

Exam Questions MuleSoft-Integration-Architect-I

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam

<https://www.2passeasy.com/dumps/MuleSoft-Integration-Architect-I/>



NEW QUESTION 1

An organization plans to use the Anypoint Platform audit logging service to log Anypoint MQ actions. What consideration must be kept in mind when leveraging Anypoint MQ Audit Logs?

- A. Anypoint MQ Audit Logs include logs for sending, receiving, or browsing messages
- B. Anypoint MQ Audit Logs include logs for failed Anypoint MQ operations
- C. Anypoint MQ Audit Logs include logs for queue create, delete, modify, and purge operations

Answer: C

NEW QUESTION 2

According to MuleSoft, which system integration term describes the method, format, and protocol used for communication between two systems?

- A. Component
- B. interaction
- C. Message
- D. Interface

Answer: D

NEW QUESTION 3

An API implementation is being designed that must invoke an Order API which is known to repeatedly experience downtime. For this reason a fallback API is to be called when the Order API is unavailable. What approach to designing invocation of the fallback API provides the best resilience?

- A. Redirect client requests through an HTTP 303 temporary redirect status code to the fallback API whenever the Order API is unavailable
- B. Set an option in the HTTP Requester component that invokes the order API to instead invoke a fallback API whenever an HTTP 4XX or 5XX response status code is received from Order API
- C. Create a separate entry for the order API in API manager and then invoke this API as a fallback API if the primary Order API is unavailable
- D. Search Anypoint Exchange for a suitable existing fallback API and then implement invocations to their fallback API in addition to the Order API

Answer: A

NEW QUESTION 4

Mule applications need to be deployed to CloudHub so they can access on-premises database systems. These systems store sensitive and hence tightly protected data, so are not accessible over the internet.

What network architecture supports this requirement?

- A. An Anypoint VPC connected to the on-premises network using an IPsec tunnel or AWS DirectConnect, plus matching firewall rules in the VPC and on-premises network
- B. Static IP addresses for the Mule applications deployed to the CloudHub Shared Worker Cloud, plus matching firewall rules and IP whitelisting in the on-premises network
- C. An Anypoint VPC with one Dedicated Load Balancer fronting each on-premises database system, plus matching IP whitelisting in the load balancer and firewall rules in the VPC and on-premises network
- D. Relocation of the database systems to a DMZ in the on-premises network, with Mule applications deployed to the CloudHub Shared Worker Cloud connecting only to the DMZ

Answer: A

NEW QUESTION 5

What is an advantage that Anypoint Platform offers by providing universal API management and Integration-Platform-as-a-Service (iPaaS) capabilities in a unified platform?

- A. Ability to use a single iPaaS to manage and integrate all API gateways
- B. Ability to use a single connector to manage and integrate all APIs
- C. Ability to use a single control plane for both full-lifecycle API management and integration
- D. Ability to use a single iPaaS to manage all API developer portals

Answer: C

NEW QUESTION 6

An API implementation is being developed to expose data from a production database via HTTP requests. The API implementation executes a database SELECT statement that is dynamically created based upon data received from each incoming HTTP request. The developers are planning to use various types of testing to make sure the Mule application works as expected, can handle specific workloads, and behaves correctly from an API consumer perspective. What type of testing would typically mock the results from each SELECT statement rather than actually execute it in the production database?

- A. Unit testing (white box)
- B. Integration testing
- C. Functional testing (black box)
- D. Performance testing

Answer: A

NEW QUESTION 7

Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improve throughput of the system. But it was observed over the period of time that few of the important exception log messages which were used to rollback transactions are not working as expected causing huge loss

to the Organization. Organization wants to avoid these losses. Application also has constraints due to which they cant compromise on throughput much. What is the possible option in this case?

- A. Logging needs to be changed from asynchronous to synchronous
- B. External log appender needs to be used in this case
- C. Persistent memory storage should be used in such scenarios
- D. Mixed configuration of asynchronous or synchronous loggers should be used to log exceptions via synchronous way

Answer: D

NEW QUESTION 8

A Mule application is being designed for deployment to a single CloudHub worker. The Mule application will have a flow that connects to a SaaS system to perform some operations each time the flow is invoked.

The SaaS system connector has operations that can be configured to request a short-lived token (fifteen minutes) that can be reused for subsequent connections within the fifteen minute time window. After the token expires, a new token must be requested and stored.

What is the most performant and idiomatic (used for its intended purpose) Anypoint Platform component or service to use to support persisting and reusing tokens in the Mule application to help speed up reconnecting the Mule application to the SaaS application?

- A. Nonpersistent object store
- B. Persistent object store
- C. Variable
- D. Database

Answer: D

NEW QUESTION 9

Which Anypoint Platform component helps integration developers discover and share reusable APIs, connectors, and templates?

- A. Anypoint Exchange
- B. API Manager
- C. Anypoint Studio
- D. Design Center

Answer: A

NEW QUESTION 10

A Mule application name Pub uses a persistence object store. The Pub Mule application is deployed to Cloudhub and it configured to use Object Store v2. Another Mule application name sub is being developed to retrieve values from the Pub Mule application persistence object Store and will also be deployed to cloudhub.

What is the most direct way for the Sub Mule application to retrieve values from the Pub Mule application persistence object store with the least latency?

- A. Use an object store connector configured to access the Pub Mule application persistence object store
- B. Use a VM connector configured to directly access the persistence queue of the Pub Mule application persistence object store.
- C. Use an Anypoint MQ connector configured to directly access the Pub Mule application persistence object store
- D. Use the Object store v2 REST API configured to access the Pub Mule application persistence object store.

Answer: D

NEW QUESTION 10

According to MuleSoft's IT delivery and operating model, which approach can an organization adopt in order to reduce the frequency of IT project delivery failures?

- A. Decouple central IT projects from the innovation that happens within each line of business
- B. Adopt an enterprise data model
- C. Prevent technology sprawl by reducing production of API assets
- D. Stop scope creep by centralizing requirements-gathering

Answer: A

NEW QUESTION 11

What is true about the network connections when a Mule application uses a JMS connector to interact with a JMS provider (message broker)?

- A. To complete sending a JMS message, the JMS connector must establish a network connection with the JMS message recipient
- B. To receive messages into the Mule application, the JMS provider initiates a network connection to the JMS connector and pushes messages along this connection
- C. The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider
- D. The AMQP protocol can be used by the JMS connector to portably establish connections to various types of JMS providers

Answer: C

NEW QUESTION 16

A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should be published in a way that is guaranteed in normal situations and also minimizes duplicate delivery in less frequent failure scenarios.

The organizational firewall is configured to only allow outbound traffic on ports 80 and 443. Some external event consumers are within the organizational network, while others are located outside the firewall.

What Anypoint Platform service is most idiomatic (used for its intended purpose) for publishing these Mule events to all external consumers while addressing the desired reliability goals?

- A. CloudHub VM queues
- B. Anypoint MQ
- C. Anypoint Exchange
- D. CloudHub Shared Load Balancer

Answer: B

NEW QUESTION 18

A rate limiting policy has been applied to a soap V1.2 API published in Clonhub. The API implementation catches errors in a global error handler on error propagate in the main flow for HTTP: RETRY_EXHAUSTED with HTTP status set to 429 and any with the HTTP status set to 500. What is the expected H1TP status when the client exceeds the quota of the API calls?

- A. HTTP status 429 as defined in the HTTP:RETRY EXHAUSTED error handler in the API
- B. HTTP status 500 as defined in the ANY error handler in the API since an API:RETRY_EXHAUSTED will be generated
- C. HTTP status 401 unauthorized for policy violation
- D. HTTP status 400 from the rate-limiting policy violation since the call does not reach the back-end


Answer: A

NEW QUESTION 22


A gaming company has implemented an API as a Mule application and deployed the API implementation to a CloudHub 2.0 private space. The API implementation must connect to a mainframe application running in the customer??s on-premises corporate data center and also to a Kafka cluster running in an Amazon AWS VPC.

What is the most efficient way to enable the API to securely connect from its private space to the mainframe application and Kafka cluster?


- A. In Runtime Manager, set up VPC peering between the CloudHub 2.0 private network and the on-premises data center. In the AWS account, set up VPC peering between the AWS VPC and the CloudHub 2.0 private network.




- B) In the AWS account, attach the CloudHub 2.0 private space to an AWS transit gateway that routes from the CloudHub 2.0 private space to the on-premises data centre
- B. InRuntime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the AWS VPC.



- C) In Runtime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the on-premises data centre
- C. In the AWS account, attach the private..



D) In the AWS account, attach the private space directly to the AWS VPC, In the AWS account, use an AWS transit gateway to route from the AWS VPC to the on-premises data center.



- D. Option A
- E. Option B
- F. Option C
- G. Option D

Answer: B

NEW QUESTION 26

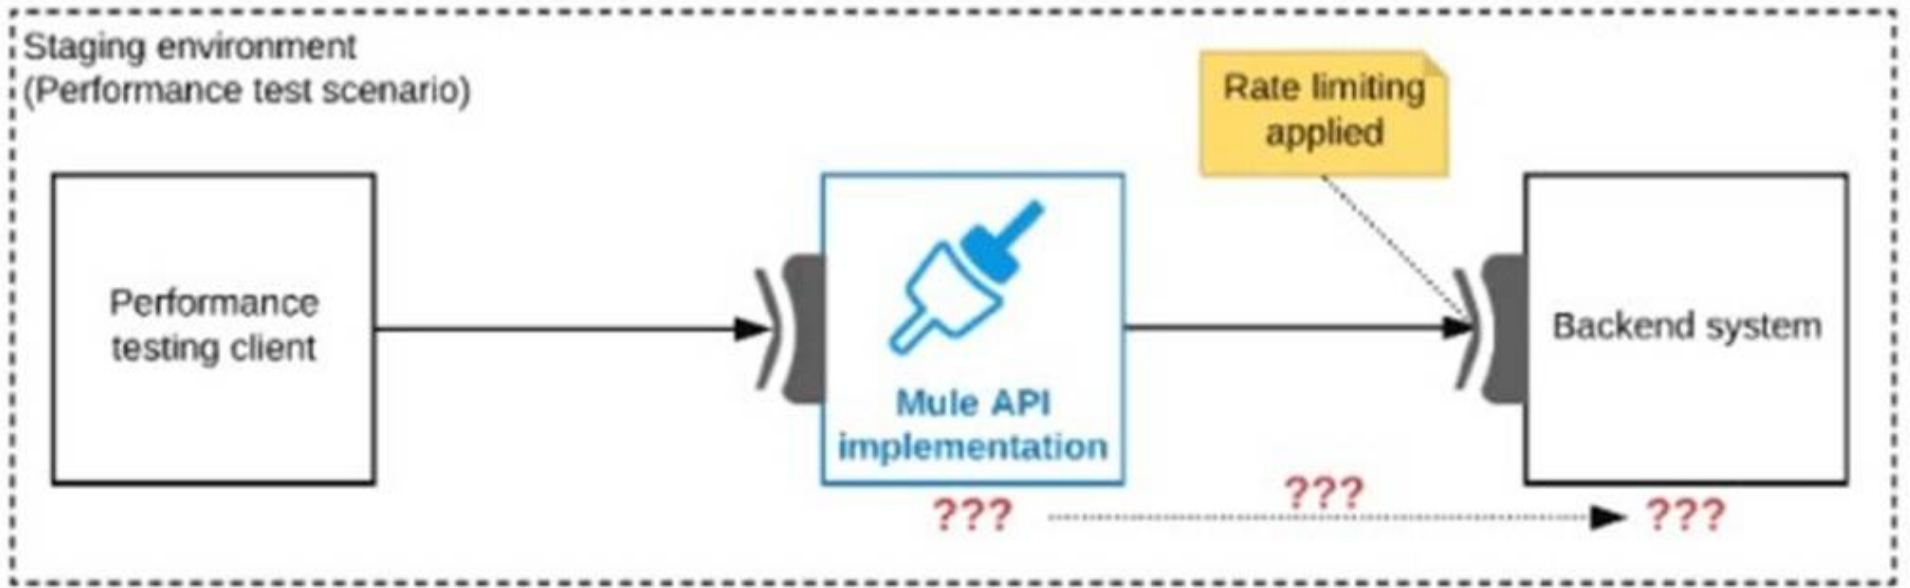
Which Mulesoft feature helps users to delegate their access without sharing sensitive credentials or giving full control of accounts to 3rd parties?

