

## MCIA-Level-1 Dumps

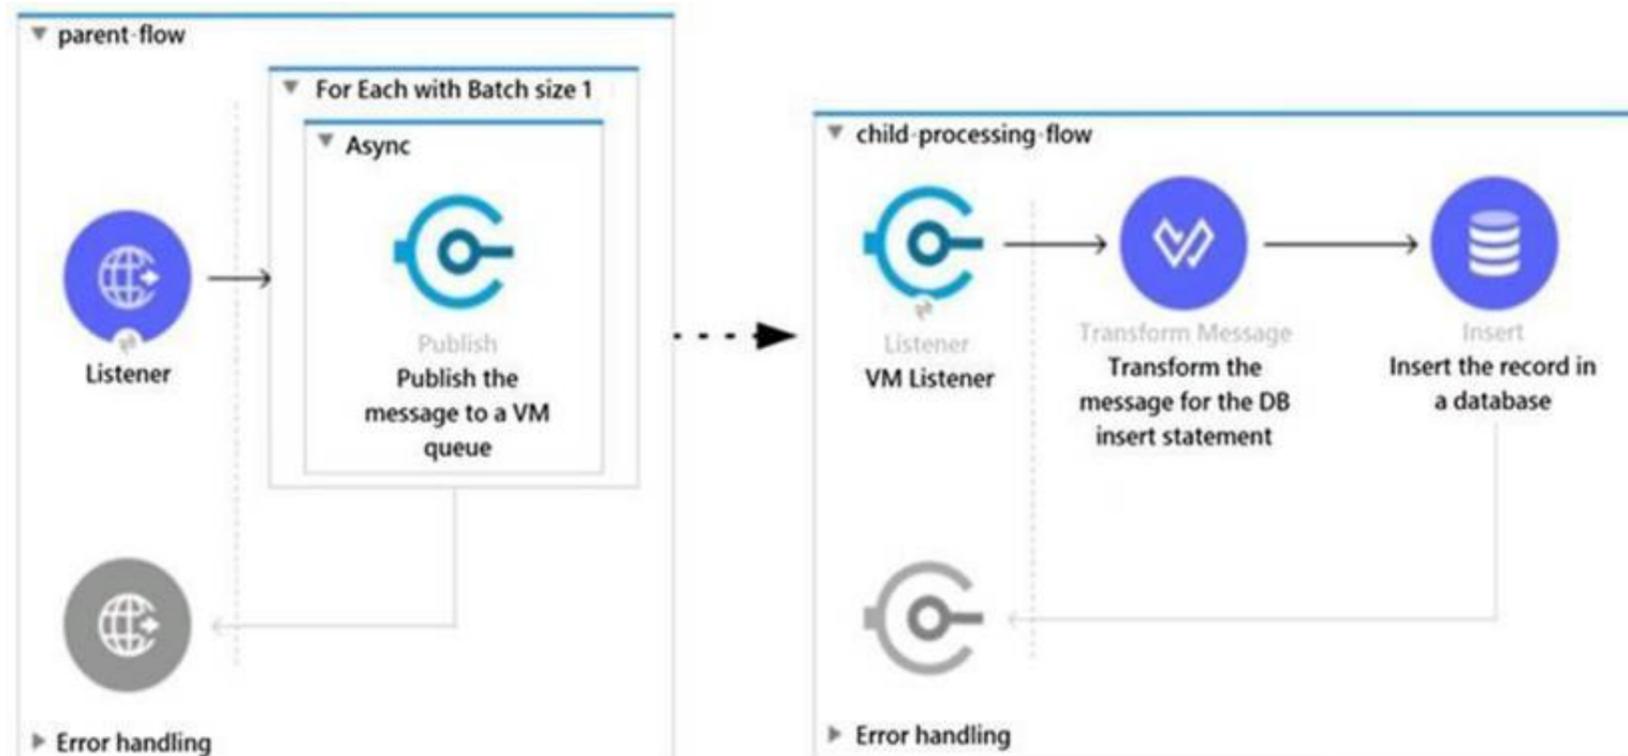
### MuleSoft Certified Integration Architect - Level 1

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



**NEW QUESTION 1**

Refer to the exhibit.



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 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.

