

MCIA-Level-1 Dumps

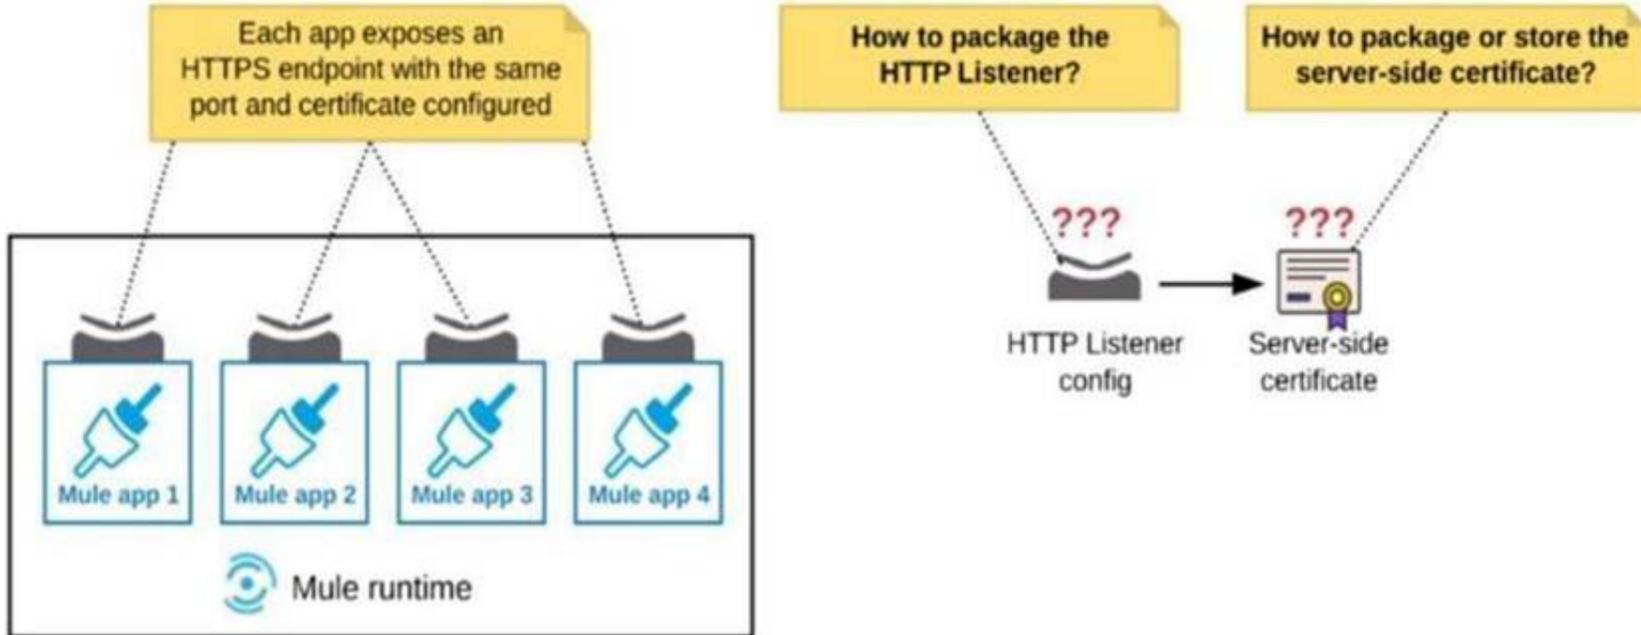
MuleSoft Certified Integration Architect - Level 1