- A. Secure Scheme
- B. client id enforcement policy
- C. Connected apps
- D. Certificates

Answer: C

NEW QUESTION 31

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. Create a mocking service that replicates the backend system's production performance characteristic
- B. Then configure the API implementation to use the mocking service and conduct the performance tests
- C. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
- D. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situation
- E. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
- F. Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale

Answer: A

NEW QUESTION 33

An organization has strict unit test requirement that mandate every mule application must have an MUnit test suit with a test case defined for each flow and a minimum test coverage of 80%.

A developer is building Munit test suit for a newly developed mule application that sends API request to an external rest API.

What is the effective approach for successfully executing the Munit tests of this new application while still achieving the required test coverage for the Munit tests?

- A. Invoke the external endpoint of the rest API from the mule flows
- B. Mark the rest API invocations in the Munits and then call the mocking service flow that simulates standard responses from the REST API
- C. Mock the rest API invocation in the Munits and return a mock response for those invocations
- D. Create a mocking service flow to simulate standard responses from the rest API and then configure the mule flows to call the marking service flow

Answer: C

NEW QUESTION 34

According to MuleSoft's API development best practices, which type of API development approach starts with writing and approving an API contract?

- A. Implement-first
- B. Catalyst
- C. Agile
- D. Design-first

Answer: D

NEW QUESTION 35

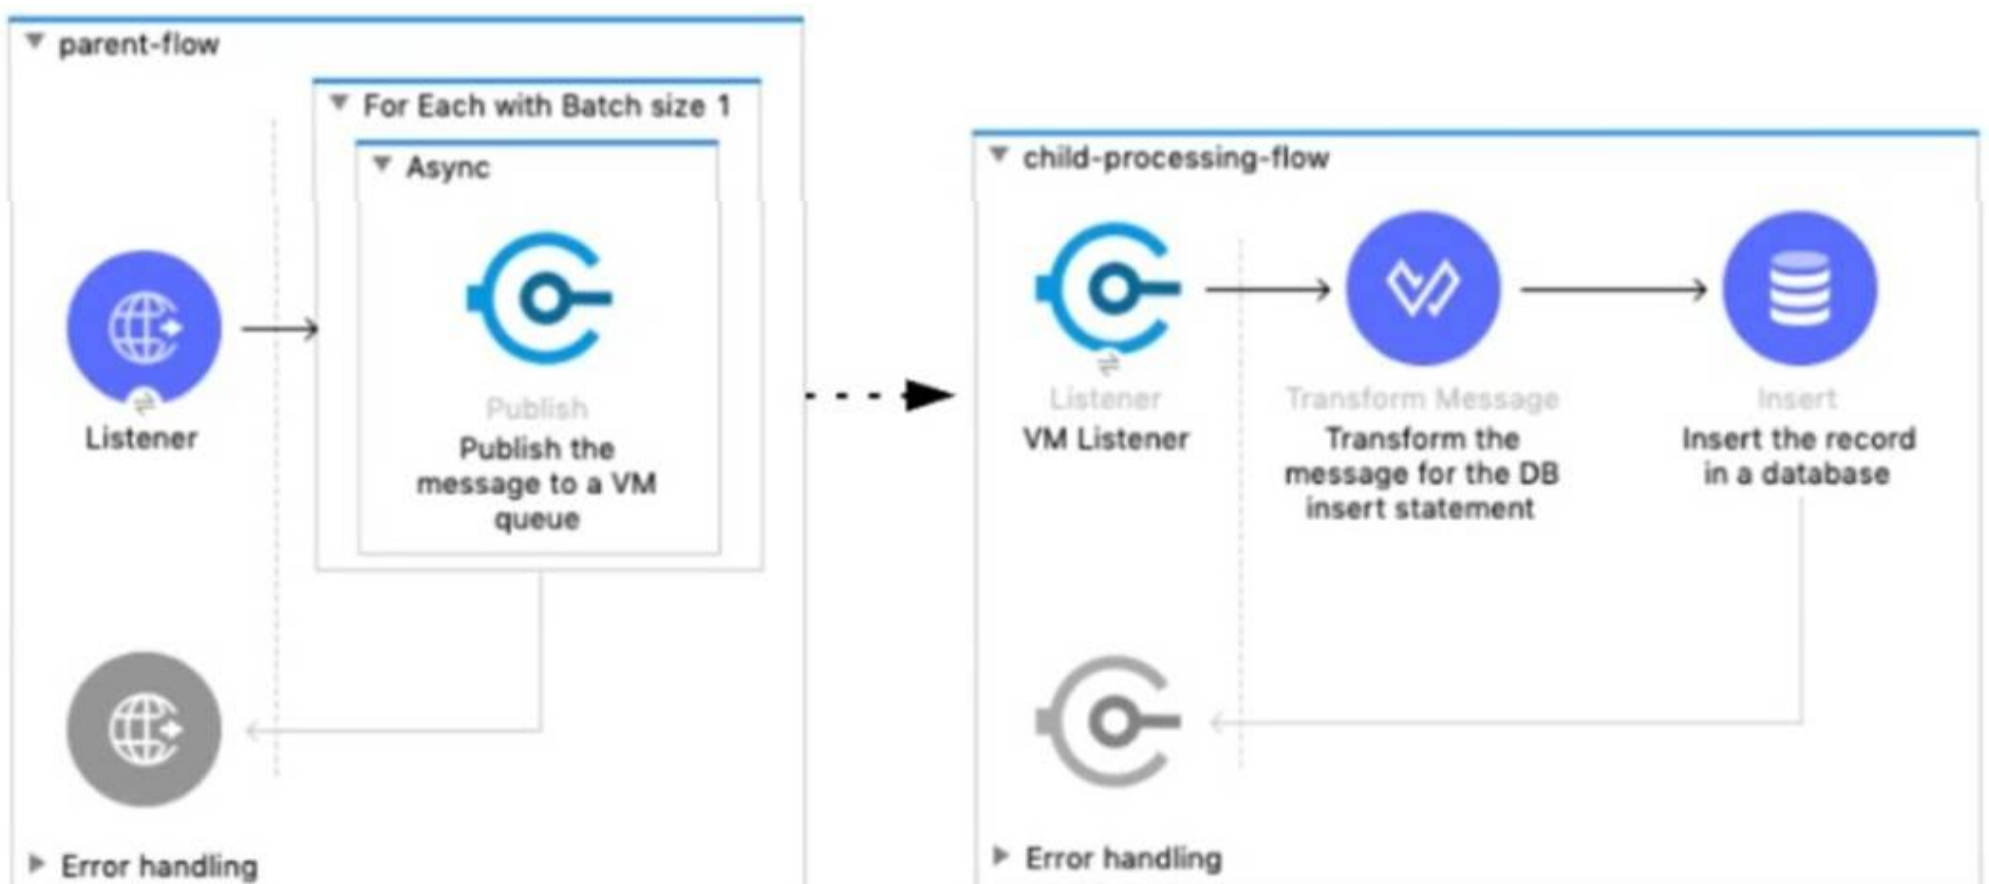
When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

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

Answer: B

NEW QUESTION 38

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 41

An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/CD pipeline. Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?

- A. Use runtime manager rest apis for API management and mavenforAPI deployment
- B. Use Maven with a custom configuration required for the API lifecycle
- C. Use Anypoint CLI or Anypoint Platform REST apis with scripting language such as groovy
- D. Use Exchange rest api's for API management and MavenforAPI deployment

Answer: C

NEW QUESTION 43

An organization has decided on a cloud migration strategy to minimize the organization's own IT resources. Currently the organization has all of its new applications running on its own premises and uses an on-premises load balancer that exposes all APIs under the base URL (<https://api.rutujar.com>).

As part of migration strategy, the organization is planning to migrate all of its new applications and load balancer CloudHub.

What is the most straightforward and cost-effective approach to Mule application deployment and load balancing that preserves the public URL's?

- A. Deploy the Mule application to Cloudhub Create a CNAME record for base URL(<https://api.rutujar.com>) in the Cloudhub shared load balancer that points to the A record of the on-premises load balancer Apply mapping rules in SLB to map URL to their corresponding Mule applications
- B. Deploy the Mule application to Cloudhub Update a CNAME record for base URL (<https://api.rutujar.com>) in the organization's DNS server to point to the A record of the Cloudhub dedicated load balancer Apply mapping rules in DLB to map URL to their corresponding Mule applications
- C. Deploy the Mule application to Cloudhub Update a CNAME record for base URL (<https://api.rutujar.com>) in the organization's DNS server to point to the A record of the CloudHub shared load balancer Apply mapping rules in SLB to map URL to their corresponding Mule applications
- D. For each migrated Mule application, deploy an API proxy application to Cloudhub with all traffic to the mule applications routed through a Cloud Hub Dedicated load balancer (DLB) Update a CNAME record for base URL (<https://api.rutujar.com>) in the organization's DNS server to point to the A record of the CloudHub dedicated load balancer Apply mapping rules in DLB to map each API proxy application who is responding new application

Answer: C

NEW QUESTION 44

An external API frequently invokes an Employees System API to fetch employee data from a MySQL database. The architect must design a caching strategy to query the database only when there is an update to the Employees table or else return a cached response in order to minimize the number of redundant transactions being handled by the database.

- A. Use an On Table Row operation configured with the Employees table, call invalidate cache, and hardcode the new Employees data to cache
- B. Use an object-store-caching- strategy and set the expiration interval to 1 hour.
- C. Use an On Table Row operation configured with the Employees table and call invalidate cache
- D. Use an object-store-caching-strategy and the default expiration interval.
- E. Use a Scheduler with a fixed frequency set to every hour to trigger an invalidate cache flow
- F. Use an object-store-caching-strategy and the default expiration interval.
- G. Use a Scheduler with a fixed frequency set to every hour, triggering an invalidate cache flow
- H. Use an object-store-caching-strategy and set the expiration interval to 1 hour.

Answer: B

NEW QUESTION 49

In a Mule Application, a flow contains two (2) JMS consume operations that are used to connect to a JMS broker and consume messages from two(2) JMS destination. The Mule application then joins the two JMS messages together.

The JMS broker does not implement high availability (HA) and periodically experiences scheduled outages of upto 10 mins for routine maintenance.

What is the most idiomatic (used for its intended purpose) way to build the mule flow so it can best recover from the expected outages?

- A. Configure a reconnection strategy for the JMS connector
- B. Enclose the two(2) JMS operation in an Until Successful scope
- C. Consider a transaction for the JMS connector
- D. Enclose the two(2) JMS operations in a Try scope with an Error Continue error handler

Answer: A

NEW QUESTION 52

What Mule application can have API policies applied by Anypoint Platform to the endpoint exposed by that Mule application?

- A. A Mule application that accepts requests over HTTP/1x
- B. A Mule application that accepts JSON requests over TCP but is NOT required to provide a response.
- C. A Mule application that accepts JSON requests over WebSocket
- D. A Mule application that accepts gRPC requests over HTTP/2

Answer: A

NEW QUESTION 54

In Anypoint Platform, a company wants to configure multiple identity providers (IdPs) for multiple lines of business (LOBs). Multiple business groups, teams, and environments have been defined for these LOBs.

What Anypoint Platform feature can use multiple IdPs across the company's business groups, teams, and environments?

- A. MuleSoft-hosted (CloudHub) dedicated load balancers
- B. Client (application) management
- C. Virtual private clouds
- D. Permissions

Answer: A

NEW QUESTION 58

Following MuleSoft best practices, what MuleSoft runtime deployment option best meets the company's goals to begin its digital transformation journey?

- A. Runtime Fabric on VMs/bare metal
- B. CloudHub runtimes
- C. Customer-hosted runtimes provisioned by a MuleSoft services partner
- D. Customer-hosted self-provisioned runtimes

Answer: B

NEW QUESTION 63

Which component of Anypoint platform belongs to the platform control plane?

- A. Runtime Fabric
- B. Runtime Replica
- C. Anypoint Connectors
- D. API Manager

Answer: D

NEW QUESTION 64

Organization wants to achieve high availability goal for Mule applications in customer hosted runtime plane. Due to the complexity involved, data cannot be shared among of different instances of same Mule application. What option best suits to this requirement considering high availability is very much critical to the organization?

- A. The cluster can be configured
- B. Use third party product to implement load balancer
- C. High availability can be achieved only in CloudHub
- D. Use persistent object store

Answer: B

NEW QUESTION 69

According to MuteSoft, which principle is common to both Service Oriented Architecture (SOA) and API-led connectivity approaches?