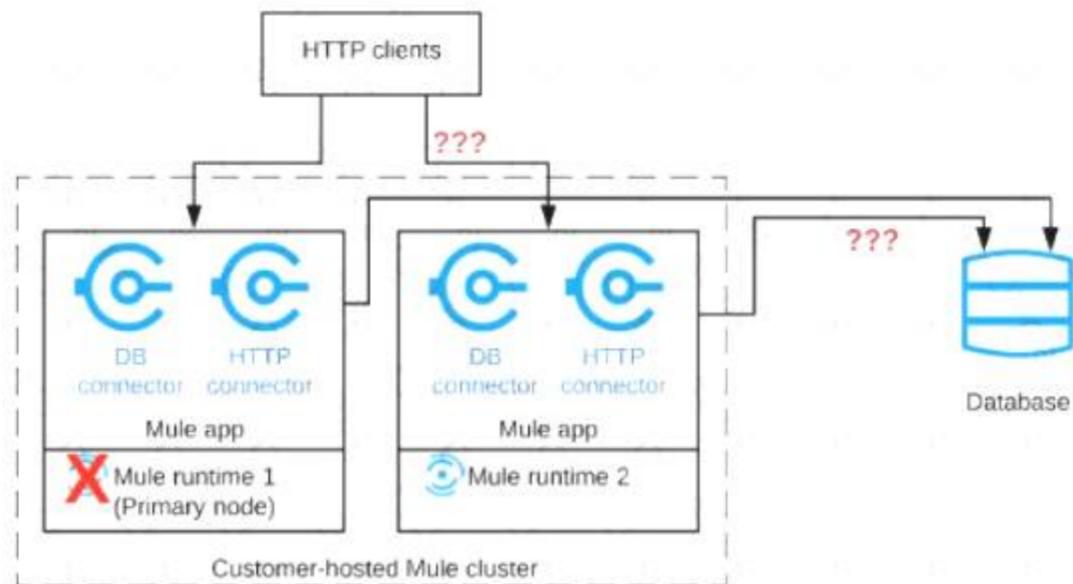
What key elements should be included in the integration design specification when requesting approval for this Mule application?

- 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**

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.

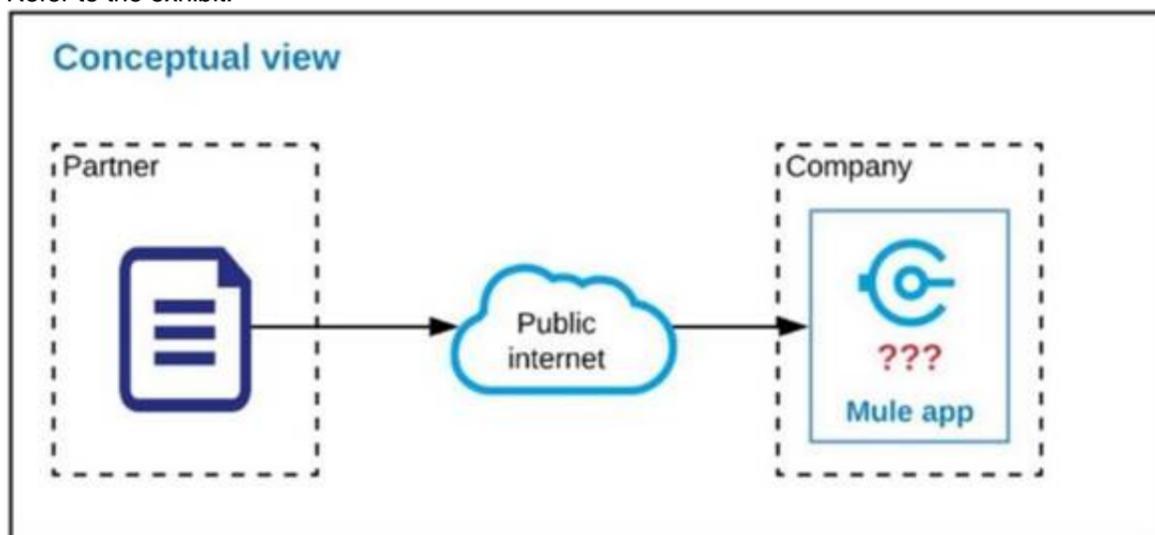
What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has failed, but before that node is restarted?

- 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 5**

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).

