

MuleSoft

Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1



NEW QUESTION 1

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 2

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 3

Refer to the exhibit.

A Mule application is deployed to a cluster of two customer-hosted Mule runtimes. The Mule application has a flow that polls a database and another flow with an HTTP Listener.

HTTP clients send HTTP requests directly to individual cluster nodes.

What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has failed, but before that node is restarted?

- A. Database polling continues Only HTTP requests sent to the remaining node continue to be accepted
- B. Database polling stops All HTTP requests continue to be accepted
- C. Database polling continues All HTTP requests continue to be accepted, but requests to the failed node incur increased latency
- D. Database polling stops All HTTP requests are rejected

Answer: A

NEW QUESTION 4

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 5

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 mlaoservices 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

Answer: D

NEW QUESTION 6

When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

- A. Identity management setup
- B. Environment setup
- C. Role and permission setup
- D. Dedicated Load Balancer setup

Answer: C

NEW QUESTION 7

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 8

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 9

Refer to the exhibit.

An organization is sizing an Anypoint VPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization currently has three non-production environments (DEV, SIT and UAT) that share this VPC. The AWS region of the VPC has two AZs.

The organization has a very mature DevOps approach which automatically progresses each application through all non-production environments before automatically deploying to production. This process results in several Mule application deployments per hour, using CloudHub's normal zero-downtime deployment feature.

What is a CIDR block for this VPC that results in the smallest usable private IP address range?

- A. 10.0.0.0/26 (64 IPs)
- B. 10.0.0.0/25 (128 IPs)
- C. 10.0.0.0/24 (256 IPs)
- D. 10.0.0.0/22 (1024 IPs)

Answer: D

NEW QUESTION 10

A set of integration Mule applications, some of which expose APIs, are being created to enable a new business process. Various stakeholders may be impacted by this. These stakeholders are a combination of

semi-technical users (who understand basic integration terminology and concepts such as JSON and XML) and technically skilled potential consumers of the Mule applications and APIs.

What is an effective way for the project team responsible for the Mule applications and APIs being built to communicate with these stakeholders using Anypoint

Platform and its supplied toolset?

- A. Use Anypoint Design Center to implement the Mule applications and APIs and give the various stakeholders access to these Design Center projects, so they can collaborate and provide feedback
- B. Create Anypoint Exchange entries with pages elaborating the integration design, including API notebooks (where applicable) to help the stakeholders understand and interact with the Mule applications and APIs at various levels of technical depth
- C. Use Anypoint Exchange to register the various Mule applications and APIs and share the RAML definitions with the stakeholders, so they can be discovered
- D. Capture documentation about the Mule applications and APIs inline within the Mule integration flows and use Anypoint Studio's Export Documentation feature to provide an HTML version of this documentation to the stakeholders

Answer: B

NEW QUESTION 10

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 exchangeOrderFulfilled messages are sent to an Anypoint MQ queueBoth 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 14

A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application.

What non-functional requirement can be expected to be negatively affected when implementing this change?

- A. Throughput
- B. Availability
- C. Response time
- D. Consistency

Answer: D

NEW QUESTION 18

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. JSON Schema
- B. RAML
- C. WSDL
- D. XSD

Answer: C

NEW QUESTION 20

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 componentOn Error Continue scope
- B. VM connectorFirst Successful scopeOn Error Propagate scope
- C. Async scopeOn Error Propagate scope
- D. For Each scopeOn Error Continue scope

Answer: D

NEW QUESTION 23

.....

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

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