- A. Service centralization
- B. Service statefulness
- C. Service reusability
- D. Service interdependence

Answer: C

NEW QUESTION 71

An external web UI application currently accepts occasional HTTP requests from client web browsers to change (insert, update, or delete) inventory pricing information in an inventory system's database. Each inventory pricing change must be transformed and then synchronized with multiple customer experience systems in near real-time (in under 10 seconds). New customer experience systems are expected to be added in the future.

The database is used heavily and limits the number of SELECT queries that can be made to the database to 10 requests per hour per user.

What is the most scalable, idiomatic (used for its intended purpose), decoupled, reusable, and maintainable integration mechanism available to synchronize each inventory pricing change with the various customer experience systems in near real-time?

- A. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the watermark attribute set to an appropriate database columnIn the same now, use a Scatter-Gather to call each customer experience system's REST API with transformed inventory-pricing records
- B. Add a trigger to the inventory-pricing database table so that for each change to the inventory pricing database, a stored procedure is called that makes a REST call to a Mule applicationWrite the Mule application to publish each Mule event as a message to an Anypoint MQ exchangeWrite other Mule applications to subscribe to the Anypoint MQ exchange, transform each received message, and then update the Mule application's corresponding customer experience system(s)

C. Replace the external web UI application with a Mule application to accept HTTP requests from client web browsers. In the same Mule application, use a Batch Job scope to test if the database request will succeed, aggregate pricing changes within a short time window, and then update both the inventory pricing database and each customer experience system using a Parallel For Each scope.
 D. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the ID attribute set to an appropriate database column. In the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queue. Write other Mule applications to subscribe to the Anypoint MQ queue, transform each received message, and then update the Mule application's corresponding customer experience system(s).

Answer: B

NEW QUESTION 74

An insurance company is implementing a MuleSoft API to get inventory details from the two vendors. Due to network issues, the invocations to vendor applications are getting timed-out intermittently. But the transactions are successful upon reprocessing. What is the most performant way of implementing this requirement?

- A. Implement a scatter-gather scope to invoke the two vendor applications on two different routes. Use the Until-Successful scope to implement the retry mechanism for timeout errors on each route.
- B. Implement a Choice scope to invoke the two vendor applications on two different routes. Use the try-catch scope to implement the retry mechanism for timeout errors on each route.
- C. Implement a For-Each scope to invoke the two vendor applications. Use until successful scope to implement the retry mechanism for the timeout errors.
- D. Implement Round-Robin scope to invoke the two vendor applications on two different routes. Use the Try-Catch scope to implement retry mechanism for timeout errors on each route.

Answer: A

NEW QUESTION 77

A leading eCommerce giant will use MuleSoft APIs on Runtime Fabric (RTF) to process customer orders. Some customer-sensitive information, such as credit card information, is required in request payloads or is included in response payloads in some of the APIs. Other API requests and responses are not authorized to access some of this customer-sensitive information but have been implemented to validate and transform based on the structure and format of this customer-sensitive information (such as account IDs, phone numbers, and postal codes).

What approach configures an API gateway to hide sensitive data exchanged between API consumers and API implementations, but can convert tokenized fields back to their original value for other API requests or responses, without having to recode the API implementations?

Later, the project team requires all API specifications to be augmented with an additional non-functional requirement (NFR) to protect the backend services from a high rate of requests, according to defined service-level agreements (SLAs). The NFR's SLAs are based on a new tiered subscription level "Gold", "Silver", or "Platinum" that must be tied to a new parameter that is being added to the Accounts object in their enterprise data model.

Following MuleSoft's recommended best practices, how should the project team now convey the necessary non-functional requirement to stakeholders?

- A. Create and deploy API proxies in API Manager for the NFR, change the baseUrl in each API specification to the corresponding API proxy implementation endpoint, and publish each modified API specification to Exchange.
- B. Update each API specification with comments about the NFR's SLAs and publish each modified API specification to Exchange.
- C. Update each API specification with a shared RAML fragment required to implement the NFR and publish the RAML fragment and each modified API specification to Exchange.
- D. Create a shared RAML fragment required to implement the NFR, list each API implementation endpoint in the RAML fragment, and publish the RAML fragment to Exchange.

Answer: C

NEW QUESTION 79

A Mule application is being designed to receive nightly a CSV file containing millions of records from an external vendor over SFTP. The records from the file need to be validated, transformed, and then written to a database. Records can be inserted into the database in any order.

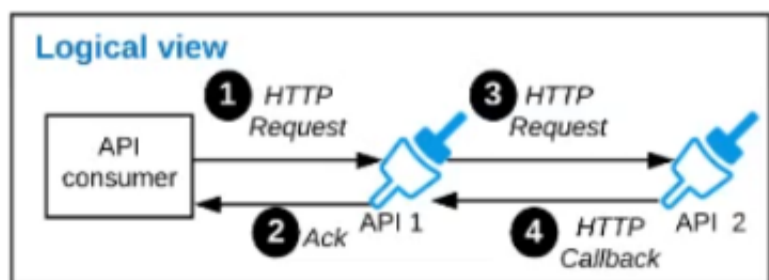
In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

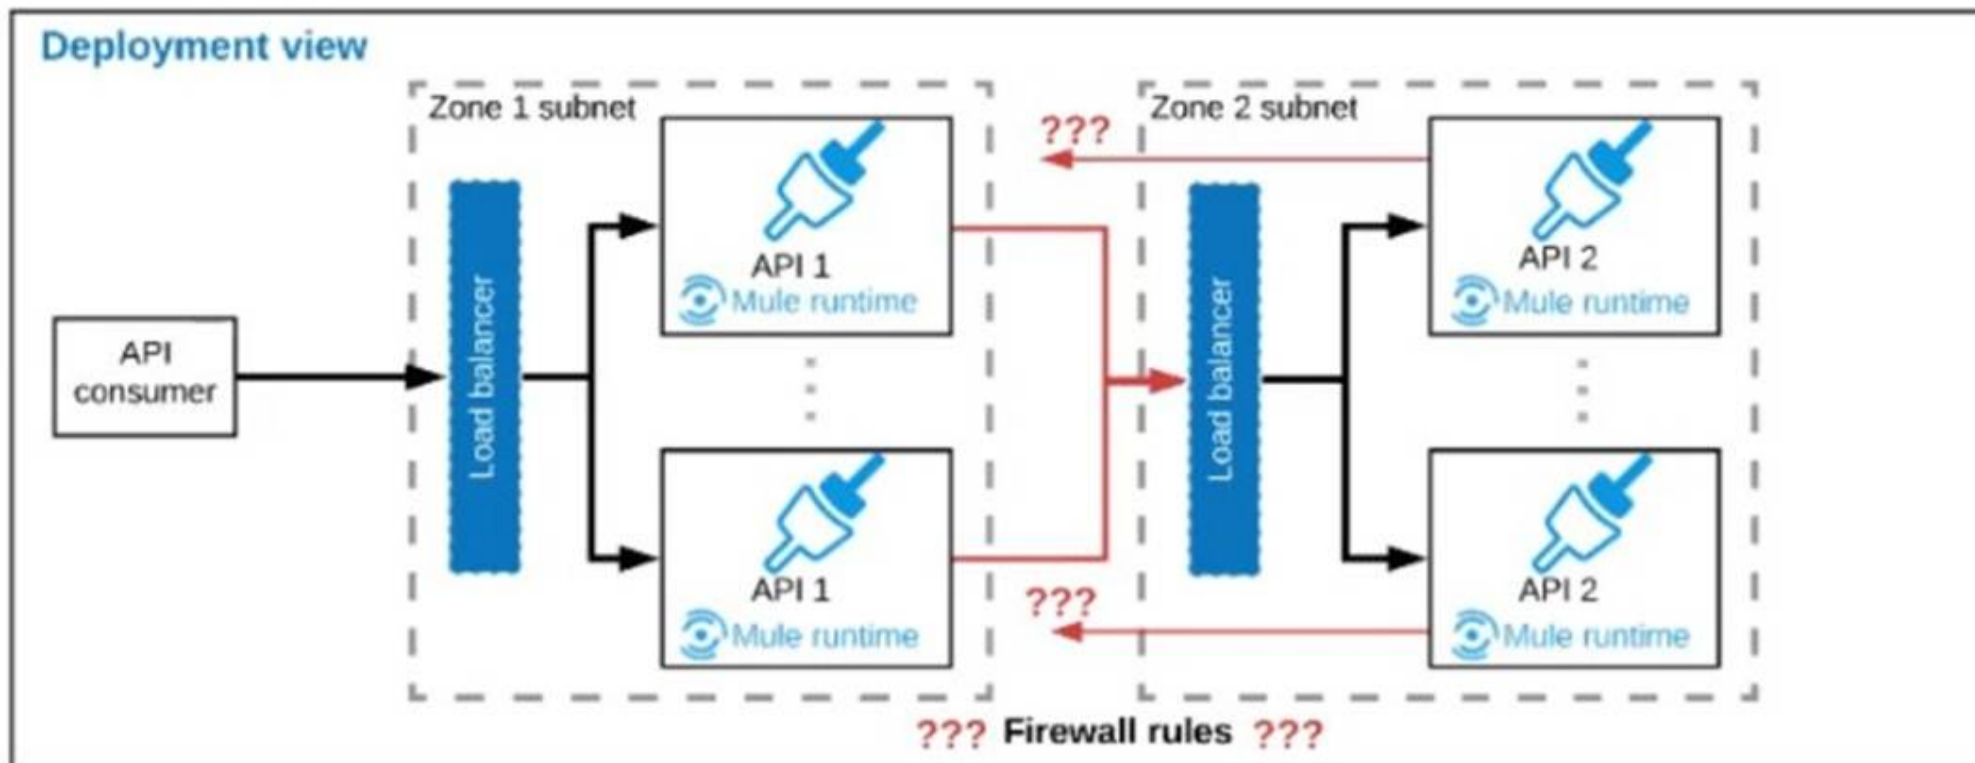
- A. Use a Parallel for Each scope to Insert records one by one into the database.
- B. Use a Scatter-Gather to bulk insert records into the database.
- C. Use a Batch job scope to bulk insert records into the database.
- D. Use a DataWeave map operation and an Async scope to insert records one by one into the database.

Answer: C

NEW QUESTION 82

Refer to the exhibit.





A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API 1 to notify about completion of the asynchronous process.

Each API is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone. In the network architecture, how must the firewall rules be configured to enable the above Interaction between API 1 and API 2?

- A. To authorize the certificate to be used both APIs
- B. To enable communication from each API's Mule Runtimes and Network zone to the load balancer of the other API
- C. To open direct two-way communication between the Mule Runtimes of both API's
- D. To allow communication between load balancers used by each API

Answer: B

NEW QUESTION 85

In Anypoint Platform, a company wants to configure multiple identity providers (IdPs) for various lines of business (LOBs). Multiple business groups and environments have been defined for these LOBs. What Anypoint Platform feature can use multiple IdPs to access the company's business groups and environment?

- A. User management
- B. Roles and permissions
- C. Dedicated load balancers
- D. Client Management

Answer: D

NEW QUESTION 86

Which type of communication is managed by a service mesh in a microservices architecture?

- A. Communication between microservices runtime administrators
- B. Communication between microservices developers
- C. Communication between microservices
- D. Communication between trading partner services

Answer: C

NEW QUESTION 88

A finance giant is planning to migrate all its Mule applications to Runtime fabric (RTF). Currently all Mule applications are deployed cloud hub using automated CI/CD scripts.

As an integration architect, which of the below steps would you suggest to ensure that the applications from cloudhub are migrated properly to Runtime Fabric (RTF) with an assumption that organization is keen on keeping the same deployment strategy.

- A. No changes need to be made to POM.xml file and CI/CD script should be modified as per the RTF configurations
- B. runtimeFabric dependency should be added as a mule plug-in to POM.xml file and CI/CD script should be modified as per the RTF configurations
- C. runtimeFabric deployment should be added to POM.xml file in all the mule applications and CI/CD script should be modified as per the RTF configurations
- D. runtimeFabric profile should be added to mule configuration files in the mule applications and CI/CD script should be modified as per the RTF configurations

Answer: C

NEW QUESTION 91

An organization is designing a Mule application to periodically poll an SFTP location for new files containing sales order records and then process those sales orders. Each sales order must be processed exactly once.

To support this requirement, the Mule application must identify and filter duplicate sales orders on the basis of a unique ID contained in each sales order record and then only send the new sales orders to the downstream system.