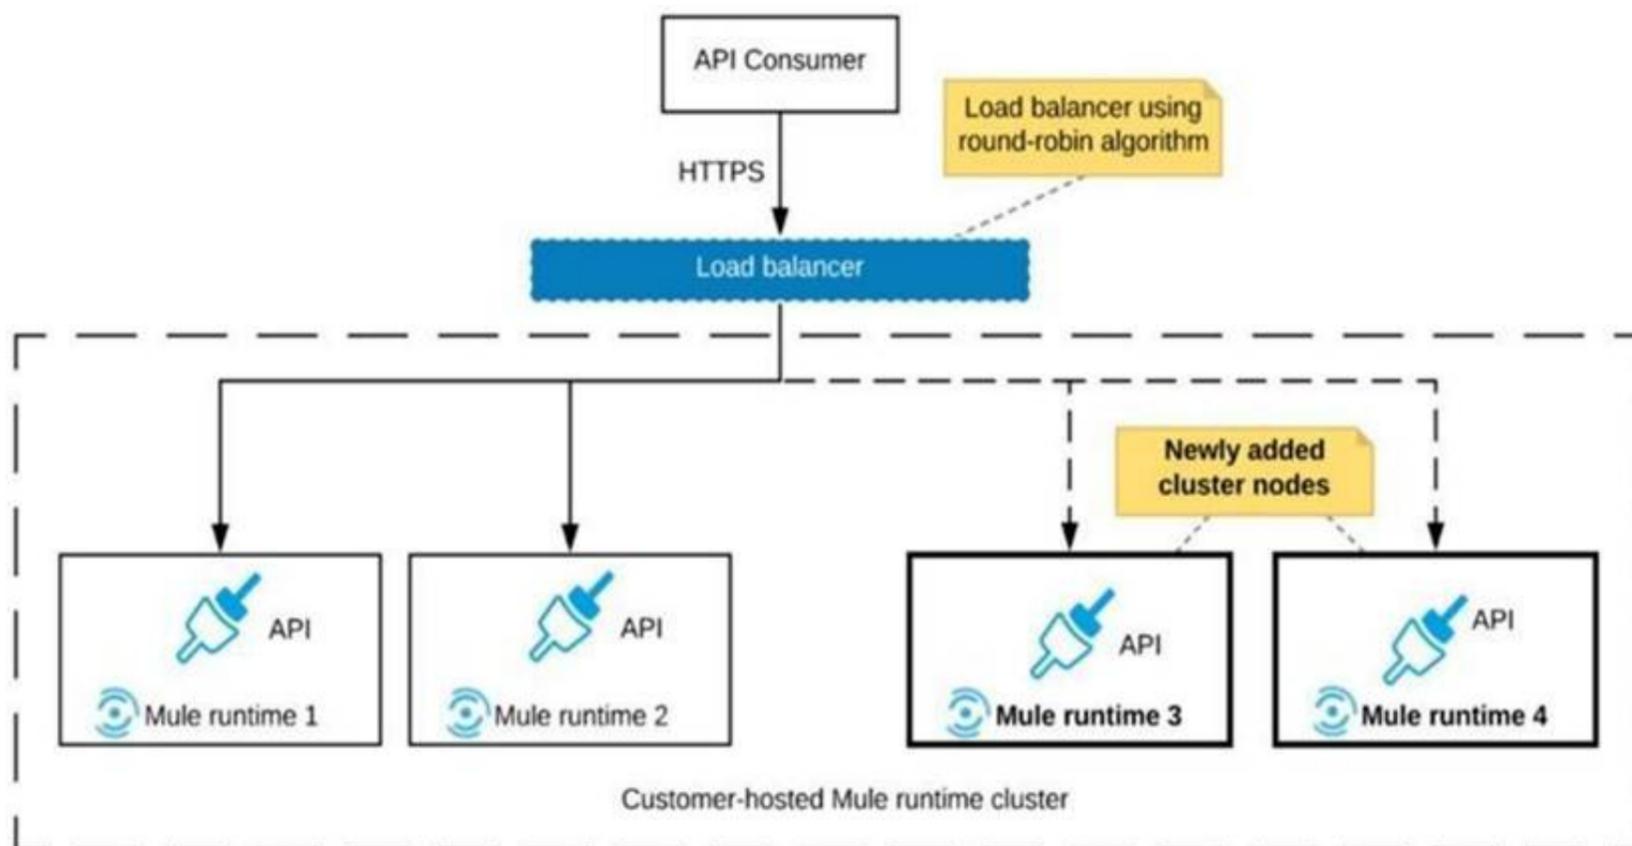
What Anypoint Connector can be used in the organization's Mule application to securely receive data from this external business partner?