<https://www.certleader.com/MCIA-Level-1-dumps.html>



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

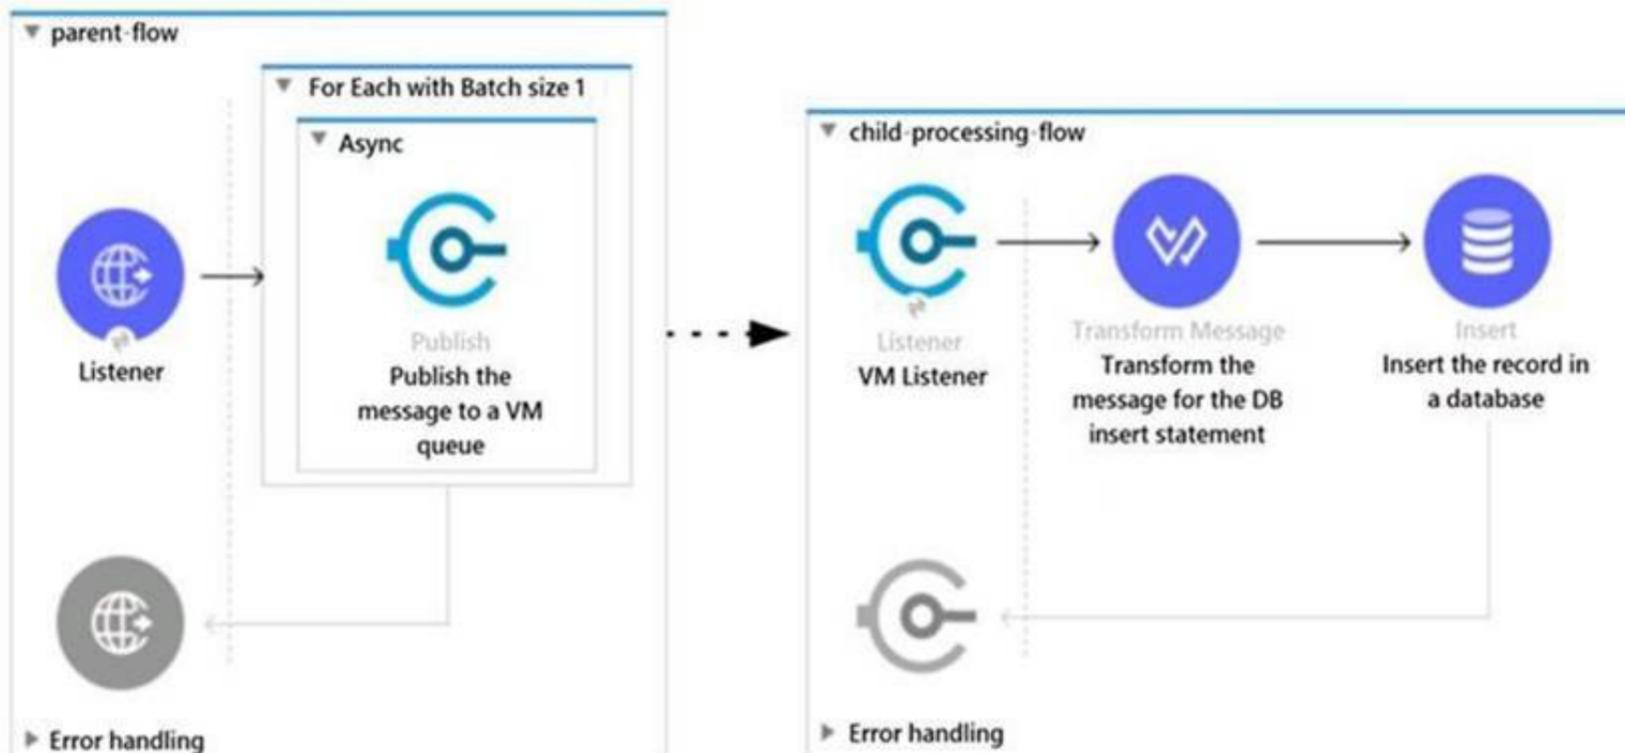
An XA transaction Is being configured that involves a JMS connector listening for Incoming JMS messages. What is the meaning of the timeout attribute of the XA transaction, and what happens after the timeout expires?

- 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 connectionAfter 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 thetimeout, a new JMS consumer thread is created

Answer: C

NEW QUESTION 3

Refer to the exhibit.