What is the most idiomatic (used for its intended purpose) Anypoint connector, validator, or scope that can be configured in the Mule application to filter duplicate sales orders on the basis of the unique ID field contained in each sales order record?

- A. Configure a Cache scope to filter and store each record from the received file by the order ID
- B. Configure a Database connector to filter and store each record by the order ID
- C. Configure an Idempotent Message Validator component to filter each record by the order ID
- D. Configure a watermark In an On New or Updated File event source to filter unique records by the order ID

Answer: C

NEW QUESTION 94

A Mule application is synchronizing customer data between two different database systems.

What is the main benefit of using eXtended Architecture (XA) transactions over local transactions to synchronize these two different database systems?

- A. An XA transaction synchronizes the database systems with the least amount of Mule configuration or coding
- B. An XA transaction handles the largest number of requests in the shortest time
- C. An XA transaction automatically rolls back operations against both database systems if any operation fails
- D. An XA transaction writes to both database systems as fast as possible

Answer: B

NEW QUESTION 99

An organization has several APIs that accept JSON data over HTTP POST. The APIs are all publicly available and are associated with several mobile applications and web applications. The organization does NOT want to use any authentication or compliance policies for these APIs, but at the same time, is worried that some bad actor could send payloads that could somehow compromise the applications or servers running the API implementations. What out-of-the-box Anypoint Platform policy can address exposure to this threat?

- A. Apply a Header injection and removal policy that detects the malicious data before it is used
- B. Apply an IP blacklist policy to all APIs; the blacklist will include all bad actors
- C. Shut out bad actors by using HTTPS mutual authentication for all API invocations
- D. Apply a JSON threat protection policy to all APIs to detect potential threat vectors

Answer: D

NEW QUESTION 104

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 108

What is an example of data confidentiality?

- A. Signing a file digitally and sending it using a file transfer mechanism
- B. Encrypting a file containing personally identifiable information (PII)
- C. Providing a server's private key to a client for secure decryption of data during a two-way SSL handshake
- D. De-masking a person's Social Security number while inserting it into a database

Answer: B

NEW QUESTION 110

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
- C. 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
- D. Order messages are sent directly to the Fulfillment microservice
- E. 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
- F. Order messages are sent to a JMS queue
- G. 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: B

NEW QUESTION 112

According to MuleSoft, which deployment characteristic applies to a microservices application architecture?

- A. Services exist as independent deployment artifacts and can be scaled -independently of other services
- B. All services of an application can be deployed together as single Java WAR file
- C. A deployment to enhance one capability requires a redeployment of all capabilities
- D. Core business capabilities are encapsulated in a single, deployable application

Answer: A

NEW QUESTION 114

An organization will deploy Mule applications to Cloudhub, Business requirements mandate that all application logs be stored ONLY in an external splunk consolidated logging service and NOT in Cloudhub.

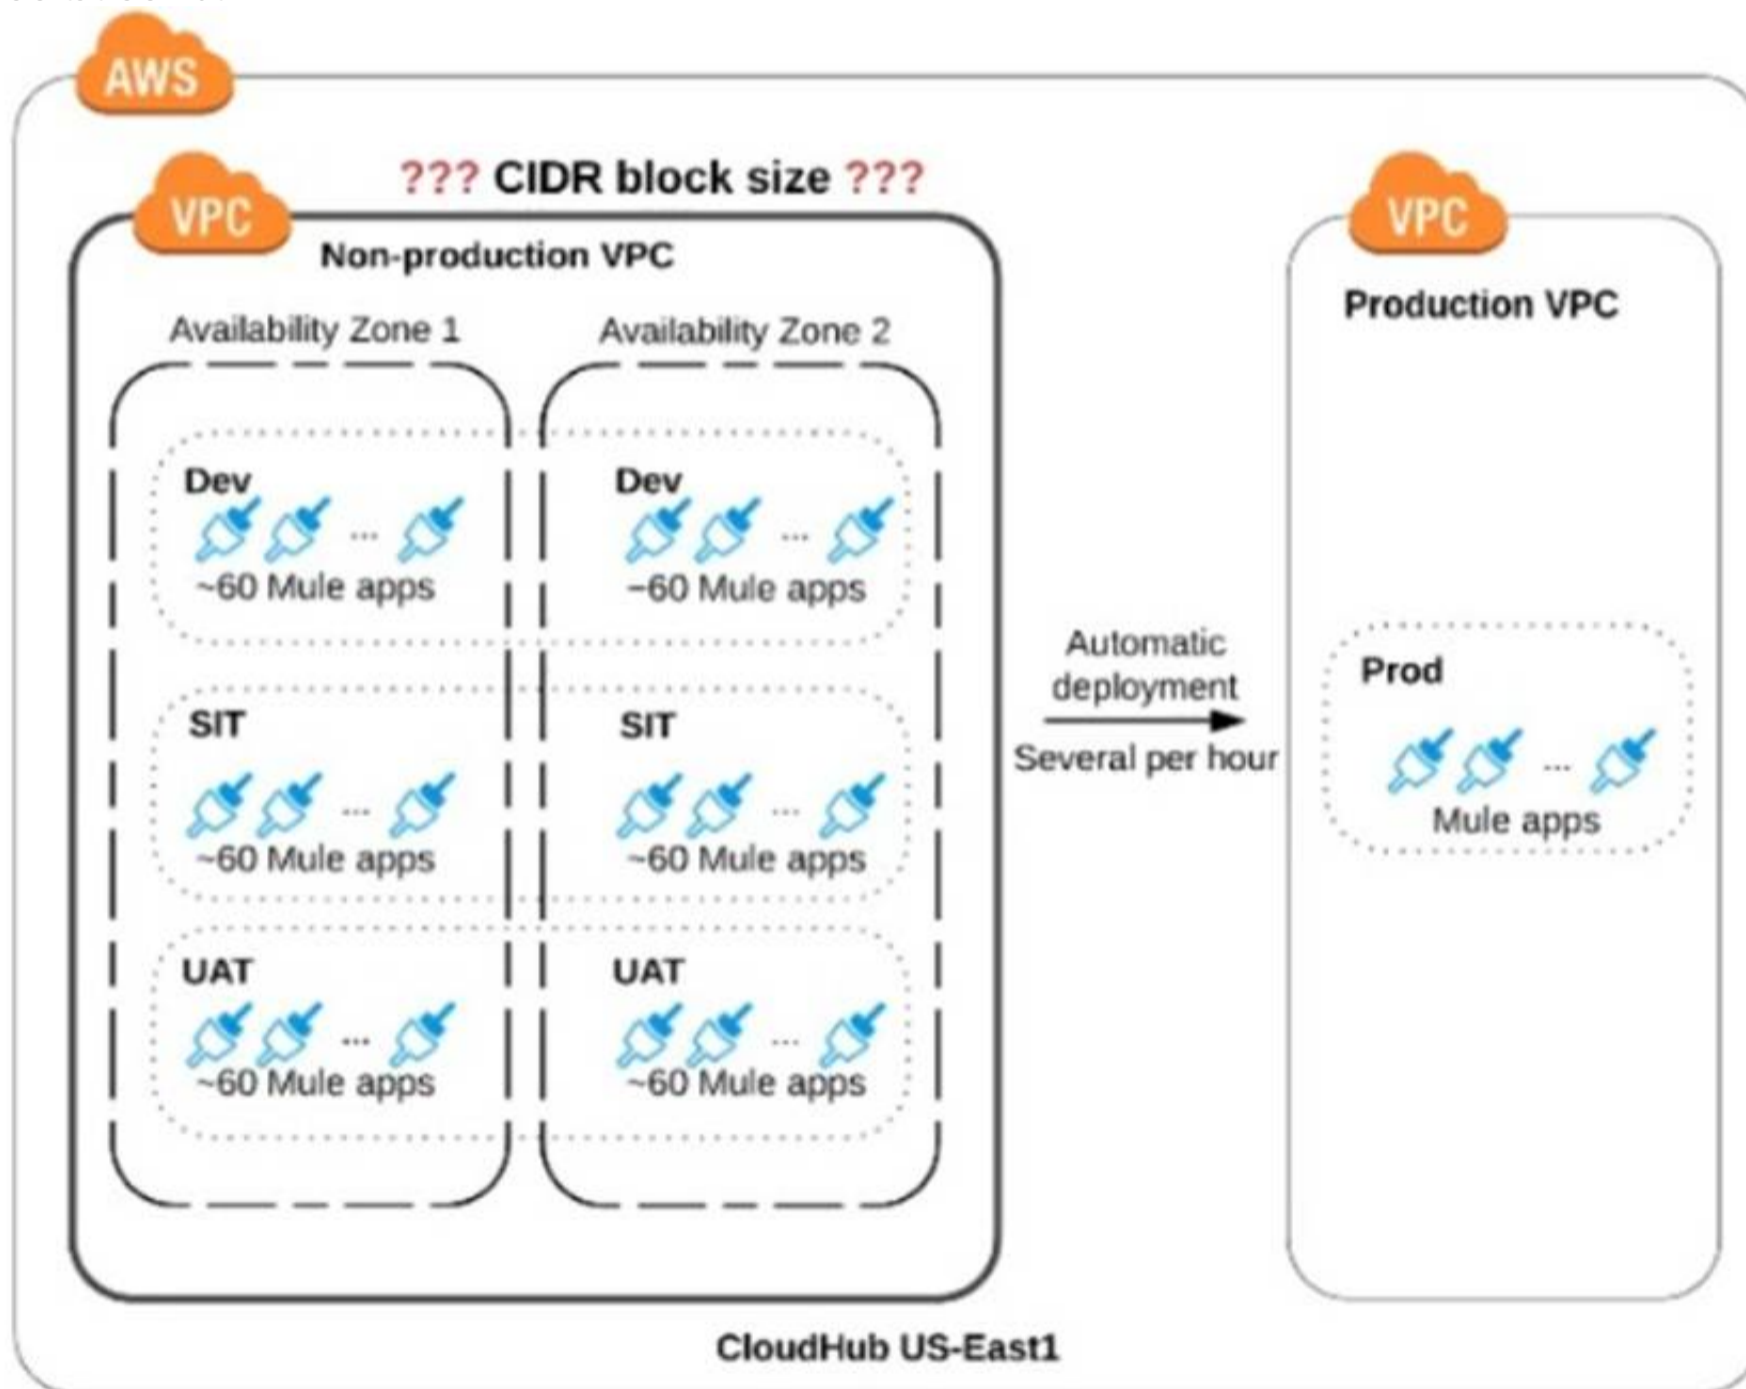
In order to most easily store Mule application logs ONLY in Splunk, how must Mule application logging be configured in Runtime Manager, and where should the log4j2 splunk appender be defined?

- A. Keep the default logging configuration in RuntimeManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support at Mule application deployments.
- B. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in EACH Mule application??s log4j2.xml file
- C. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manger to support at Mule application deployments.
- D. Keep the default logging configuration in Runtime Manager Define the Splunk appender in EACH Mule application log4j2.xml file

Answer: B

NEW QUESTION 119

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 123

In which order are the API Client, API Implementation, and API interface components called in a typical REST request?

- A. API Client > API implementation > API Interface
- B. API interface > API Client > API Implementation
- C. API Client > API Interface > API implementation
- D. API Implementation > API Interface > API Client

Answer: C

NEW QUESTION 125

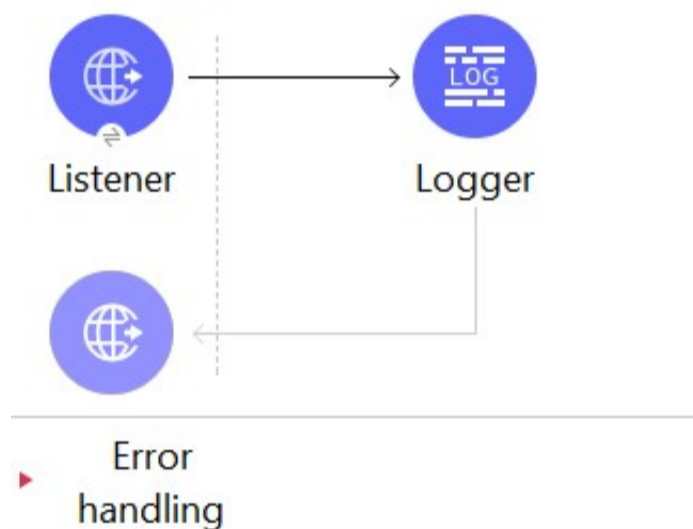
Why would an Enterprise Architect use a single enterprise-wide canonical data model (CDM) when designing an integration solution using Anypoint Platform?