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

**Answer: C**

**NEW QUESTION 6**

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.

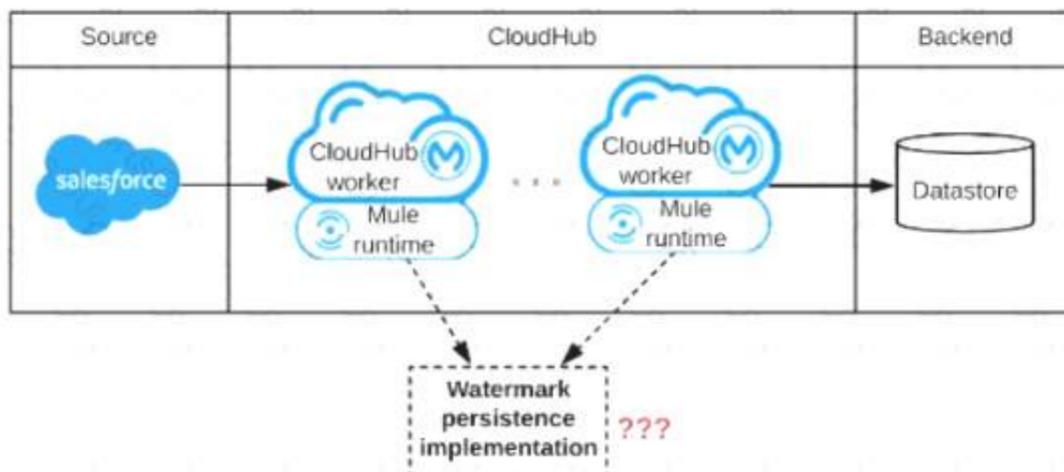
What average performance change is guaranteed to happen, assuming all cluster nodes are fully operational?

- 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 7**

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 8**

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 9**

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 10**

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.

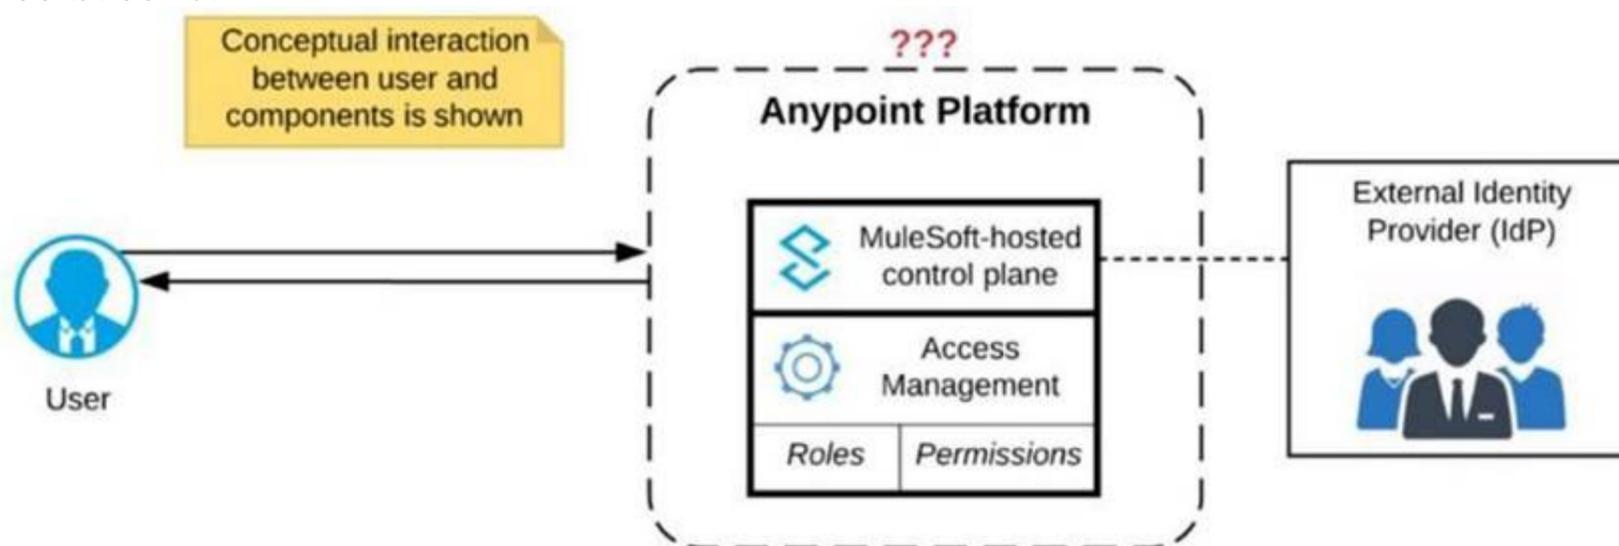
What combination of deployment options of the Anypoint Platform control plane and runtime plane(s) best serves this organization at the start of this strategic journey?

