



# **MuleSoft**

## **Exam Questions MCIA-Level-1**

MuleSoft Certified Integration Architect - Level 1

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

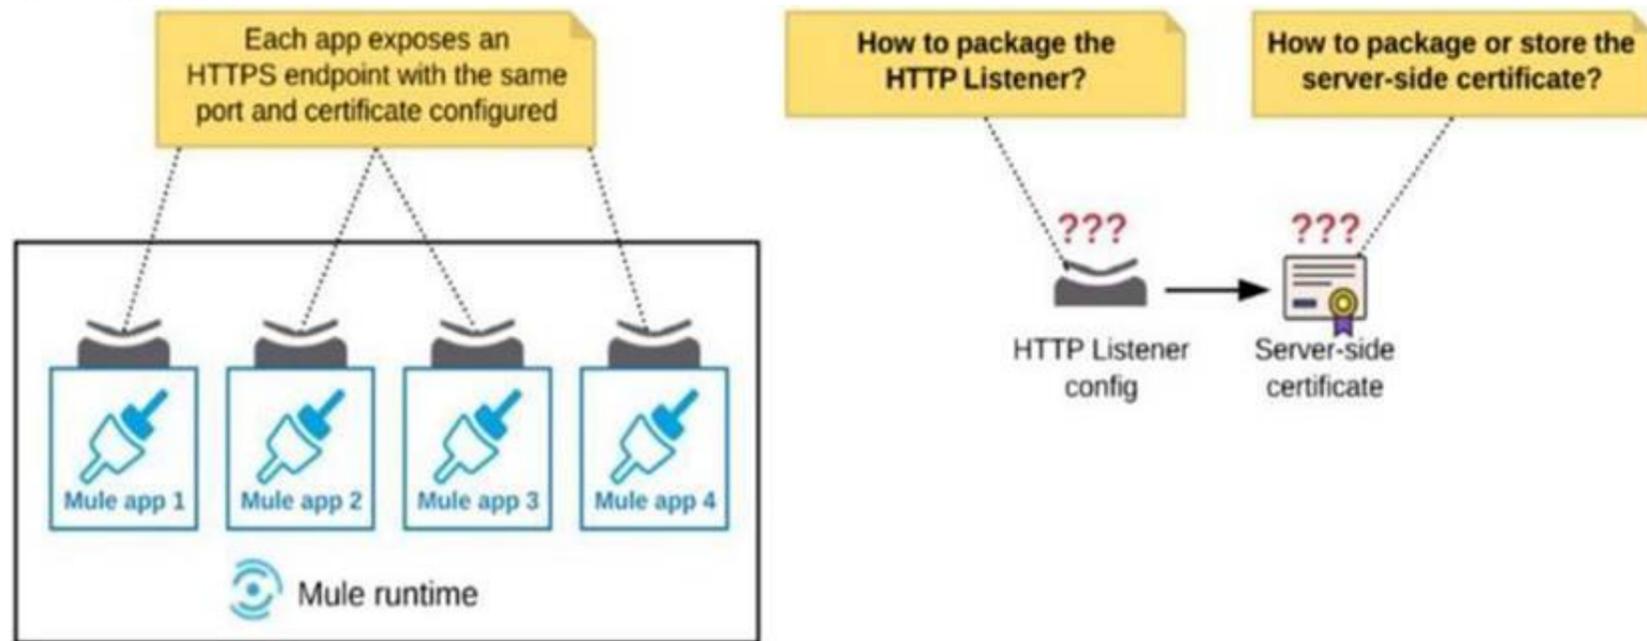
Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

**NEW QUESTION 1**

Refer to the exhibit.



An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a serverside certificate that rotates often.

What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

- A. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Muleapplications that need to expose an HTTPS endpoint Package the server-side certificate in ALL Mule APPLICATIONS that need to expose an HTTPS endpoint
- B. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Muleapplications that need to expose an HTTPS endpointStore the server-side certificate in a shared filesystem location in the Mule runtime's classpath,OUTSIDE the Mule DOMAIN or any Mule APPLICATION
- C. Package an HTTPS Listener configuration In all Mule APPLICATIONS that need to expose an HTTPSendpoint Package the server-side certificate in a NEW Mule DOMAIN project
- D. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing It from all Muleapplications that need to expose an HTTPS endpoint Package the server-side certificate in the SAME Mule DOMAIN project Go to Set

**Answer: B**

**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**

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 4**

An integration Mute 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 falls, then a log entry must be written for that row, but processing of other rows must not be affected.