A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items,

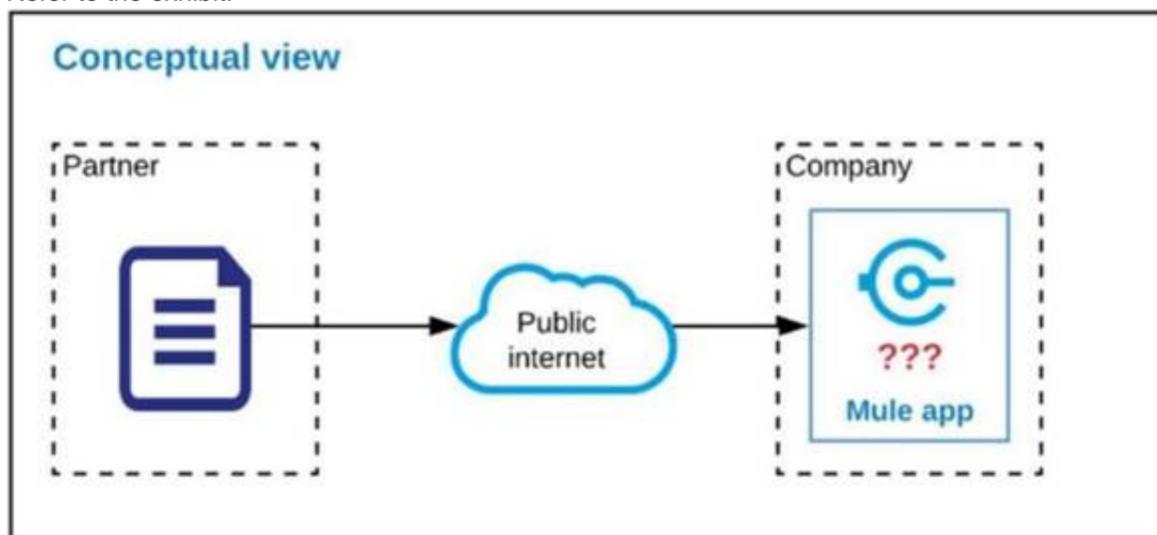
then sends each item one at a time inside an Async scope to a VM queue.
A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing the item to a database.
This Mule application is deployed to four CloudHub workers with persistent queues enabled.
What message processing guarantees are provided by the VM queue and the CloudHub workers, and how are VM messages routed among the CloudHub workers for each invocation of the parent flow under normal operating conditions where all the CloudHub workers remain online?

- A. EACH item VM message is processed AT MOST ONCE by ONE CloudHub worker, with workers chosen in a deterministic round-robin fashion Each of the four CloudHub workers can be expected to process 1/4 of the Item VM messages (about 50 items)
- B. EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker Each of the four CloudHub workers can be expected to process some item VM messages
- C. ALL Item VM messages are processed AT LEAST ONCE by the SAME CloudHub worker where the parent flow was invoked This one CloudHub worker processes ALL 200 item VM messages
- D. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub worker This one CloudHub worker processes ALL 200 item VM messages

Answer: B

NEW QUESTION 4

Refer to the exhibit.



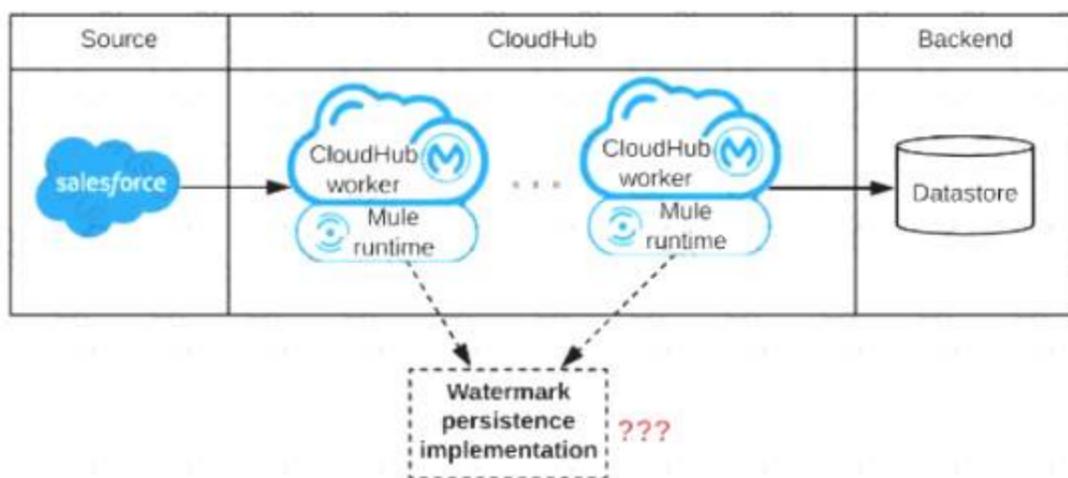
An organization is designing a Mule application to receive data from one external business partner. The two companies currently have no shared IT infrastructure and do not want to establish one. Instead, all communication should be over the public internet (with no VPN).
What Anypoint Connector can be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

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.