- 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 10**

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 15**

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 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 21**

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 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.

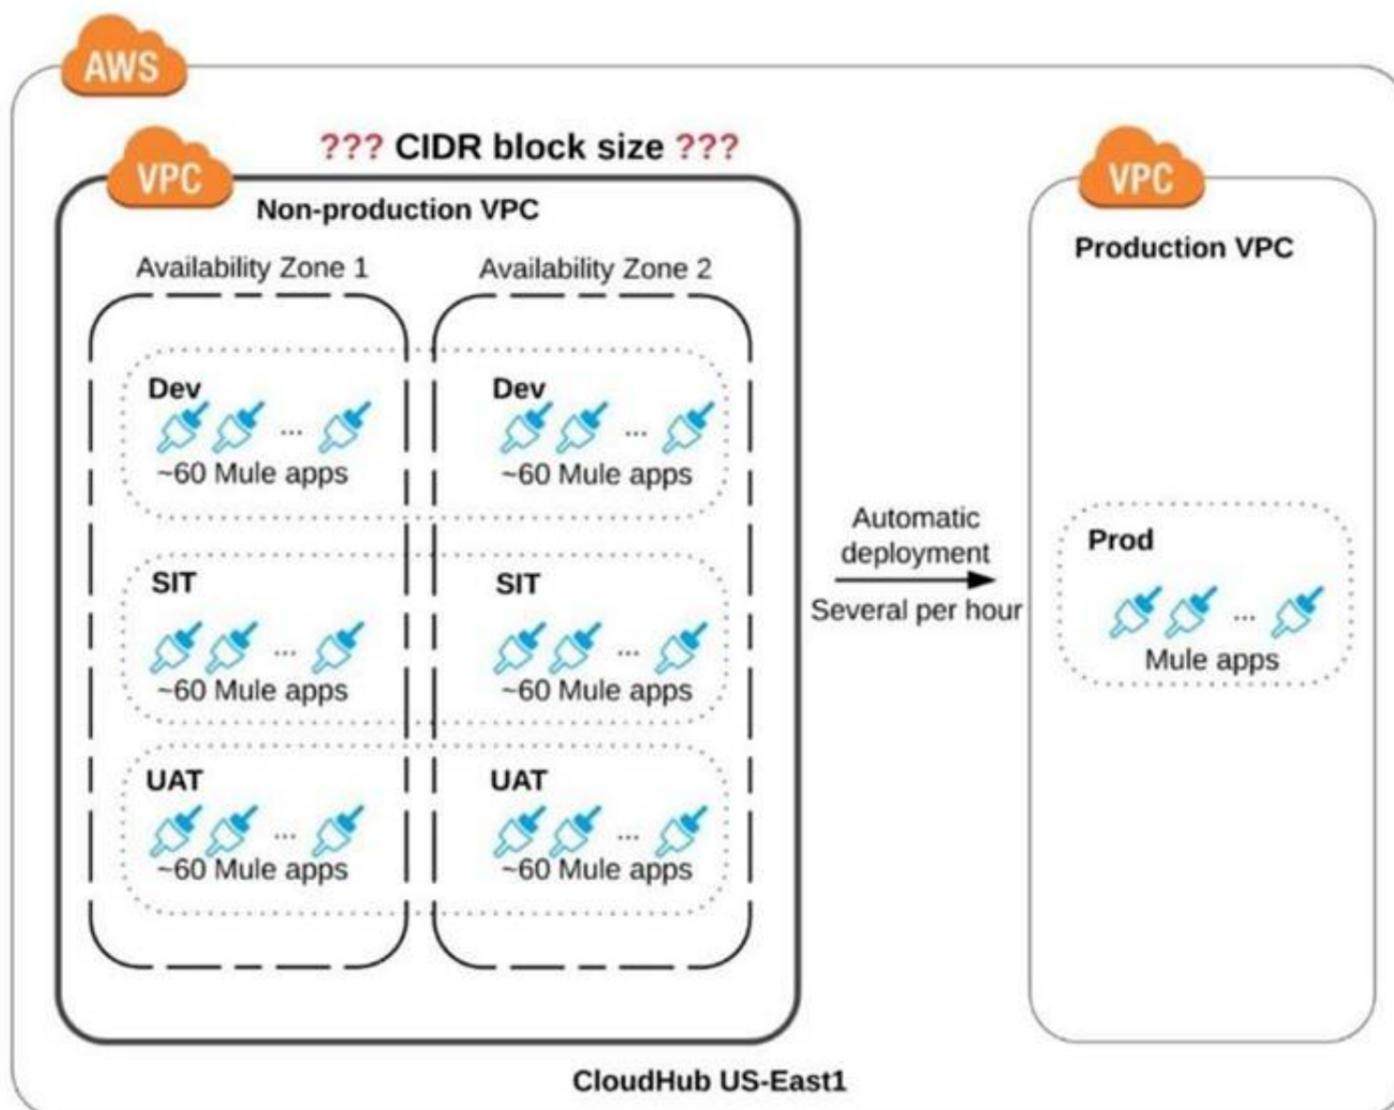
How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access thebackend systems?