- A. To reduce dependencies when integrating multiple systems that use different data formats
- B. To automate AI-enabled API implementation generation based on normalized backend databases from separate vendors
- C. To leverage a data abstraction layer that shields existing Mule applications from nonbackward compatible changes to the model's data structure
- D. To remove the need to perform data transformation when processing message payloads in Mule applications

Answer: A

NEW QUESTION 128

Refer to the exhibit.



The HTTP Listener and the Logger are being handled from which thread pools respectively?

- A. CPU_INTENSIVE and Dedicated Selector pool
- B. UBER and NONBLOCKING
- C. Shared Selector Pool and CPU LITE
- D. BLOCKING_IO and UBER

Answer: C

NEW QUESTION 133

A Mule application contains a Batch Job scope with several Batch Step scopes. The Batch Job scope is configured with a batch block size of 25. A payload with 4,000 records is received by the Batch Job scope.

When there are no errors, how does the Batch Job scope process records within and between the Batch Step scopes?

- A. The Batch Job scope processes multiple record blocks in parallel, and a block of 25 records can jump ahead to the next Batch Step scope over an earlier block of records. Each Batch Step scope is invoked with one record in the payload of the received Mule event. For each Batch Step scope, all 25 records within a block are processed in parallel. All the records in a block must be completed before the block of 25 records is available to the next Batch Step scope.
- B. The Batch Job scope processes each record block sequentially, one at a time. Each Batch Step scope is invoked with one record in the payload of the received Mule event. For each Batch Step scope, all 25 records within a block are processed sequentially, one at a time. All 4000 records must be completed before the blocks of records are available to the next Batch Step scope.
- C. The Batch Job scope processes multiple record blocks in parallel, and a block of 25 records can jump ahead to the next Batch Step scope over an earlier block of records. Each Batch Step scope is invoked with one record in the payload of the received Mule event. For each Batch Step scope, all 25 records within a block are processed sequentially, one record at a time. All the records in a block must be completed before the block of 25 records is available to the next Batch Step scope.
- D. The Batch Job scope processes multiple record blocks in parallel. Each Batch Step scope is invoked with a batch of 25 records in the payload of the received Mule event. For each Batch Step scope, all 4000 records are processed in parallel. Individual records can jump ahead to the next Batch Step scope before the rest of the records finish processing in the current Batch Step scope.

Answer: A

NEW QUESTION 136

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

Answer: D

NEW QUESTION 138

An integration Mule application is deployed to a customer-hosted multi-node Mule 4 runtime cluster. 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 142

An organization is evaluating using the CloudHub shared Load Balancer (SLB) vs creating a CloudHub dedicated load balancer (DLB). They are evaluating how this choice affects the various types of certificates used by CloudHub deployed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates. What type of restrictions exist on the types of certificates for the service that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

- A. Underlying Mule applications need to implement own certificates
- B. Only MuleSoft provided certificates can be used for server side certificate
- C. Only self signed certificates can be used
- D. All certificates which can be used in shared load balancer need to get approved by raising support ticket

Answer: B

NEW QUESTION 147

What requirement prevents using Anypoint MQ as the messaging broker for a Mule application?

- A. When the payload sent through the message broker must use XML format
- B. When the payload sent through the message broker must be encrypted
- C. When the messaging broker must support point-to-point messaging
- D. When the messaging broker must be deployed on-premises

Answer: D

NEW QUESTION 152

Which role is primarily responsible for building API implementation as part of a typical MuleSoft integration project?

- A. API Developer
- B. API Designer
- C. Integration Architect
- D. Operations

Answer: A

NEW QUESTION 153

As a part of business requirement, old CRM system needs to be integrated using Mule application. CRM system is capable of exchanging data only via SOAP/HTTP protocol. As an integration architect who follows API led approach, what is the below step you will perform so that you can share document with CRM team?

- A. Create RAML specification using Design Center
- B. Create SOAP API specification using Design Center
- C. Create WSDL specification using text editor
- D. Create WSDL specification using Design Center

Answer: C

NEW QUESTION 157

An architect is designing a Mule application to meet the following two requirements:

* 1. The application must process files asynchronously and reliably from an FTPS server to a back-end database using VM intermediary queues for load-balancing Mule events.

* 2. The application must process a medium rate of records from a source to a target system using a Batch Job scope.

To make the Mule application more reliable, the Mule application will be deployed to two CloudHub 1.0 workers.

Following MuleSoft-recommended best practices, how should the Mule application deployment typically be configured in Runtime Manager to best support the performance and reliability goals of both the Batch Job scope and the file processing VM queues?

- A. Check the Persistent VM queues checkbox in the application deployment configuration
- B. Check the Non-persistent VM queues checkbox in the application deployment configuration
- C. In the Runtime Manager Properties tab, disable persistent VM queues for Batch Job scopes
- D. In the Runtime Manager Properties tab, enable persistent VM queues for the FTPSconnector

Answer: A

NEW QUESTION 161

The company's FTPS server login username and password

- A. TLS context trust store containing a public certificate for the company
- B. The company's PGP public key that was used to sign the files
- C. The partner's PGP public key used by the company to login to the FTPS server
- D. A TLS context key store containing the private key for the companyThe partner's PGP private key that was used to sign the files
- E. The company's FTPS server login username and password
- F. A TLS context trust store containing a public certificate for ftps.partner.comThe partner's PGP public key that was used to sign the files
- G. The partner's PGP public key used by the company to login to the FTPS server
- H. A TLS context key store containing the private key for ftps.partner.comThe company's PGP private key that was used to sign the files

Answer: C

NEW QUESTION 163

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: B

NEW QUESTION 165

An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders. The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS). At a minimum, what must be stored in each truststore and keystore of these two Mule applications to properly support two-way TLS between the two Mule applications while properly protecting each Mule application's keys?

- A. Orders API truststore: The Orders API public keyProcess Orders keystore: The Process Orders private key and public key
- B. Orders API truststore: The Orders API private key and public key Process Orders keystore: The Process Orders private key public key
- C. Orders API truststore: The Process Orders public keyOrders API keystore: The Orders API private key and public key Process Orders truststore: The Orders API public keyProcess Orders keystore: The Process Orders private key and public key
- D. Orders API truststore: The Process Orders public key Orders API keystore: The Orders API private key Process Orders truststore: The Orders API public keyProcess Orders keystore: The Process Orders private key

Answer: C

NEW QUESTION 166

A marketing organization is designing a Mule application to process campaign data. The Mule application will periodically check for a file in a SFTP location and process the records in the file. The size of the file can vary from 10MB to 5GB. Due to the limited availability of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.

The application must transform and send different formats of this file to three different downstream SFTP locations.

What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

- A. Use an in-memory repeatable stream
- B. Use a file-stored non-repeatable stream
- C. Use an in-memory non-repeatable stream
- D. Use a file-stored repeatable stream

Answer: A

NEW QUESTION 171

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. EDIFACT 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: C

NEW QUESTION 176

How should the developer update the logging configuration in order to enable this package specific debugging?

- A. In Anypoint Monitoring, define a logging search query with class property set to org.apache.cxf and level set to DEBUG
- B. In the Mule application's log4j2.xml file, add an AsyncLogger element with name property set to org.apache.cxf and level set to DEBUG, then redeploy the Mule application in the CloudHub production environment
- C. In the Mule application's log4j2.xml file, change the root logger's level property to DEBUG, then redeploy the Mule application to the CloudHub production environment
- D. In Anypoint Runtime Manager, in the Deployed Application Properties tab for the Mule application, add a line item with DEBUG level for package org.apache.cxf

and apply the changes

Answer: A

NEW QUESTION 177

A set of integration Mule applications, some of which expose APIs, are being created to enable a new business process. Various stakeholders may be impacted by this. These stakeholders are a combination of semi-technical users (who understand basic integration terminology and concepts such as JSON and XML) and technically skilled potential consumers of the Mule applications and APIs.

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

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

Answer: B

NEW QUESTION 181

An integration team follows MuleSoft's recommended approach to full lifecycle API development. Which activity should this team perform during the API implementation phase?

- A. Validate the API specification
- B. Use the API specification to build the MuleSoft application
- C. Design the API specification
- D. Use the API specification to monitor the MuleSoft application

Answer: B

NEW QUESTION 186

An airline is architecting an API connectivity project to integrate its flight data into an online aggregation website. The interface must allow for secure communication high-performance and asynchronous message exchange.

What are suitable interface technologies for this integration assuming that Mulesoft fully supports these technologies and that Anypoint connectors exist for these interfaces?

- A. AsyncAPI over HTTPS AMQP with RabbitMQ JSON/REST over HTTPS
- B. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
- C. CSV over FTP YAM L over TLS JSON over HTTPS
- D. SOAP over HTTPS HOP over TLS gRPC over HTTPS

Answer: A

NEW QUESTION 187

An organization is designing an integration Mule application to process orders by submitting them to a back-end 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 back-end system. Orders that cannot be successfully submitted due to rejections from the back-end system will need to be processed manually (outside the back-end system).

The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization's firewall. The back-end system has a track record of unreliability due to both 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 back-end system while supporting but minimizing manual order processing?

- A. An Until Successful scope to call the back-end system One or more ActiveMQ long-retry queues One or more ActiveMQ dead-letter queues for manual processing
- B. One or more On Error scopes to assist calling the back-end system An Until Successful scope containing VM components for long retries A persistent dead-letter VM queue configured in CloudHub
- C. One or more On Error scopes to assist calling the back-end system One or more ActiveMQ long-retry queues A persistent dead-letter object store configured in the CloudHub Object Store service
- D. A Batch Job scope to call the back-end system An Until Successful scope containing Object Store components for long retries A dead-letter object store configured in the Mule application

Answer: A

NEW QUESTION 190

An API client makes an HTTP request to an API gateway with an Accept header containing the value application/json.

What is a valid HTTP response payload for this request in the client requested data format?

- A. <status>healthy</status>
- B. {"status": "healthy"}
- C. status(health)
- D. status: healthy

Answer: B

NEW QUESTION 192

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

Answer: A

NEW QUESTION 193

An organization's IT team must secure all of the internal APIs within an integration solution by using an API proxy to apply required authentication and authorization policies.

Which integration technology, when used for its intended purpose, should the team choose to meet these requirements if all other relevant factors are equal?

- A. API Management (APIM)
- B. Robotic Process Automation (RPA)
- C. Electronic Data Interchange (EDI)
- D. Integration Platform-as-a-service (PaaS)

Answer: A

NEW QUESTION 194

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 BOTH DB inserts in ONE DB transaction 3) Acknowledge the JMS message
- B. 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
- 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 and acknowledge the JMS message (NOT in an XA transaction) 2) In a NEW XA transaction, perform BOTH DB inserts

Answer: A

NEW QUESTION 199

A company is designing a mule application to consume batch data from a partner's ftps server. The data files have been compressed and then digitally signed using PGP. What inputs are required for the application to securely consume these files?

- A. ATLS context Key Store requiring the private key and certificate for the company PGP public key of partner PGP private key for the company
- B. ATLS context first store containing a public certificate for partner ftps server and the PGP public key of the partner TLS contact Key Store containing the FTP credentials
- C. TLS context trust or containing a public certificate for the ftps server The FTP username and password The PGP public key of the partner
- D. The PGP public key of the partner The PGP private key for the company The FTP username and password

Answer: D

NEW QUESTION 203

An organization has chosen Mulesoft for their integration and API platform.

According to the Mulesoft catalyst framework, what would an integration architect do to create achievement goals as part of their business outcomes?

- A. Measure the impact of the centre for enablement
- B. build and publish foundational assets
- C. agree upon KPI's and help develop and overall success plan
- D. evangelize API's

Answer: C

NEW QUESTION 207

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 RECORDS are processed IN PARALLEL within and between the two Batch Steps
- B. 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
- C. 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
- D. 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

Answer: A

NEW QUESTION 211

What approach configures an API gateway to hide sensitive data exchanged between API consumers and API implementations, but can convert tokenized fields

back to their original value for other API requests or responses, without having to recode the API implementations?