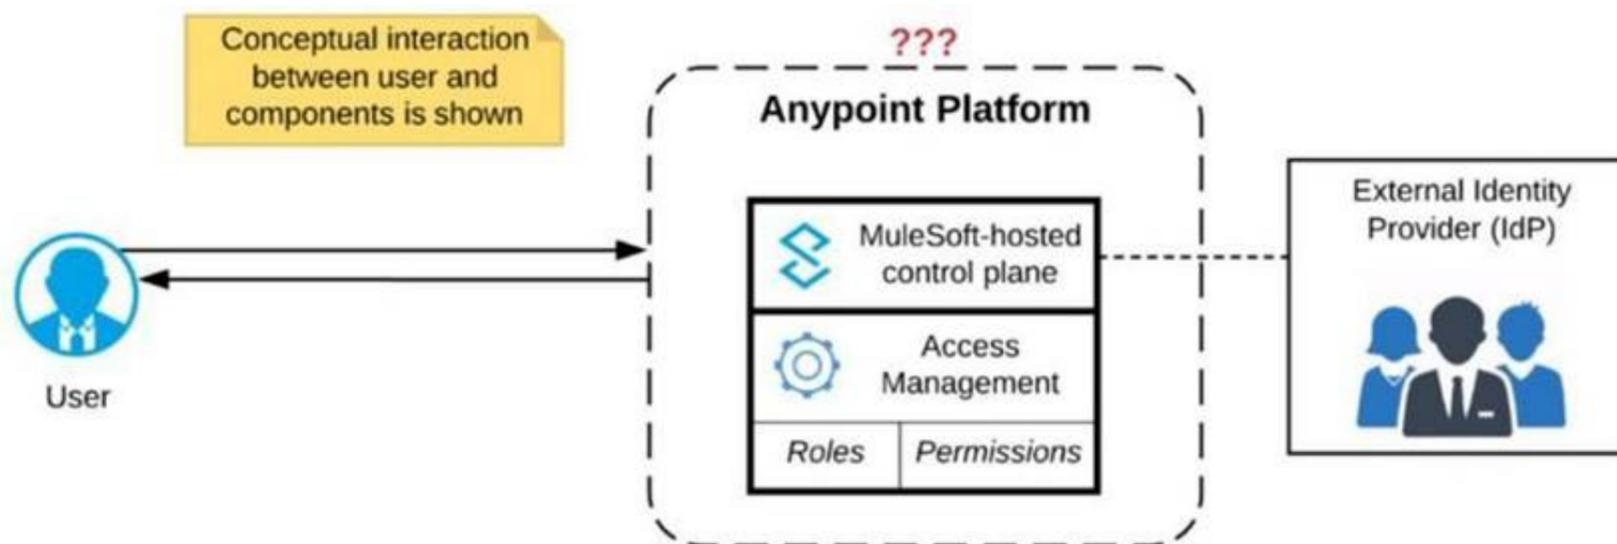
What combination of deployment options of the Anypoint Platform control plane and runtime plane(s) best serves this organization at the start of this strategic journey?

- 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 7

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 8

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 9

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 10

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 10

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 12

An integration Mule application is deployed to a customer-hosted multi-node Mule 4 runtime duster. The Mule application uses a Listener operation of a JMS connector to receive incoming messages from a JMS queue.

How are the messages consumed by the Mule application?

- A. Depending on the JMS provider's configuration, either all messages are consumed by ONLY the primary cluster node or else ALL messages are consumed by ALL cluster nodes
- B. Regardless of the Listener operation configuration, all messages are consumed by ALL cluster nodes
- C. Depending on the Listener operation configuration, either all messages are consumed by ONLY the primary cluster node or else EACH message is consumed by ANY ONE cluster node
- D. Regardless of the Listener operation configuration, all messages are consumed by ONLY the primary cluster node

Answer: C

NEW QUESTION 17

An integration Mute application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mute 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 componentMuleSoft 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 20

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.

