

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 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

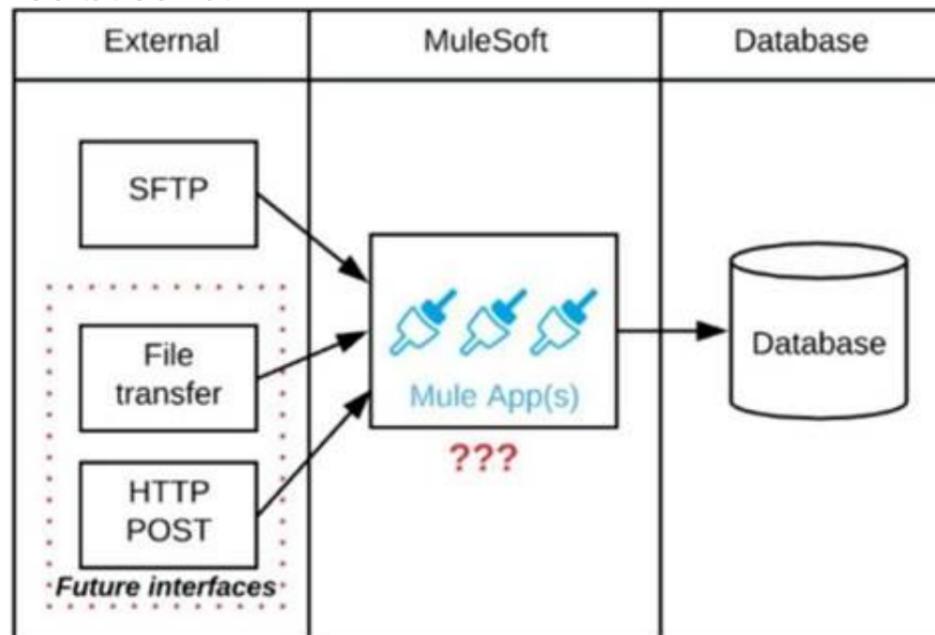
An XA transaction is being configured that involves a JMS connector listening for incoming JMS messages. What is the meaning of the timeout attribute of the XA transaction, and what happens after the timeout expires?

- A. The time that is allowed to pass between committing the transaction and the completion of the Mule flow. After the timeout, flow processing triggers an error.
- B. The time that is allowed to pass between receiving JMS messages on the same JMS connection. After the timeout, a new JMS connection is established.
- C. The time that is allowed to pass without the transaction being ended explicitly. After the timeout, the transaction is forcefully rolled-back.
- D. The time that is allowed to pass for state JMS consumer threads to be destroyed. After the timeout, a new JMS consumer thread is created.

Answer: C

NEW QUESTION 3

Refer to the exhibit.



A business process involves the receipt of a file from an external vendor over SFTP. The file needs to be parsed and its content processed, validated, and ultimately persisted to a database. The delivery mechanism is expected to change in the future as more vendors send similar files using other mechanisms such as file transfer or HTTP POST.

What is the most effective way to design for these requirements in order to minimize the impact of future change?

- A. Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- B. Create a Process API to receive the file and process it using a MuleSoft Batch Job while delegating the data save process to a System API
- C. Create an API that receives the file and invokes a Process API with the data contained in the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed
- D. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing

Answer: C

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.