- A. Create both masking and tokenization formats and use both to apply a tokenization policy in an API gateway to mask sensitive values in message payloads with characters, and apply a corresponding detokenization policy to return the original values to other APIs
- B. Create a masking format and use it to apply a tokenization policy in an API gateway to mask sensitive values in message payloads with characters, and apply a corresponding detokenization policy to return the original values to other APIs
- C. Use a field-level encryption policy in an API gateway to replace sensitive fields in message payload with encrypted values, and apply a corresponding field-level decryption policy to return the original values to other APIs
- D. Create a tokenization format and use it to apply a tokenization policy in an API gateway to replace sensitive fields in message payload with similarly formatted tokenized values, and apply a corresponding detokenization policy to return the original values to other APIs

Answer: D

NEW QUESTION 214

A leading bank implementing new mule API.

The purpose of API to fetch the customer account balances from the backend application and display them on the online platform the online banking platform. The online banking platform will send an array of accounts to Mule API get the account balances.

As a part of the processing the Mule API needs to insert the data into the database for auditing purposes and this process should not have any performance related implications on the account balance retrieval flow

How should this requirement be implemented to achieve better throughput?

- A. Implement the Async scope fetch the data from the backend application and to insert records in the Audit database
- B. Implement a for each scope to fetch the data from the back-end application and to insert records into the Audit database
- C. Implement a try-catch scope to fetch the data from the back-end application and use the Async scope to insert records into the Audit database
- D. Implement parallel for each scope to fetch the data from the backend application and use Async scope to insert the records into the Audit database

Answer: C

NEW QUESTION 217

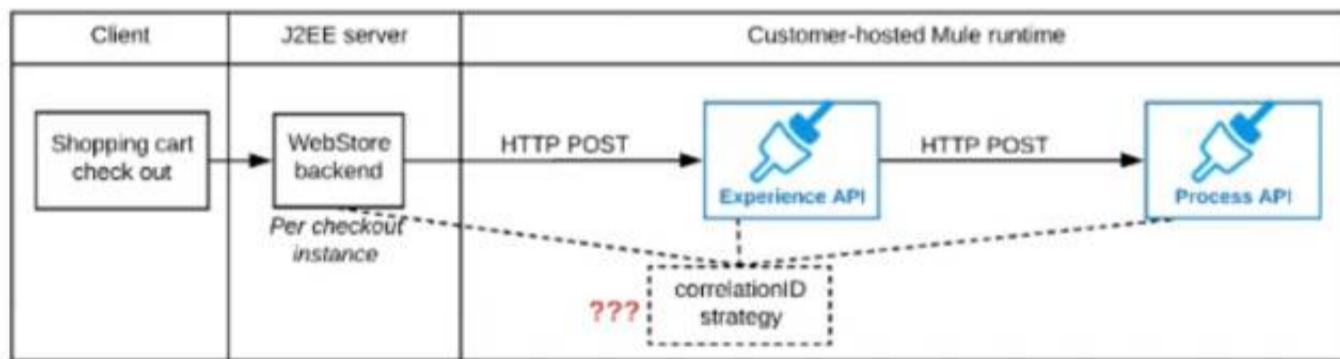
What operation can be performed through a JMX agent enabled in a Mule application?

- A. View object store entries
- B. Replay an unsuccessful message
- C. Set a particular tog4J2 log level to TRACE
- D. Deploy a Mule application

Answer: A

NEW QUESTION 221

Refer to the exhibit.

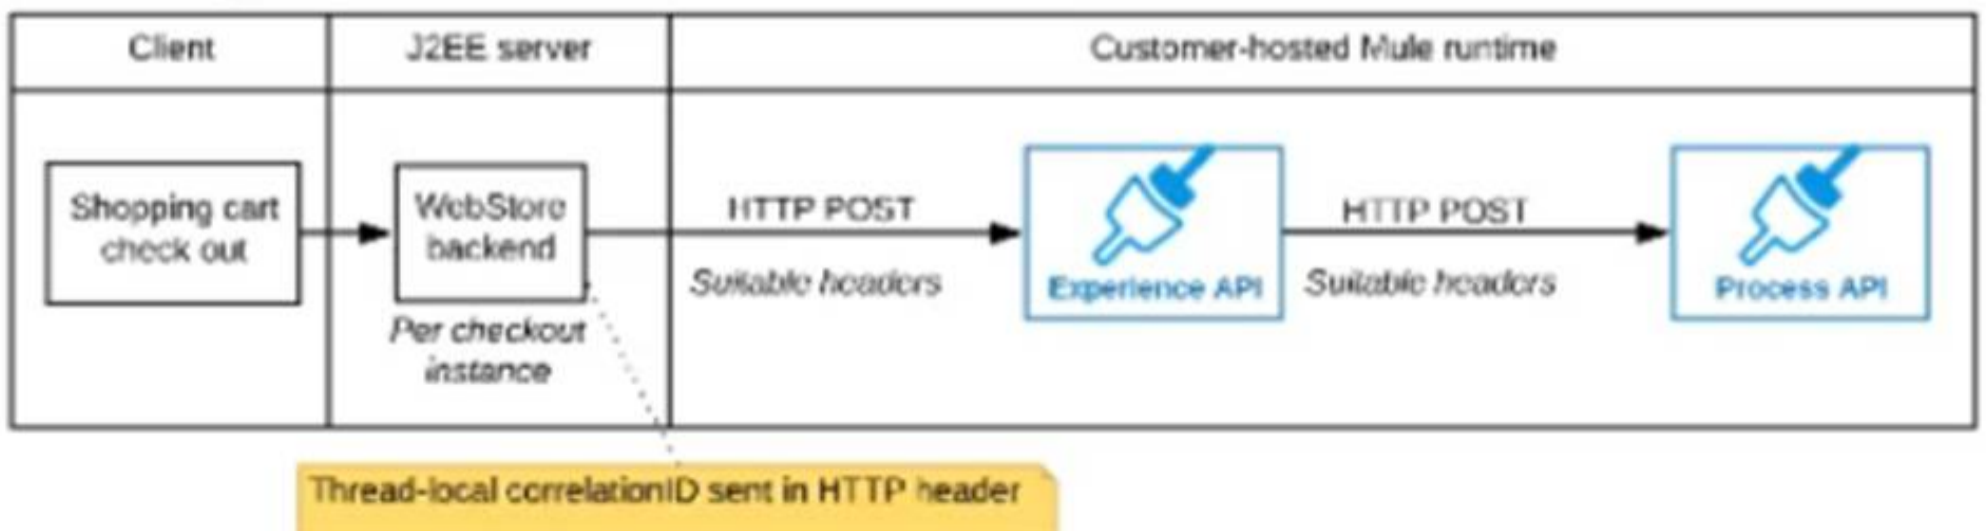


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

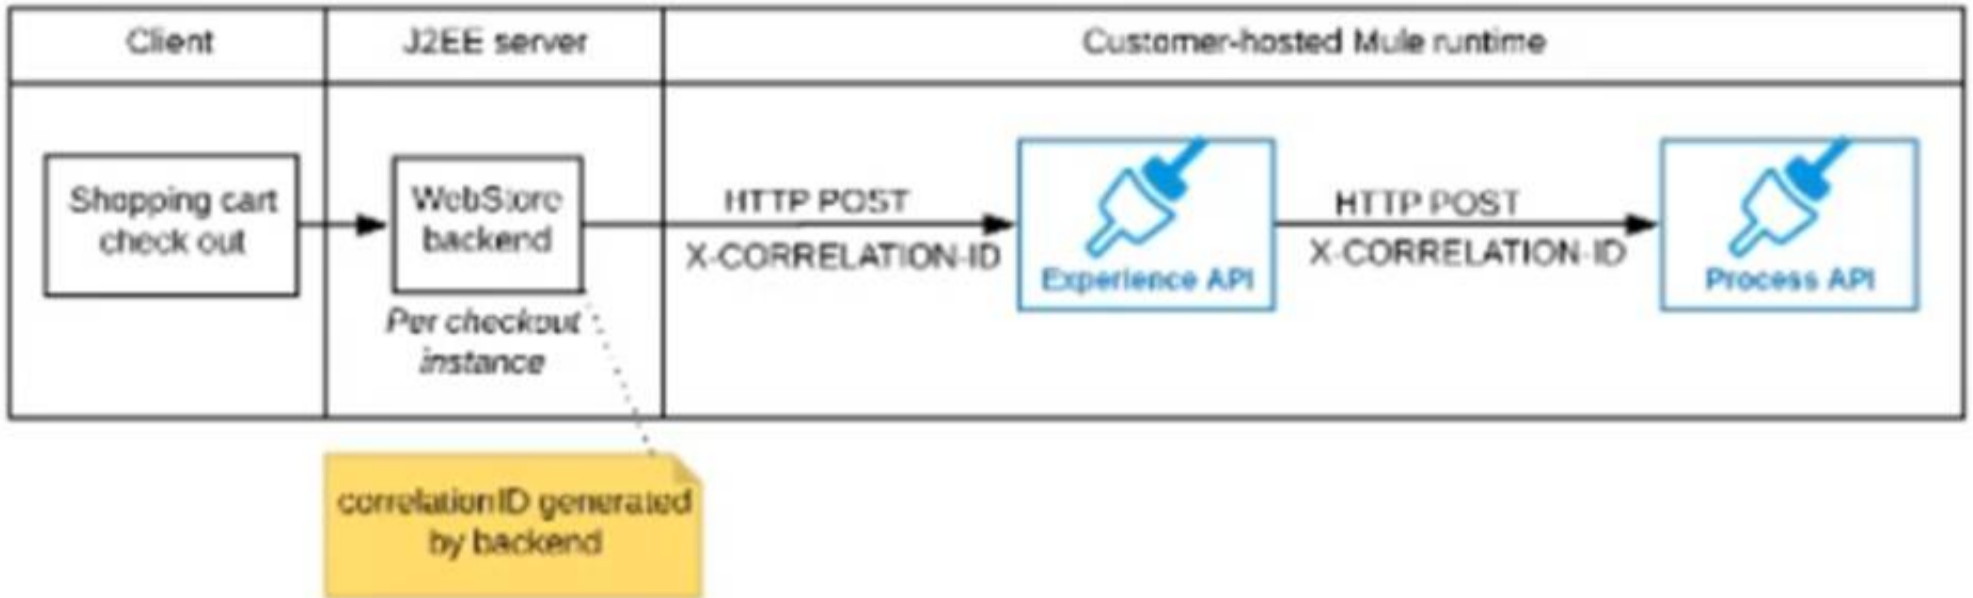
End-to-end correlation of all HTTP requests and responses belonging to each individual checkout Instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the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

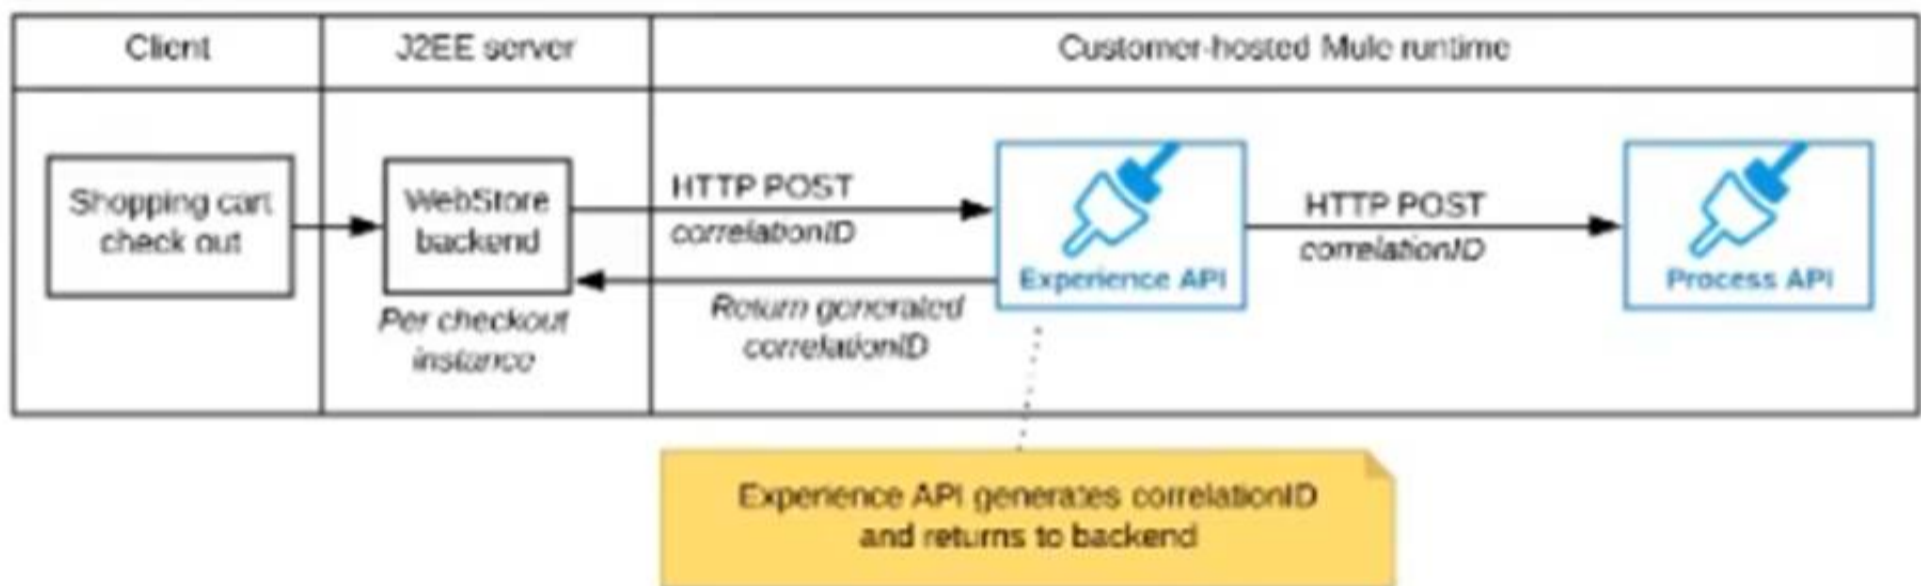
- A) The web store backend, being a Java EE application, automatically makes use of the thread-local correlation ID generated by the Java EE application server and automatically transmits that to the Experience API using HTTP-standard headers
- No special code or configuration is included in the web store backend, Experience API, and Process API implementations to generate and manage the correlation ID