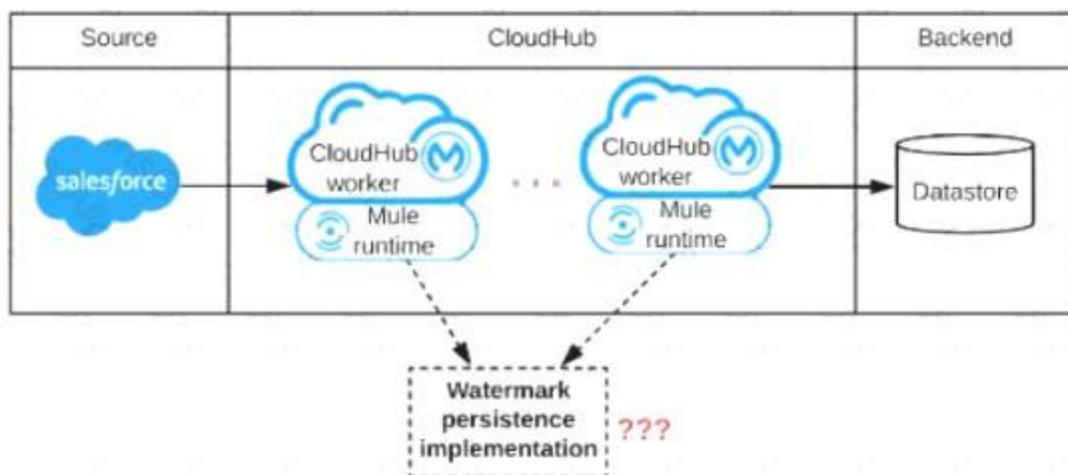
What combination of Mute components is most idiomatic (used according to their intended purpose) when Implementing the above requirements?

- 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

**Answer: C**

**NEW QUESTION 5**

Refer to the exhibit.



A Mule application is being designed to be deployed to several CloudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.

A watermark will be used to only retrieve those Satesforce Accounts that have been modified since the last time the integration logic ran.

What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?

- A. Persistent Anypoint MQ Queue
- B. Persistent Object Store
- C. Persistent Cache Scope
- D. Persistent VM Queue

**Answer: B**

**NEW QUESTION 6**

What Is a recommended practice when designing an integration Mule 4 application that reads a large XML payload as a stream?

- A. The payload should be dealt with as a repeatable XML stream, which must only be traversed (iterated-over) once and CANNOT be accessed randomly from DataWeave expressions and scripts
- B. The payload should be dealt with as an XML stream, without converting it to a single Java object (POJO)
- C. The payload size should NOT exceed the maximum available heap memory of the Mute runtime on which the Mule application executes
- D. The payload must be cached using a Cache scope If It Is to be sent to multiple backend systems

**Answer: B**

**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**

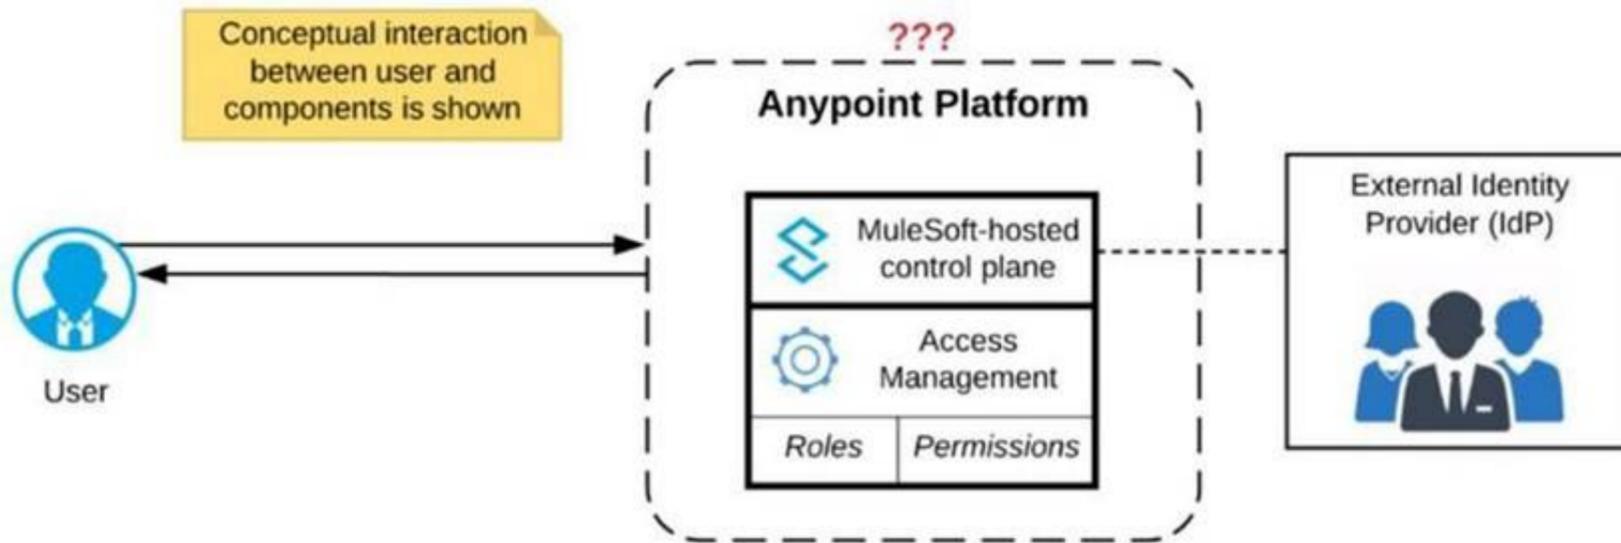
An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform. What feature of Anypoint Platform should be used to meet this requirement?