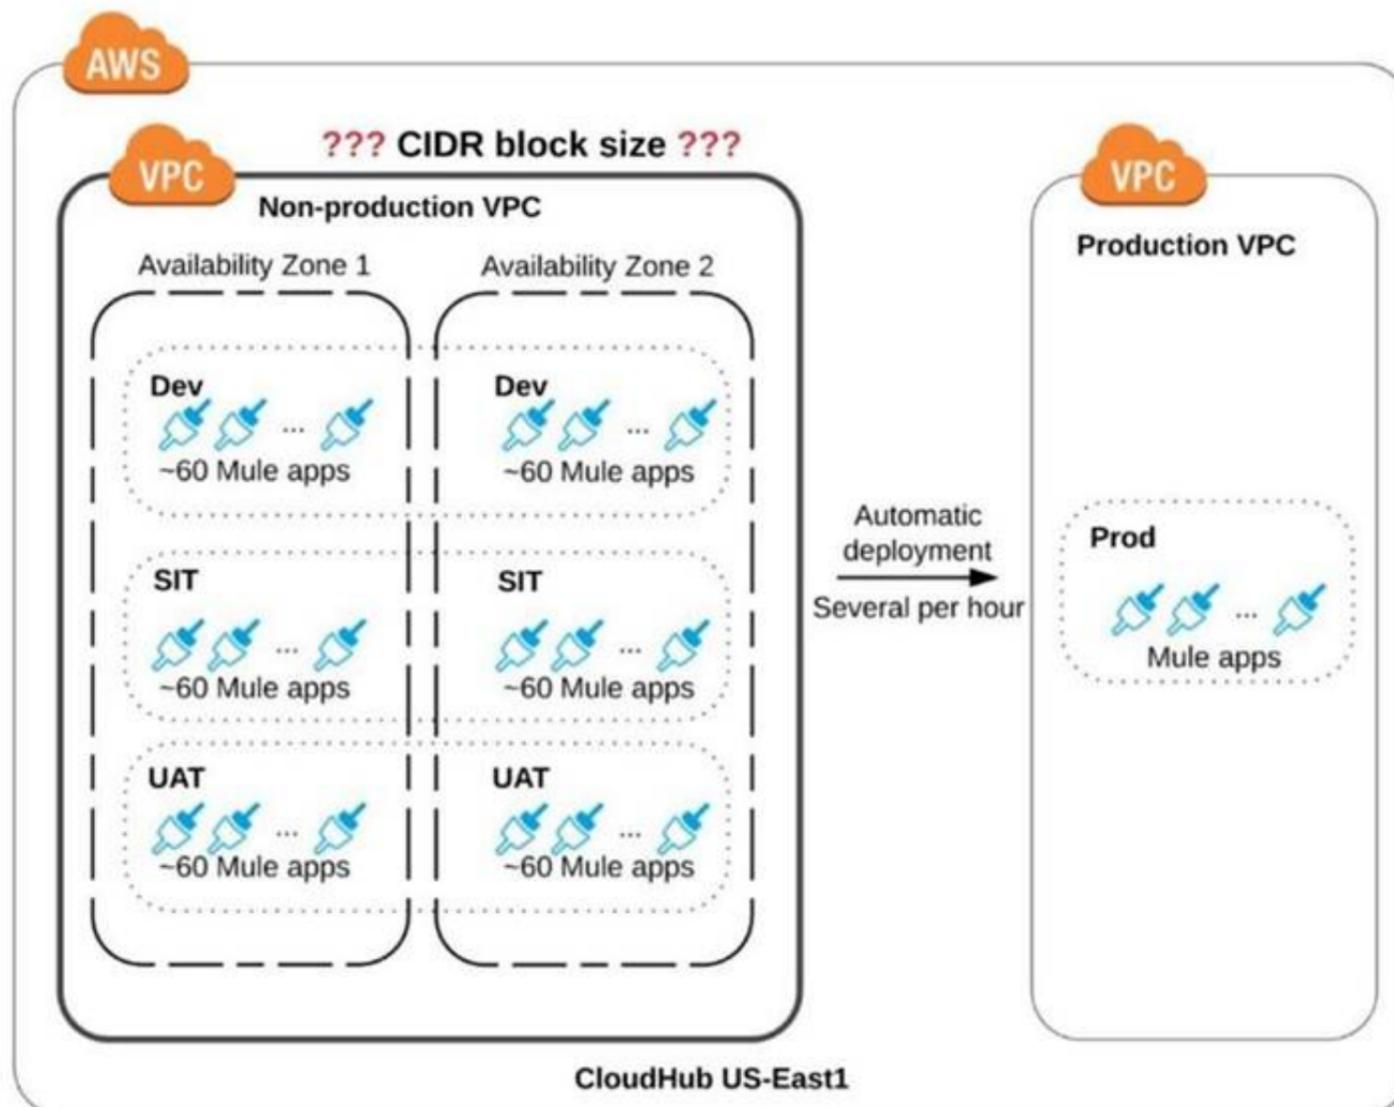
What are suitable interface technologies for this data exchange that are secure, cross-platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

- 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 21

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.