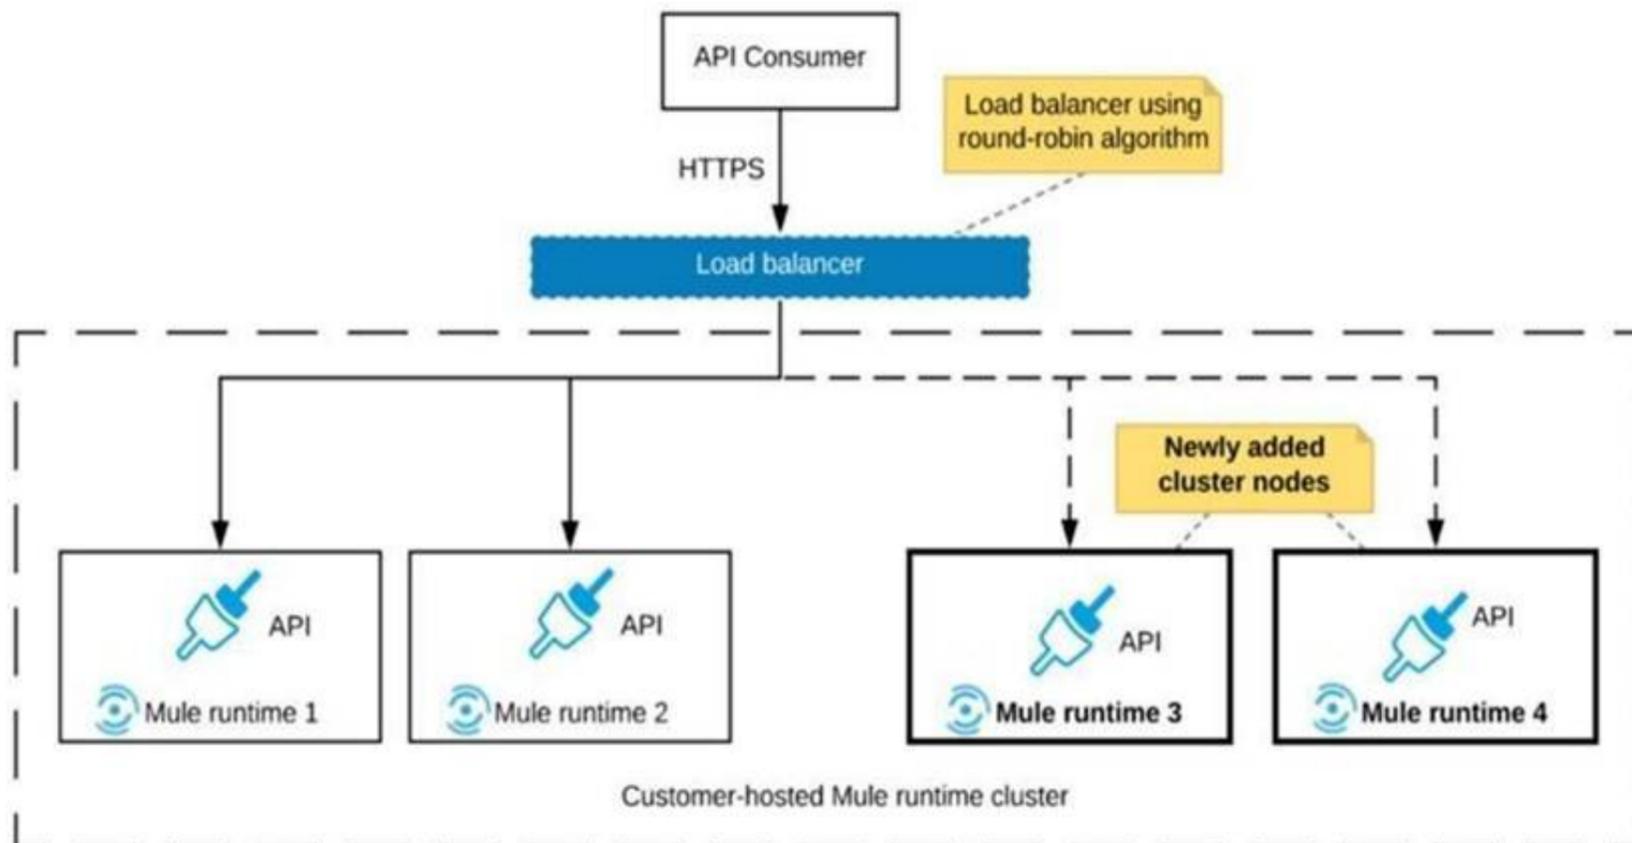
What key elements should be included in the integration design specification when requesting approval for this Mule application?

- 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

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 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 code for 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

An organization currently uses a multi-node Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications. The organization is planning to transition to a deployment model based on Docker containers in a Kubernetes cluster. The organization has already created a standard Docker image containing a Mule runtime and all required dependencies (including a JVM), but excluding the Mule application itself.