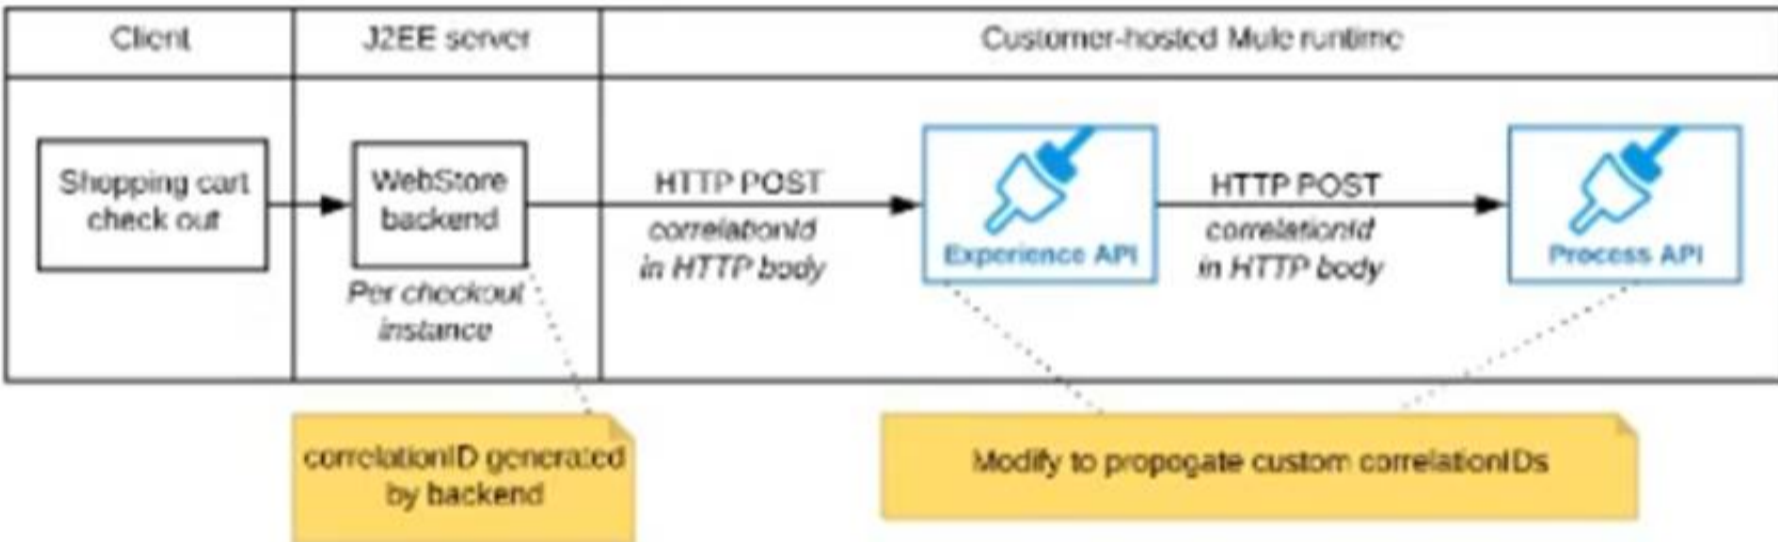
B) The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header In each API invocation belonging to that checkout
 No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID



C) The 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



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



- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

NEW QUESTION 224

An integration architect is designing an API that must accept requests from API clients for both XML and JSON content over HTTP/1.1 by default. Which API architectural style, when used for its intended and typical purposes, should the architect choose to meet these requirements?

- A. SOAP
- B. GraphQL
- C. REST
- D. grRPC

Answer: C

NEW QUESTION 227

A project team is working on an API implementation using the RAML definition as a starting point. The team has updated the definition to include new operations and has published a new version to exchange. Meanwhile another team is working on a mule application consuming the same API implementation. During the development what has to be performed by the mule application team to take advantage of the newly added operations?

- A. Scaffold the client application with the new definition
- B. Scaffold API implementation application with the new definition
- C. Update the REST connector from exchange in the client application
- D. Update the API connector in the API implementation and publish to exchange

Answer: C

NEW QUESTION 228

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: Active MQ CRM: REST
- D. IBM: CICS CRM: SOAP

Answer: D

NEW QUESTION 233

A company is implementing a new Mule application that supports a set of critical functions driven by a rest API enabled, claims payment rules engine hosted on oracle ERP. As designed the mule application requires many data transformation operations as it performs its batch processing logic.

The company wants to leverage and reuse as many of its existing java-based capabilities (classes, objects, data model etc.) as possible

What approach should be considered when implementing required data mappings and transformations between Mule application and Oracle ERP in the new Mule application?

- A. Create a new metadata RAML classes in Mule from the appropriate Java objects and then perform transformations via Dataweave
- B. From the mule application, transform via theXSLT model
- C. Transform by calling any suitable Java class from Dataweave
- D. Invoke any of the appropriate Java methods directly, create metadata RAML classes and then perform required transformations via Dataweave

Answer: C

NEW QUESTION 238

What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

- A. To send Mule application log entries to Splunk
- B. To change tog4j2 tog levels in Anypoint Runtime Manager without having to restart the Mule application
- C. To log certain messages to a custom log category
- D. To directly reference one shared and customized log4j2.xml file from multiple Mule applications

Answer: D

NEW QUESTION 241

An organization is choosing between API-led connectivity and other integration approaches.

According to MuleSoft, which business benefits is associated with an API-led connectivity approach using Anypoint Platform?

- A. improved security through adoption of monolithic architectures
- B. Increased developer productivity through self-service of API assets
- C. Greater project predictability through tight coupling of systems
- D. Higher outcome repeatability through centralized development

Answer: B

NEW QUESTION 242

As a part of project requirement, Java Invoke static connector in a mule 4 application needs to invoke a static method in a dependency jar file. What are two ways to add the dependency to be visible by the connectors class loader?

(Choose two answers)

- A. In the Java Invoke static connector configuration, configure a path and name of the dependency jar file
- B. Add the dependency jar file to the java classpath by setting the JVM parameters
- C. Use Maven command to include the dependency jar file when packaging the application
- D. Configure the dependency as a shared library in the project POM
- E. Update mule-artefact.json to export the Java package

Answer: CD

NEW QUESTION 243

A mule application designed to fulfil two requirements

a) Processing files are synchronously from an FTPS server to a back-end database using VM intermediary queues for load balancing VM events


b) Processing a medium rate of records from a source to a target system using batch job scope
 Considering the processing reliability requirements for FTPS files, how should VM queues be configured for processing files as well as for the batch job scope if the application is deployed to Cloudhub workers?

- A. Use Cloud hub persistent queues for FTPS files processing There is no need to configure VM queues for the batch jobs scope as it uses by default the worker's disc for VM queueing
- B. Use Cloud hub persistent VM queue for FTPS file processing There is no need to configure VM queues for the batch jobs scope as it uses by default the worker's JVM memory for VM queueing
- C. Use Cloud hub persistent VM queues for FTPS file processing Disable VM queue for the batch job scope
- D. Use VM connector persistent queues for FTPS file processing Disable VM queue for the batch job scope

Answer: A

NEW QUESTION 247

An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?



- A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
- B. One set of APIs (Experience API, Process API, and System API) for the web app, and another set for the mobile app.
- C. Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system
- D. A common Experience API used by both the web and mobile apps, but separate Process APIs for the web and mobile apps that interact with the database and the CRM System.

Answer: C

NEW QUESTION 252

A stock trading company handles millions of trades a day and requires excellent performance and reliability within its stock trading system. The company operates a number of event-driven APIs Implemented as Mule applications that are hosted on various customer-hosted Mule clusters and needs to enable message exchanges between the APIs within their internal network using shared message queues.

What is an effective way to meet the cross-cluster messaging requirements of its event- driven APIs?

- A. Non-transactional JMS operations with a reliability pattern and manual acknowledgements
- B. Persistent VM queues with automatic acknowledgements
- C. JMS transactions with automatic acknowledgements
- D. extended Architecture (XA) transactions and XA connected components with manual acknowledgements

Answer: C

NEW QUESTION 255

As a part of design , Mule application is required call the Google Maps API to perform a distance computation. The application is deployed to cloudhub. At the minimum what should be configured in the TLS context of the HTTP request configuration to meet these requirements?

- A. The configuration is built-in and nothing extra is required for the TLS context
- B. Request a private key from Google and create a PKCS12 file with it and add it in keyStore as a part of TLS context
- C. Download the Google public certificate from a browser, generate JKS file from it and add it in key store as a part of TLS context
- D. Download the Google public certificate from a browser, generate a JKS file from it and add it in Truststore as part of the TLS context

Answer: D

NEW QUESTION 256

A new upstream API Is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity. The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms. If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API's desired SLA?

- A. Set a timeout of 100 ms; that leaves 400 ms for the other two downstream APIs to complete
- B. Do not set a timeout; the Invocation of this API Is mandatory and so we must wait until it responds
- C. Set a timeout of 50 ms; this times out more invocations of that API but gives additional room for retries
- D. No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

Answer: D

NEW QUESTION 260

An insurance company is using a CloudHub runtime plane. As a part of requirement, email alert should be sent to internal operations team every time of policy applied to an API instance is deleted. As an integration architect suggest on how this requirement be met?

- A. Use audit logs in Anypoint platform to detect a policy deletion and configure the Audit logs alert feature to send an email to the operations team
- B. Use Anypoint monitoring to configure an alert that sends an email to the operations team every time a policy is deleted in API manager
- C. Create a custom connector to be triggered every time of policy is deleted in API manager
- D. Implement a new application that uses the Audit log REST API to detect the policy deletion and send an email to operations team the SMTP connector

Answer: D

NEW QUESTION 264

A project team uses RAML specifications to document API functional requirements and deliver API definitions. As per the current legal requirement, all designed API definitions to be augmented with an additional non-functional requirement to protect the services from a high rate of requests according to define service level agreements.

Assuming that the project is following Mulesoft API governance and policies, how should the project team convey the necessary non-functional requirement to stakeholders?

- A. Create proxies in API manager for the non functional requirement and publish to exchange
- B. Add all non functional requirements as comments to RAML specification and publish to exchange
- C. Create various SLA's in API manager for the non functional requirement and publish to exchange
- D. Update API definitions with the fragment for the appropriate policy and publish to exchange

Answer: D

NEW QUESTION 266

Mule application A receives a request Anypoint MQ message REQU with a payload containing a variable-length list of request objects. Application A uses the For Each scope to split the list into individual objects and sends each object as a message to an Anypoint MQ queue.

Service S listens on that queue, processes each message independently of all other messages, and sends a response message to a response queue.