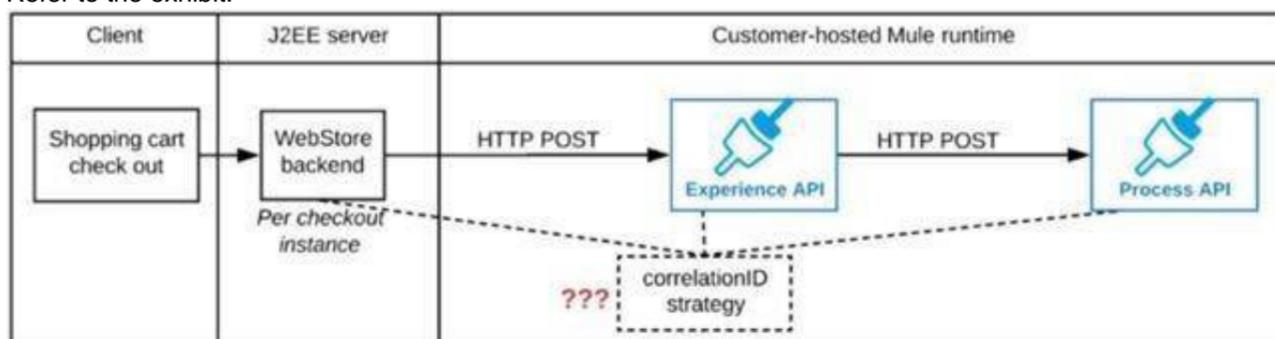
What is a CIDR block for this VPC that results in the smallest usable private IP address range?

- 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 23

Refer to the exhibit.

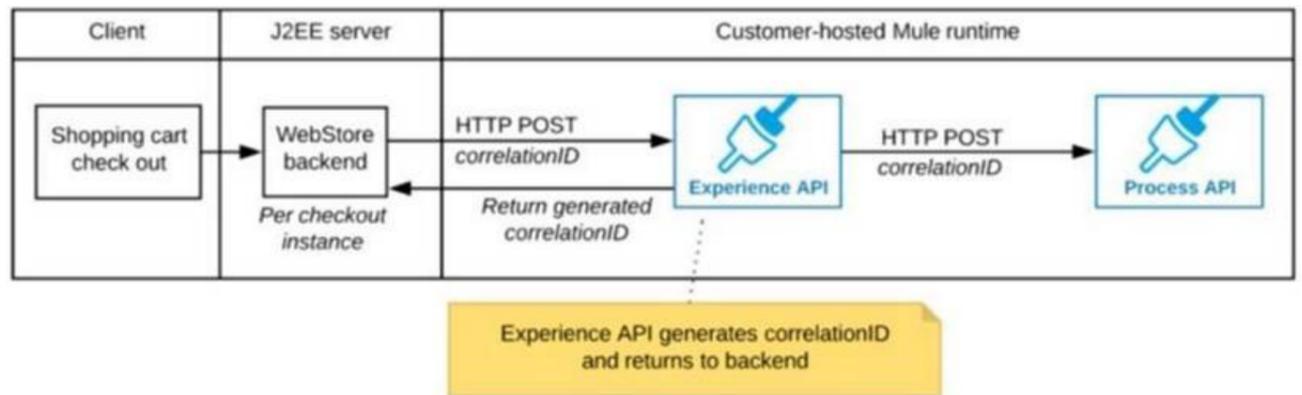


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

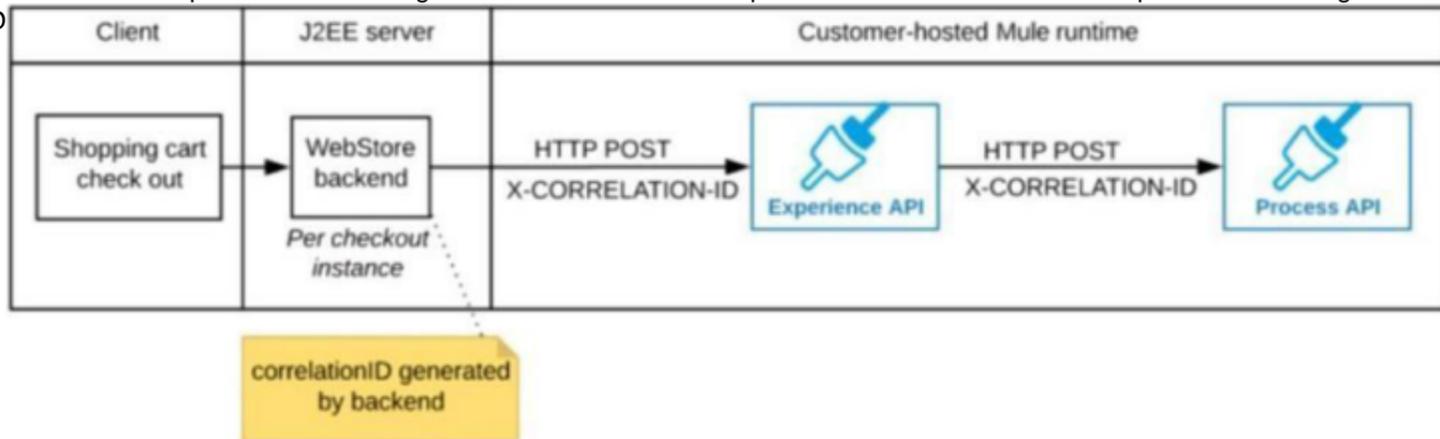
End-to-end correlation of all HTTP requests and responses belonging to each individual checkout instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the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