- A. Enterprise Security Module
- B. Client ID Secret
- C. Federated Identity Management
- D. Federated Client Management

**Answer: C**

**NEW QUESTION 9**

Refer to the exhibit.



Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.

What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled via the external Identity Provider?

- A. Controlling the business group within Anypoint Platform to which the user belongs
- B. Assigning Anypoint Platform permissions to a role
- C. Assigning Anypoint Platform role(s) to a user
- D. Removing a user's access to Anypoint Platform when they no longer work for the organization

**Answer: B**

**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/partners/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 [api.acme.com](https://api.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
- B. Add the X.509 certificate to the CloudHub Shared Load Balancer (SLB), not to the Mule application
- C. Add the X.509 certificate to a CloudHub Dedicated Load Balancer (DLB), not to the Mule application
- 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

**Answer: C**

**NEW QUESTION 10**

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 12**

A Mule application is being designed to do the following:

Step 1: Read a SalesOrder message from a JMS queue, where each SalesOrder consists of a header and a list of SalesOrderLineItems.

Step 2: Insert the SalesOrder header and each SalesOrderLineItem into different tables in an RDBMS.

Step 3: Insert the SalesOrder header and the sum of the prices of all its SalesOrderLineItems into a table in a different RDBMS.

No SalesOrder message can be lost and the consistency of all SalesOrder-related information in both RDBMSs must be ensured at all times.

What design choice (including choice of transactions) and order of steps addresses these requirements?

- A. \* 1. Read the JMS message (NOT in an XA transaction)\* 2. Perform EACH DB insert in a SEPARATE DB transaction\* 3. Acknowledge the JMS message
- B. \* 1. Read and acknowledge the JMS message (NOT in an XA transaction)\* 2. In a NEW XA transaction, perform BOTH DB inserts
- C. \* 1. Read the JMS message in an XA transaction\* 2. In the SAME XA transaction, perform BOTH DB inserts but do NOT acknowledge the JMS message
- D. \* 1. Read the JMS message (NOT in an XA transaction)\* 2. Perform BOTH DB inserts in ONE DB transaction\* 3. Acknowledge the JMS message

**Answer: C**

**NEW QUESTION 13**

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 has been updated and hence the database driver library/JAR needs a minor version upgrade

- C. The database server was unavailable for four hours due to a major outage but is now fully operational again
- D. The credentials for accessing the database have been updated and the previous credentials are no longer valid

**Answer:** D

#### NEW QUESTION 14

What operation can be performed through a JMX agent enabled in a Mule application?

- A. View object store entries
- B. Replay an unsuccessful message
- C. Deploy a Mule application
- D. Set a particular log4j2 log level to TRACE

**Answer:** D

#### NEW QUESTION 18

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 Mule 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 22

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 26

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 theMuleSoft-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 thebackend 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 toad the properties at startup by invoking that credentials service