What is an expected outcome of this transition to container-based Mule application deployments?

- A. Required redesign of Mule applications to follow microservice architecture principles
- B. Required migration to the Docker and Kubernetes-based Anypoint Platform - Private Cloud Edition
- C. Required change to the URL endpoints used by clients to send requests to the Mule applications
- D. Guaranteed consistency of execution environments across all deployments of a Mule application

Answer: A

NEW QUESTION 8

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 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

49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58. <https://api.acme.com/partnefs/v1>. The API implementation is estimated to require deployment to 5 CloudHub workers. The retailer has obtained a public X.509 certificate for the name `apl.acme.com`, signed by a reputable CA, to be used as the server certificate. Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

- A. Add the X.509 certificate to the Mule application's deployable archive, then configure a CloudHub Dedicated Load Balancer (DLB) for each of the Mule application's CloudHub workers Create a CNAME for `api.acme.com` pointing to the DLB's A record
- B. Add the X.509 certificate to the CloudHub Shared Load Balancer (SLB), not to the Mule application Create a CNAME for `api.acme.com` pointing to the SLB's A record
- C. Add the X.509 certificate to a CloudHub Dedicated Load Balancer (DLB), not to the Mule application Create a CNAME for `api.acme.com` pointing to the DLB's A record
- D. Add the x.509 certificate to the Mule application's deployable archive, then configure the CloudHub Shared Load Balancer (SLB) for each of the Mule application's CloudHub workers Create a CNAME for `api.acme.com` pointing to the SLB's A record

Answer: C

NEW QUESTION 10

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 13

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 16

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 Mute 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 20

What Anypoint Connectors support transactions?

- A. Database, JMS, VM
- B. Database, 3MS, HTTP
- C. Database, JMS, VM, SFTP
- D. Database, VM, File

Answer: A

NEW QUESTION 23

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 26

What limits if a particular Anypoint Platform user can discover an asset in Anypoint Exchange?

- A. The type of the asset in Anypoint Exchange
- B. The business groups to which the user belongs
- C. If Design Center and RAML were both used to create the asset
- D. The existence of a public Anypoint Exchange portal to which the asset has been published

Answer: A

NEW QUESTION 30

A retailer is designing a data exchange interface to be used by its suppliers. The interface must support secure communication over the public internet. The interface must also work with a wide variety of programming languages and IT systems used by suppliers.