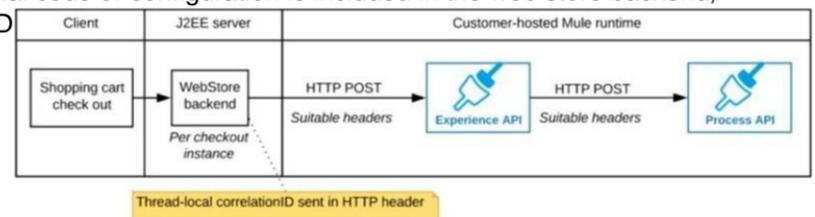
- A. The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API. The Experience API implementation must be coded to also propagate the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

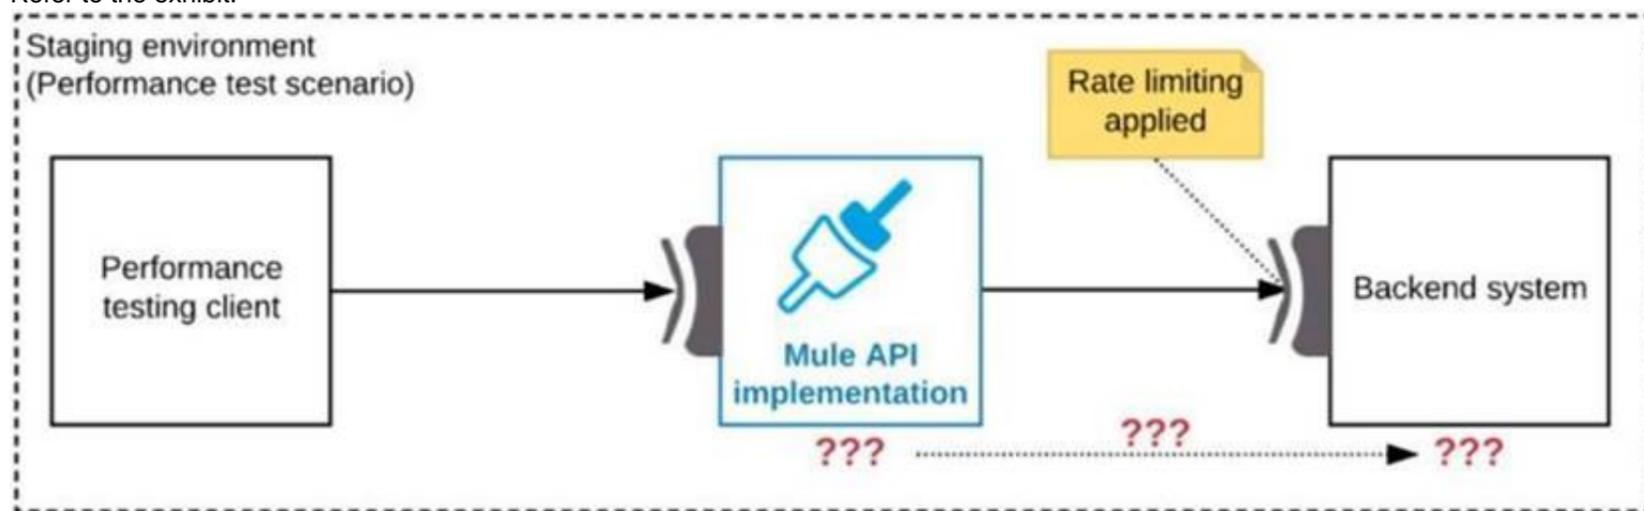


D. The web store backend sends a correlation ID value in the HTTP request body in the way required by the Experience API. The Experience API and Process API implementations must be coded to receive the custom correlation ID in the HTTP requests and propagate it in suitable HTTP request headers.