- 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 25**

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 28**

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.

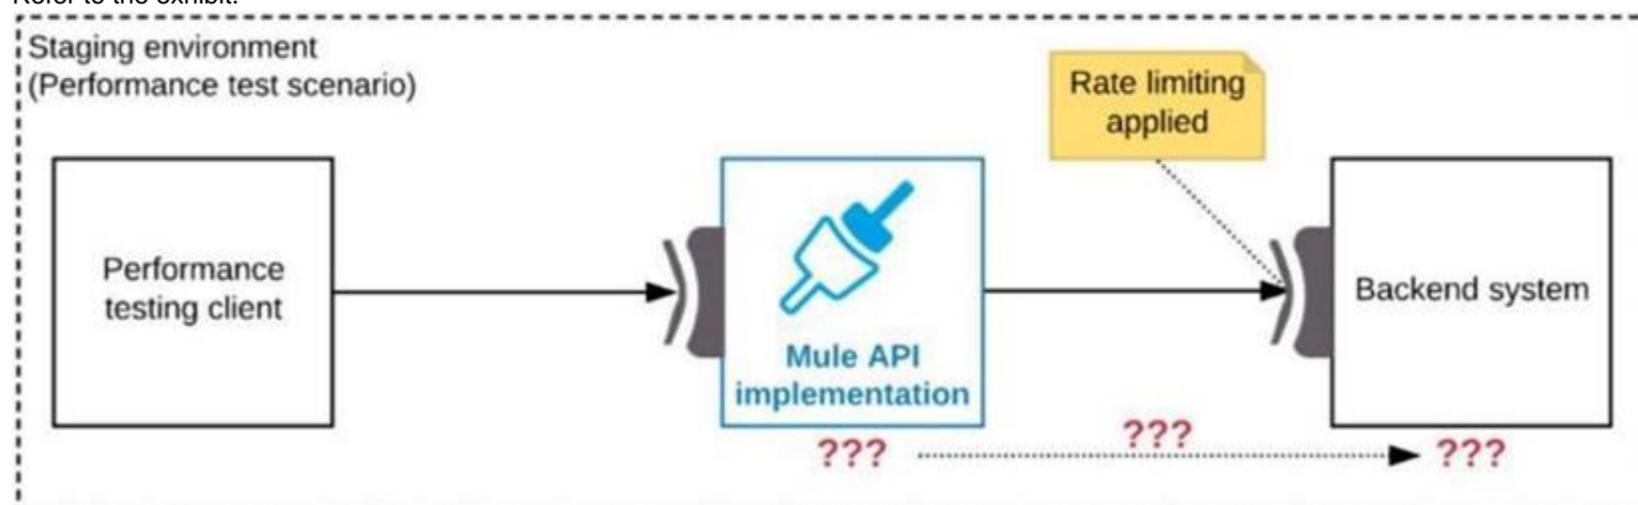
What is the most appropriate integration style for an integration solution that meets the organization's current requirements?