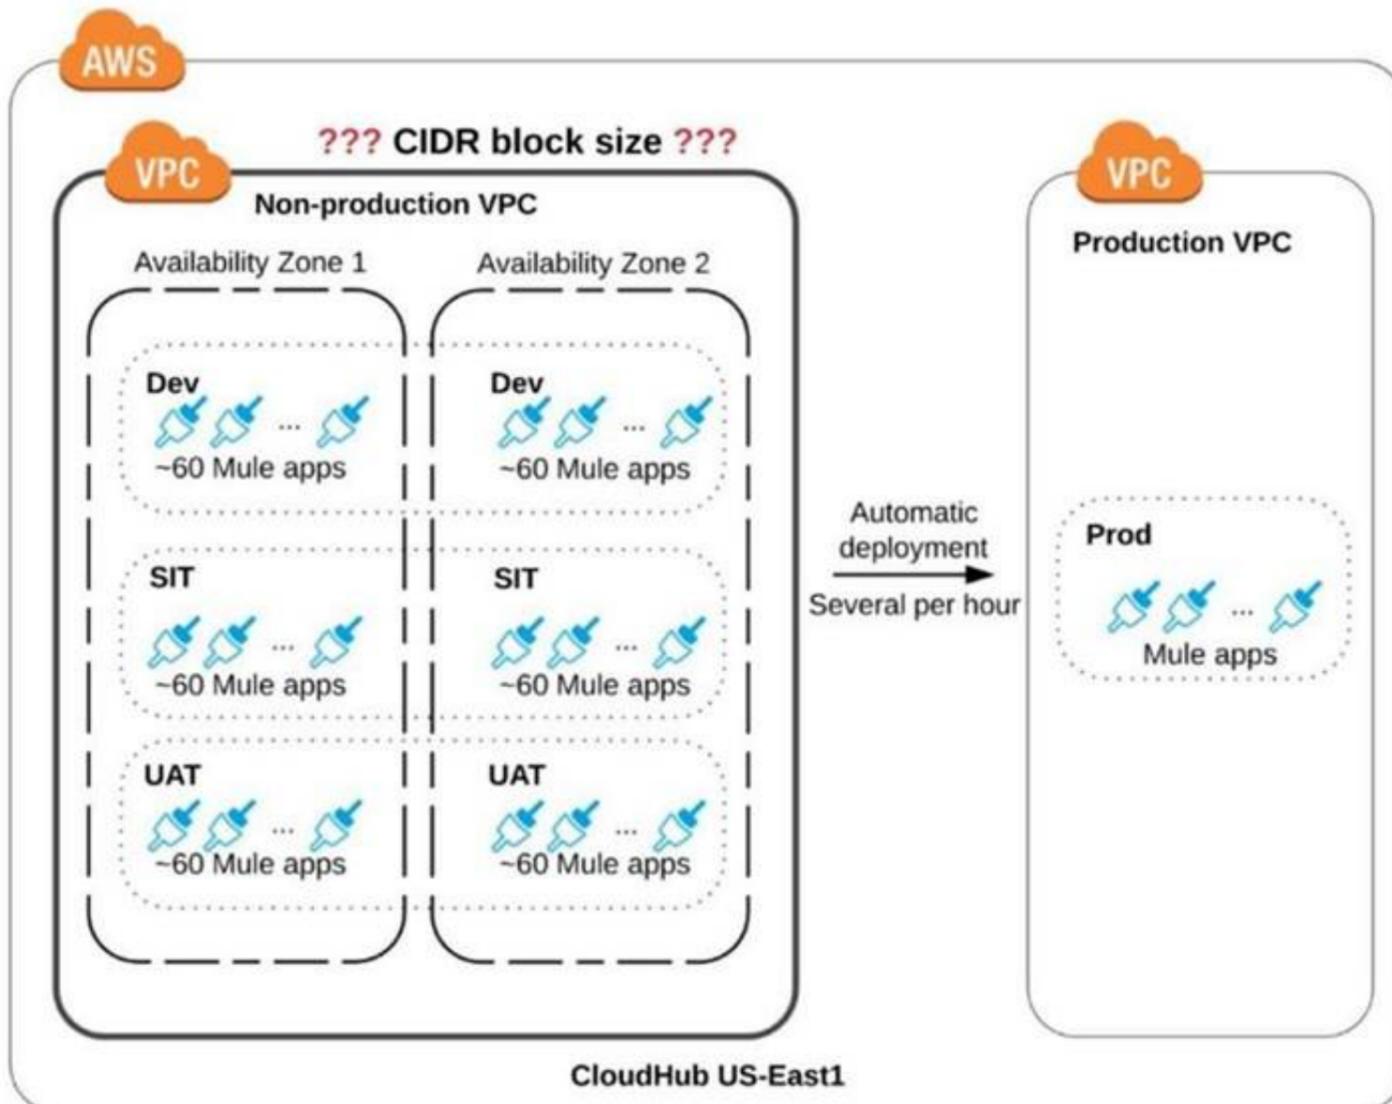
What are suitable interface technologies for this data exchange that are secure, cross-platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

- A. EDJFACT XML over SFTP JSON/REST over HTTPS
- B. SOAP over HTTPS HOP over TLS gRPC over HTTPS
- C. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
- D. CSV over FTP YAML over TLS JSON over HTTPS

Answer: B

NEW QUESTION 31

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 36

Additional nodes are being added to an existing customer-hosted Mule runtime cluster to improve performance. Mule applications deployed to this cluster are invoked by API clients through a load balancer.

What is also required to carry out this change?