Answer: B

NEW QUESTION 25

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 30

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 StepsEach 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 StepsEach 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 StepsEach 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 recordEach Batch Step instance receives A BLOCK OF records as the payload, and BLOCKS of records are processed IN ORDER

Answer: A

NEW QUESTION 33

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 36

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 41

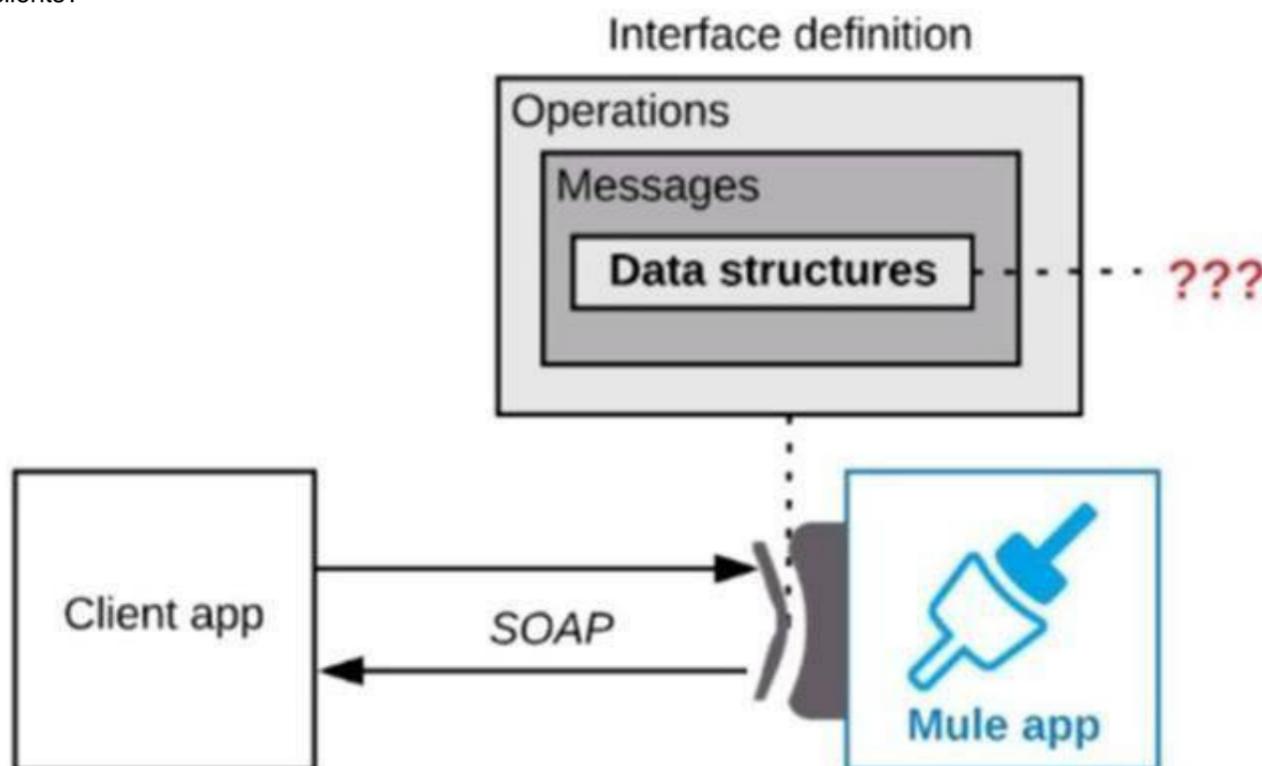
A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application. What non-functional requirement can be expected to be negatively affected when implementing this change?

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

Answer: D

NEW QUESTION 46

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 50

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 both HTTP request and HTTP response bodies for all HTTP clients
- C. Encryption of requests to both subdomains and API resource endpoints (<https://api.customer.com/> and <https://customer.com/api>)
- D. Encryption of both HTTP request header and HTTP request body for all HTTP clients

Answer: A

NEW QUESTION 52

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your MCIA-Level-1 Exam with Our Prep Materials Via below:

<https://www.certleader.com/MCIA-Level-1-dumps.html>