?

- A. Create a Mule domain project that maintains the credentials as Mule domain-shared resources Deploy the Mule applications to the Mule domain, so the credentials are available to the Mule applications
- B. Store the credentials in properties files in a shared folder within the organization's data center Have the Mule applications load properties files from this shared location at startup
- C. Segregate the credentials for each backend system into environment-specific properties files Package these properties files in each Mule application, from where they are loaded at startup
- D. Configure or create a credentials service that returns the credentials for each backend system, and that is accessible from customer-hosted and MuleSoft-hosted Mule runtimes Have the Mule applications load the properties at startup by invoking that credentials service

Answer: D

#### NEW QUESTION 8

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 9


An organization is designing an integration solution to replicate financial transaction data from a legacy system into a data warehouse (DWH). The DWH must contain a daily snapshot of financial transactions, to be delivered as a CSV file. Daily transaction volume exceeds tens of millions of records, with significant spikes in volume during popular shopping periods. What is the most appropriate integration style for an integration solution that meets the organization's current requirements?

- A. API-led connectivity
- B. Batch-triggered ETL
- C. Event-driven architecture
- D. Microservice architecture

Answer: D

**NEW QUESTION 10**

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. Use MUnit to simulate standard responses from the backend system Then conduct performance tests to identify other bottlenecks in the system
- B. 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
- C. Conduct scaled-down performance tests in the staging environment against the rate limited backend system Then upscale performance results to full production scale
- D. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situation, instead invoking local stubs that replicate typical backend system responses Then conduct performance tests using this API implementation

Answer: C

**NEW QUESTION 10**

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 11**

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

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

Answer: A

**NEW QUESTION 15**

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 queue OrderFulfilled messages are sent to a JMS topic Both 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 microservices OrderFulfilled messages are sent directly to the Order microservice The 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 queue OrderFulfilled messages are sent to a JMS topic The 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 18**

An organization uses a set of customer-hosted Mule runtimes that are managed using the Mulesoft-hosted control plane. What is a condition that can be alerted on from Anypoint Runtime Manager without any custom components or custom coding?

- A. When an SSL certificate used by one of the deployed Mule applications is about to expire
- B. When a Mule runtime on a given customer-hosted server is experiencing high memory consumption during certain periods
- C. When a Mule runtime's customer-hosted server is about to run out of disk space
- D. When the Mule runtime license installed on a Mule runtime is about to expire

**Answer:** A

**NEW QUESTION 23**

.....

## 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 First Try
- \* MCIA-Level-1 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year