**Answer:** D

#### NEW QUESTION 29

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 30

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 33**

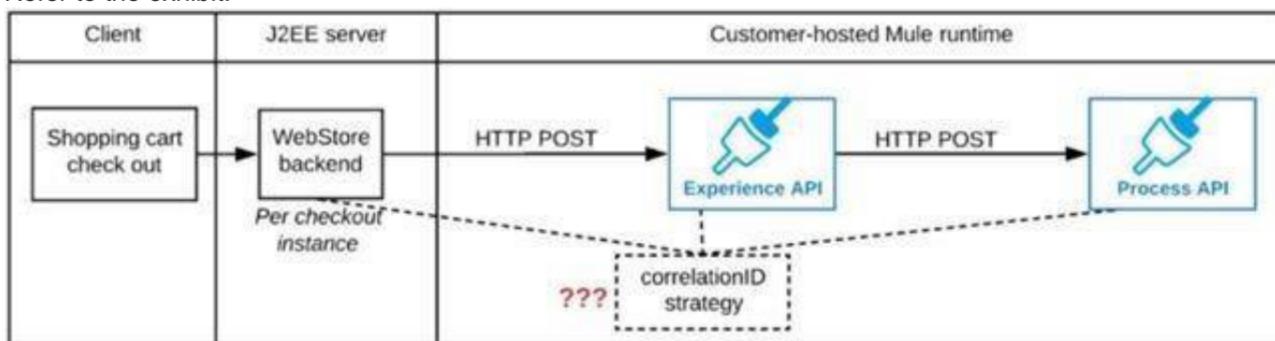
An organization is designing an integration solution to replicate financial transaction data from a legacy system into a data warehouse (DWH). The DWH must contain a daily snapshot of financial transactions, to be delivered as a CSV file. Daily transaction volume exceeds tens of millions of records, 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 37**

Refer to the exhibit.

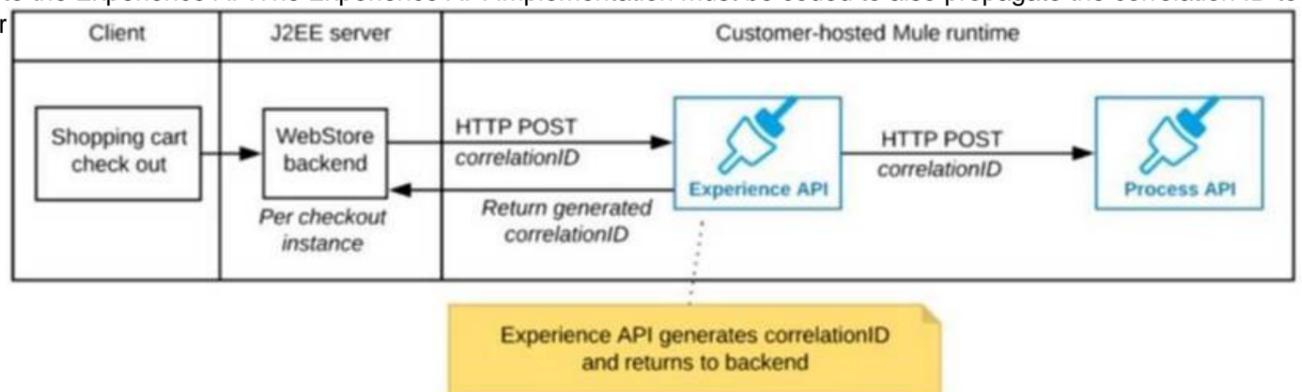


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations are Mule applications executing in a customer-hosted Mule runtime.

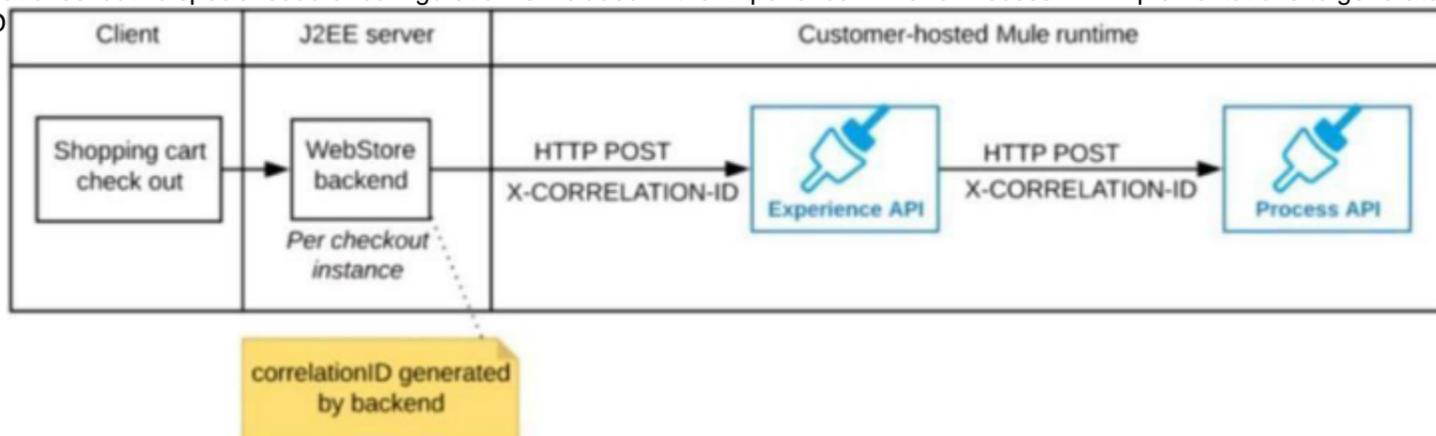
End-to-end correlation of all HTTP requests and responses belonging to each individual checkout instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the 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?