- A. A new load balancer must be provisioned to allow traffic to the new nodes in a round-robin fashion
- B. External monitoring tools or log aggregators must be configured to recognize the new nodes
- C. API implementations using an object store must be adjusted to recognize the new nodes and persist to them
- D. New firewall rules must be configured to accommodate communication between API clients and the new nodes

Answer: C

NEW QUESTION 38

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 39

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.

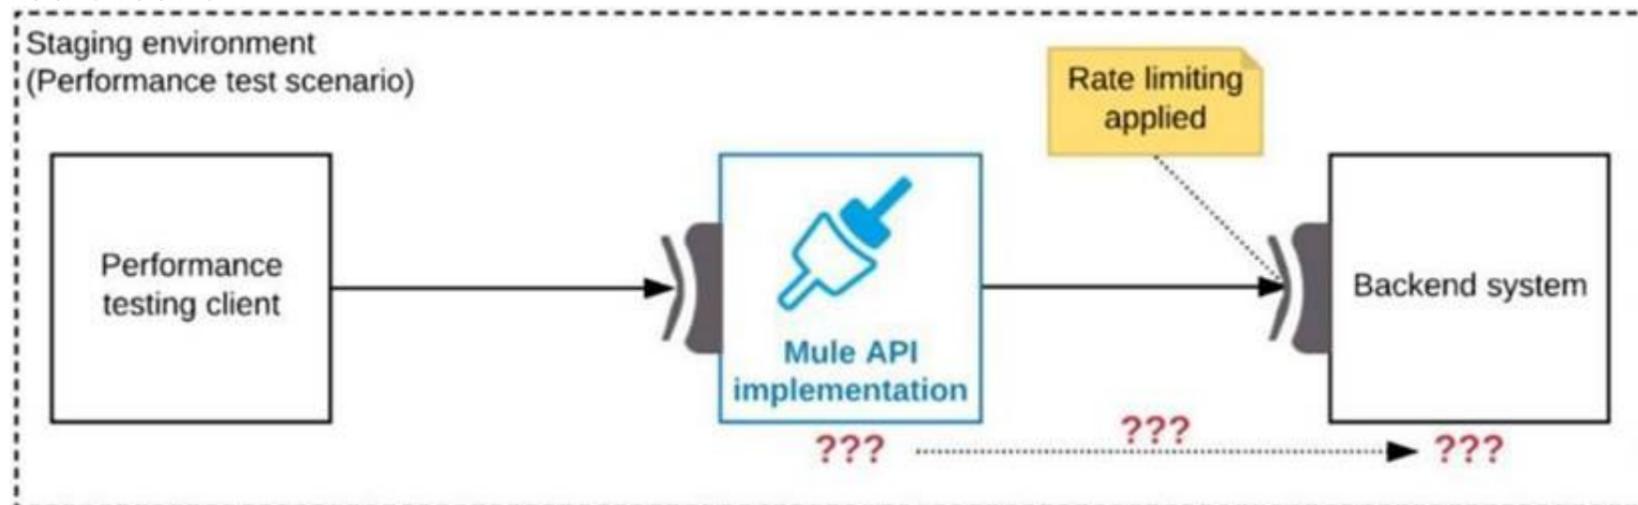
What is the most appropriate integration style for an integration solution that meets the organization's current requirements?

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

Answer: D

NEW QUESTION 43

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 48

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 52

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 56

An organization is designing the following two Mule applications that must share data via a common persistent object store instance:

- Mule application P will be deployed within their on-premises datacenter. - Mule application C will run on CloudHub in an Anypoint VPC.

The object store implementation used by CloudHub is the Anypoint Object Store v2 (OSv2).

What type of object store(s) should be used, and what design gives both Mule applications access to the same object store instance?

- A. Application C and P both use the Object Store connector to access the Anypoint Object Store v2
- B. Application C and P both use the Object Store connector to access a persistent object store
- C. Application C uses the Object Store connector to access a persistent object Application P accesses the persistent object store via the Object Store REST API
- D. Application P uses the Object Store connector to access a persistent object store Application C accesses this persistent object store via the Object Store REST API through an IPsec tunnel

Answer: A

NEW QUESTION 59

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 62

.....

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