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

**Answer: D**

**NEW QUESTION 34**

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 38**

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 41**

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 45**

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 47**

An organization is creating a set of new services that are critical for their business. The project team prefers using REST for all services but is willing to use SOAP with common WS-\* standards if a particular service requires it.

What requirement would drive the team to use SOAP/WS-\* for a particular service?

- A. Must secure the service, requiring all consumers to submit a valid SAML token
- B. Must support message acknowledgement and retry as part of the protocol
- C. Must publish and share the service specification (including data formats) with the consumers of the service
- D. Must use XML payloads for the service and ensure that it adheres to a specific schema

**Answer:** B

**NEW QUESTION 49**

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 50**

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 51**

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?

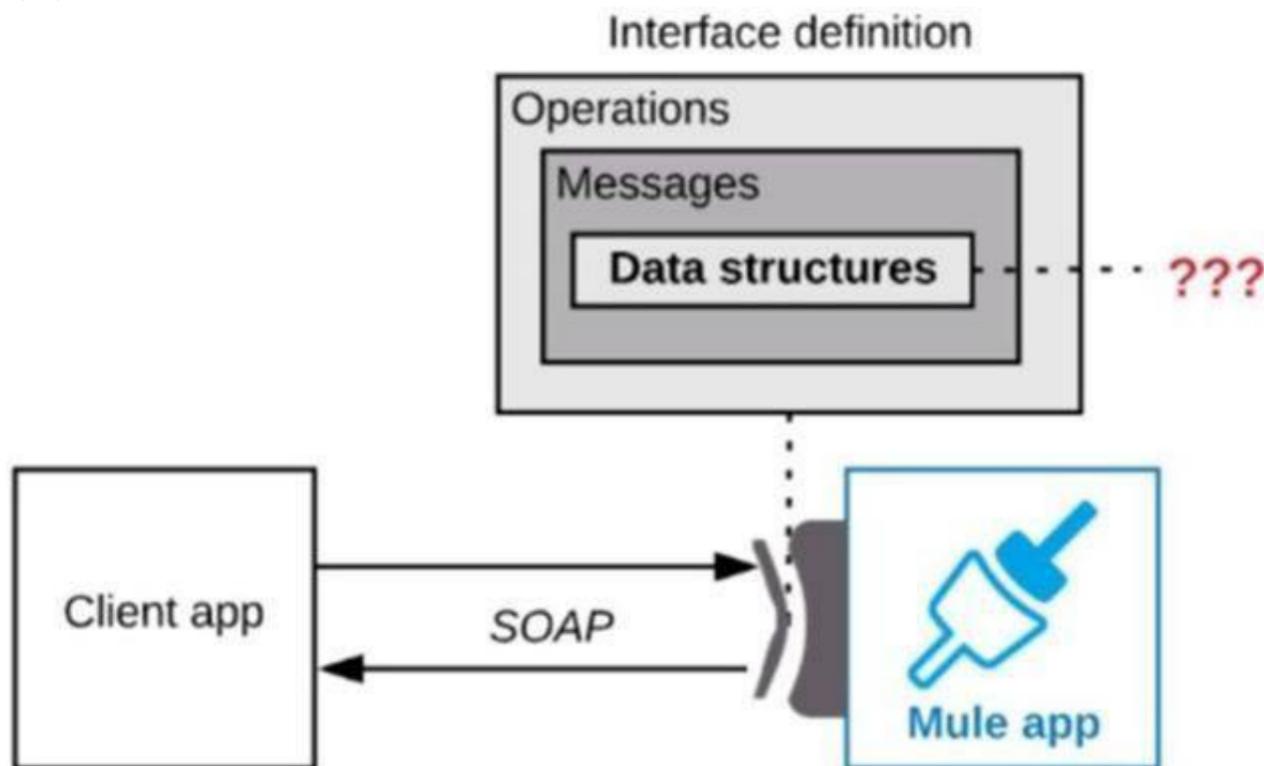
- A. Anypoint Runtime Manager
- B. Anypoint Exchange
- C. Anypoint API Manager
- D. The Mule runtimes

**Answer:** A

**NEW QUESTION 52**

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 56**

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 60**

.....

## 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>