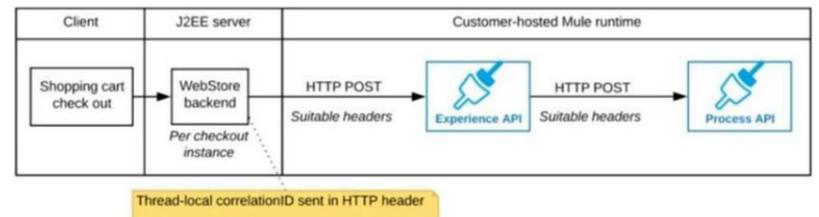
A. The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API. The Experience API implementation must be coded to also propagate the correlation ID to the Process API in a suitable HTTP request header.



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 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.

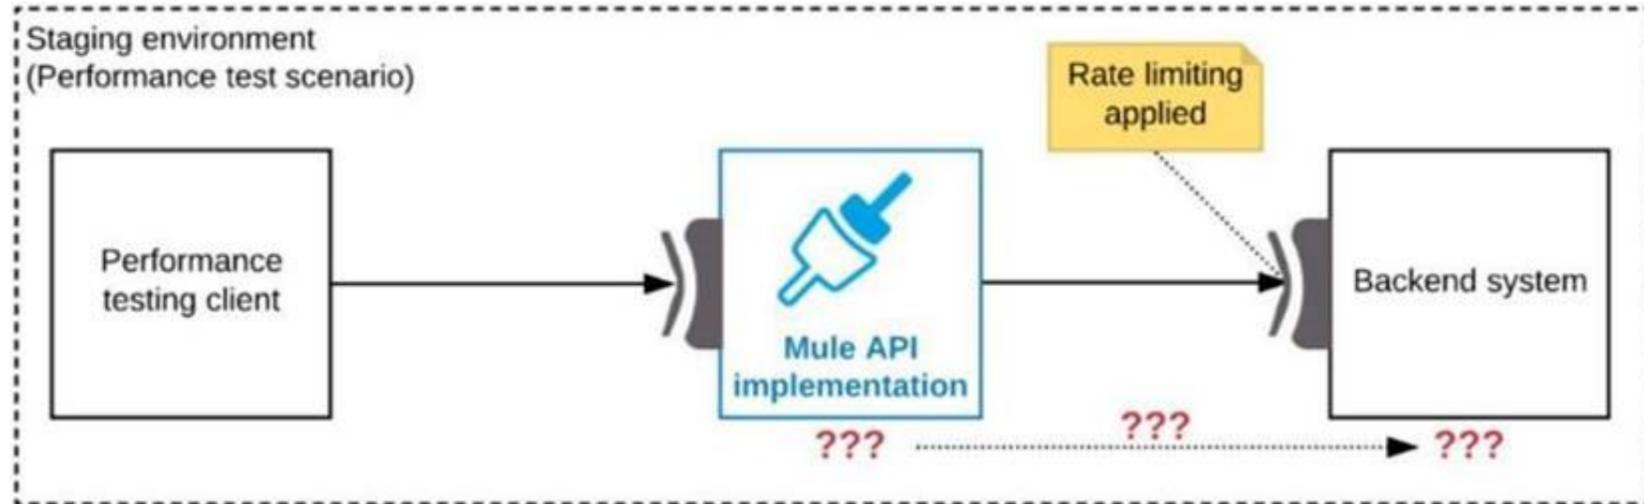


D. The web store backend sends a correlation ID value in the HTTP request body in the way required by the Experience API. The Experience API and Process API implementations must be coded to receive the custom correlation ID in the HTTP requests and propagate it in suitable HTTP request headers.

Answer: B

**NEW QUESTION 40**

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 45**

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 47**

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 48**

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 51

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 objectApplication 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 storeApplication C accesses this persistent object store via the Object Store REST API through an IPsec tunnel

**Answer: A**

#### NEW QUESTION 52

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 57

.....

## Relate Links

**100% Pass Your MCIA-Level-1 Exam with Exam Bible Prep Materials**

<https://www.exambible.com/MCIA-Level-1-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>