Application A listens on that response queue and must in turn create and publish a response Anypoint MQ message RESP with a payload containing the list of responses sent by service S in the same order as the request objects originally sent in REQU.

Assume successful response messages are returned by service S for all request messages.

What is required so that application A can ensure that the length and order of the list of objects in RESP and REQU match, while at the same time maximizing message throughput?

- A. Use a Scatter-Gather within the For Each scope to ensure response message order. Configure the Scatter-Gather with a persistent object store
- B. Perform all communication involving service S synchronously from within the For Each scope, so objects in RESP are in the exact same order as request objects in REQU
- C. Use an Async scope within the For Each scope and collect response messages in a second For Each scope in the order in which they arrive, then send RESP using this list of responses
- D. Keep track of the list length and all object indices in REQU, both in the For Each scope and in all communication involving service S. Use persistent storage when creating RESP

Answer: D

NEW QUESTION 267

In preparation for a digital transformation initiative, an organization is reviewing related IT integration projects that failed for various reasons.

According to MuleSoft's surveys of global IT leaders, what is a common cause of IT project failure that this organization may likely discover in its assessment?

- A. Following an Agile delivery methodology
- B. Reliance on an Integration-Platform-as-a-Service (iPaaS)
- C. Spending too much time on enablement
- D. Lack of alignment around business outcomes

Answer: D

NEW QUESTION 269

A payment processing company has implemented a Payment Processing API Mule application to process credit card and debit card transactions. Because the Payment Processing API handles highly sensitive information, the payment processing company requires that data must be encrypted both In-transit and at-rest. To meet these security requirements, consumers of the Payment Processing API must create request message payloads in a JSON format specified by the API, and the message payload values must be encrypted.

How can the Payment Processing API validate requests received from API consumers?

- A. A Transport Layer Security (TLS) - Inbound policy can be applied in API Manager to decrypt the message payload and the Mule application implementation can then use the JSON Validation module to validate the JSON data
- B. The Mule application implementation can use the APIkit module to decrypt and then validate the JSON data
- C. The Mule application implementation can use the Validation module to decrypt and then validate the JSON data
- D. The Mule application implementation can use DataWeave to decrypt the message payload and then use the JSON Scheme Validation module to validate the JSON data

Answer: A

NEW QUESTION 272

An organization is building out a test suite for their application using MUnit.

The Integration Architect has recommended using Test Recorder in Anypoint Studio to record the processing flows and then configure unit tests based on the captured events.

What is a core consideration that must be kept in mind while using Test Recorder?

- A. The Recorder supports loops where the structure of the data being tested changes inside the Iteration
- B. Mocking values resulting from parallel processes are possible and will not affect the execution of the processors that follow in the test
- C. The Recorder supports mocking a message before or inside a Foreach processor
- D. Tests for flows cannot be created if Mule errors are raised inside the flows, even if the errors are handled by On-Error Continue error handlers

Answer: D

NEW QUESTION 275

An IT integration team followed an API-led connectivity approach to implement an order- fulfillment business process. It created an order processing AP that coordinates stateful interactions with a variety of microservices that validate, create, and fulfill new product orders. Which interaction composition pattern did the integration architect who designed this order processing AP use?

- A. Orchestration
- B. Streaming
- C. Aggregation
- D. Multicasting

Answer: A

NEW QUESTION 280

A system API EmployeeSAPI is used to fetch employee's data from an underlying SQL database.

The architect must design a caching strategy to query the database only when there is an update to the employees table or else return a cached response in order to minimize the number of redundant transactions being handled by the database.

What must the architect do to achieve the caching objective?

- A. Use an On Table Row on employees table and call invalidate cache Use an object store caching strategy and expiration interval to empty
- B. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and expiration interval to empty
- C. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and set expiration interval to 1-hour
- D. Use an on table rule on employees table call invalidate cache and save new employees data to cache Use an object store caching strategy and set expiration interval to 1-hour

Answer: A

NEW QUESTION 285

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

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

Answer: D

NEW QUESTION 286

A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction.

What is the behavior of the Scatter-Gather when running within this local transaction?

- A. Execution of each route within the Scatter-Gather occurs sequentially Any error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
- B. Execution of all routes within the Scatter-Gather occurs in parallel Any error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
- C. Execution of each route within the Scatter-Gather occurs sequentially Any error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations
- D. Execution of each route within the Scatter-Gather occurs in parallel Any error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations

Answer: A

NEW QUESTION 289

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

Answer: D

NEW QUESTION 291

What Anypoint Connectors support transactions?

- A. Database, JMS, VM
- B. Database, JMS, HTTP
- C. Database, JMS, VM, SFTP

D. Database, VM, File

Answer: A

NEW QUESTION 292

A travel company wants to publish a well-defined booking service API to be shared with its business partners. These business partners have agreed to ONLY consume SOAP services and they want to get the service contracts in an easily consumable way before they start any development. The travel company will publish the initial design documents to Anypoint Exchange, then share those documents with the business partners. When using an API-led approach, what is the first design document the travel company should deliver to its business partners?

- A. Create a WSDL specification using any XML editor
- B. Create a RAML API specification using any text editor
- C. Create an OAS API specification in Design Center
- D. Create a SOAP API specification in Design Center

Answer: A

NEW QUESTION 295

An Organization has previously provisioned its own AWS VPC hosting various servers. The organization now needs to use Cloudhub to host a Mule application that will implement a REST API once deployed to Cloudhub, this Mule application must be able to communicate securely with the customer-provisioned AWS VPC resources within the same region, without being interceptable on the public internet.

What Anypoint Platform features should be used to meet these network communication requirements between Cloudhub and the existing customer-provisioned AWS VPC?

- A. Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC
- B. Configure an external identity provider (IDP) in Anypoint Platform with certificates from the customer provisioned AWS VPC
- C. Add a default API Whitelisting policy to API Manager to automatically whitelist the customer provisioned AWS VPC IP ranges needed by the Mule application
- D. Use VM queues in the Mule application to allow any non-mule assets within the customer provisioned AWS VPC to subscribed to and receive messages

Answer: A

NEW QUESTION 299

An organization has just developed a Mule application that implements a REST API. The mule application will be deployed to a cluster of customer hosted Mule runtimes.

What additional infrastructure component must the customer provide in order to distribute inbound API requests across the Mule runtimes of the cluster?

- A. A message broker
- B. An HTTP Load Balancer
- C. A database
- D. An Object Store

Answer: B

NEW QUESTION 301

What API policy would LEAST likely be applied to a Process API?

- A. Custom circuit breaker
- B. Client ID enforcement
- C. Rate limiting
- D. JSON threat protection

Answer: D

NEW QUESTION 302

A mule application is being designed to perform product orchestration. The Mule application needs to join together the responses from an inventory API and a Product Sales History API with the least latency.

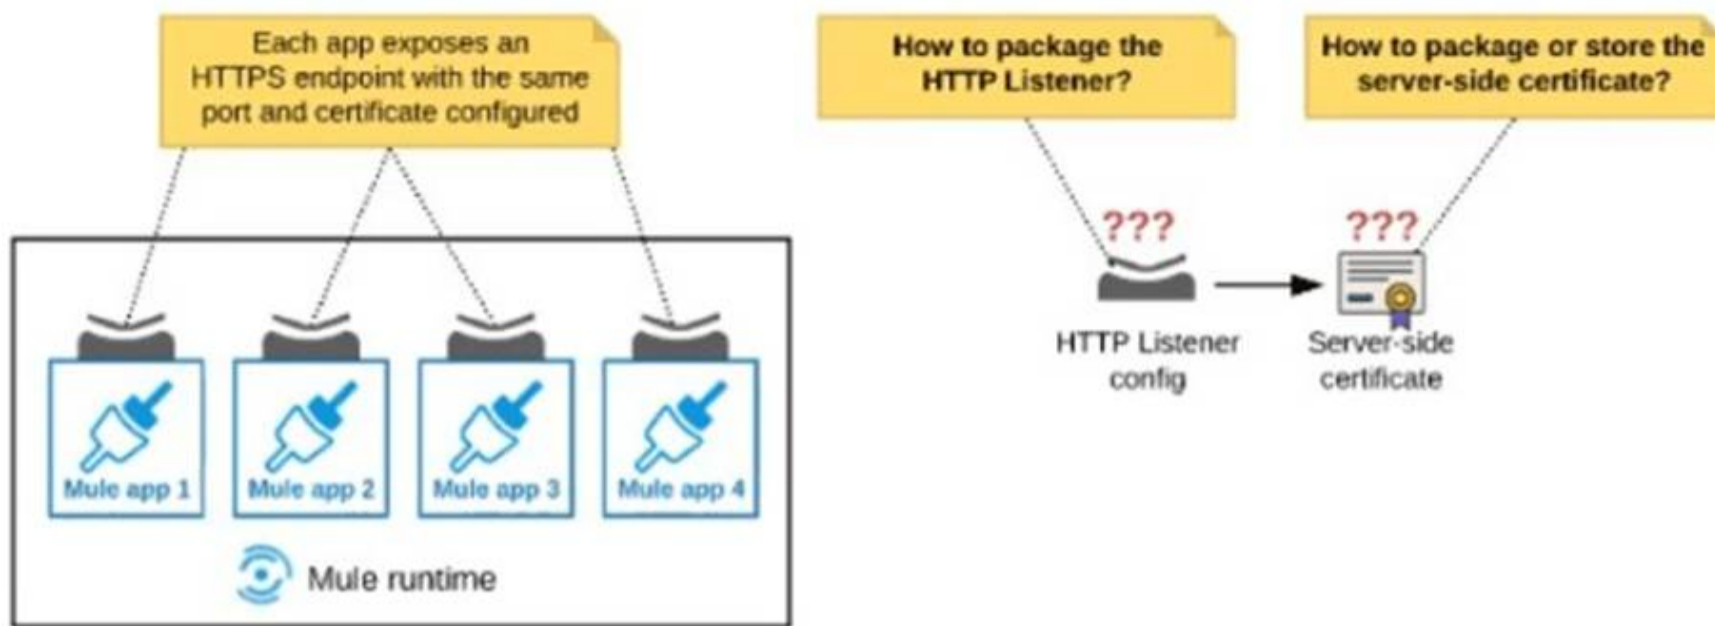
To minimize the overall latency. What is the most idiomatic (used for its intended purpose) design to call each API request in the Mule application?

- A. Call each API request in a separate lookup call from Dataweave reduce operator
- B. Call each API request in a separate route of a Scatter-Gather
- C. Call each API request in a separate route of a Parallel For Each scope
- D. Call each API request in a separate Async scope

Answer: B

NEW QUESTION 303

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

Answer: B

NEW QUESTION 308

A mule application is deployed to a Single Cloudhub worker and the public URL appears in Runtime Manager as the APP URL. Requests are sent by external web clients over the public internet to the mule application App url. Each of these requests routed to the HTTPS Listener event source of the running Mule application. Later, the DevOps team edits some properties of this running Mule application in Runtime Manager. Immediately after the new property values are applied in runtime manager, how is the current Mule application deployment affected and how will future web client requests to the Mule application be handled?

- A. Cloudhub will redeploy the Mule application to the OLD Cloudhub workerNew web client requests will RETURN AN ERROR until the Mule application is redeployed to the OLD Cloudhub worker
- B. CloudHub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests will RETURN AN ERROR until the NEW Cloudhub worker is available
- C. Cloudhub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker until the NEW Cloudhub worker is available.
- D. Cloudhub will redeploy the mule application to the OLD Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.

Answer: C

NEW QUESTION 311

Which Salesforce API is invoked to deploy, retrieve, create, update, or delete customization information, such as custom object definitions using Mule Salesforce Connectors in a Mule application?

- A. sObject Platform Action API
- B. User Interface API
- C. Metadata API
- D. Process Rules API

Answer: C

NEW QUESTION 313

What is maximum vCores can be allocated to application deployed to CloudHub?

- A. 1 vCores
- B. 2 vCores
- C. 4 vCores
- D. 16 vCores

Answer: D

NEW QUESTION 317

A company is using Mulesoft to develop API's and deploy them to Cloudhub and on premises targets. Recently it has decided to enable Runtime Fabric

deployment option as well and infrastructure is set up for this option.
 What can be used to deploy Runtime Fabric?

- A. AnypointCLI
- B. Anypoint platform REST API's
- C. Directly uploading ajar file from the Runtime manager
- D. Mule maven plug-in

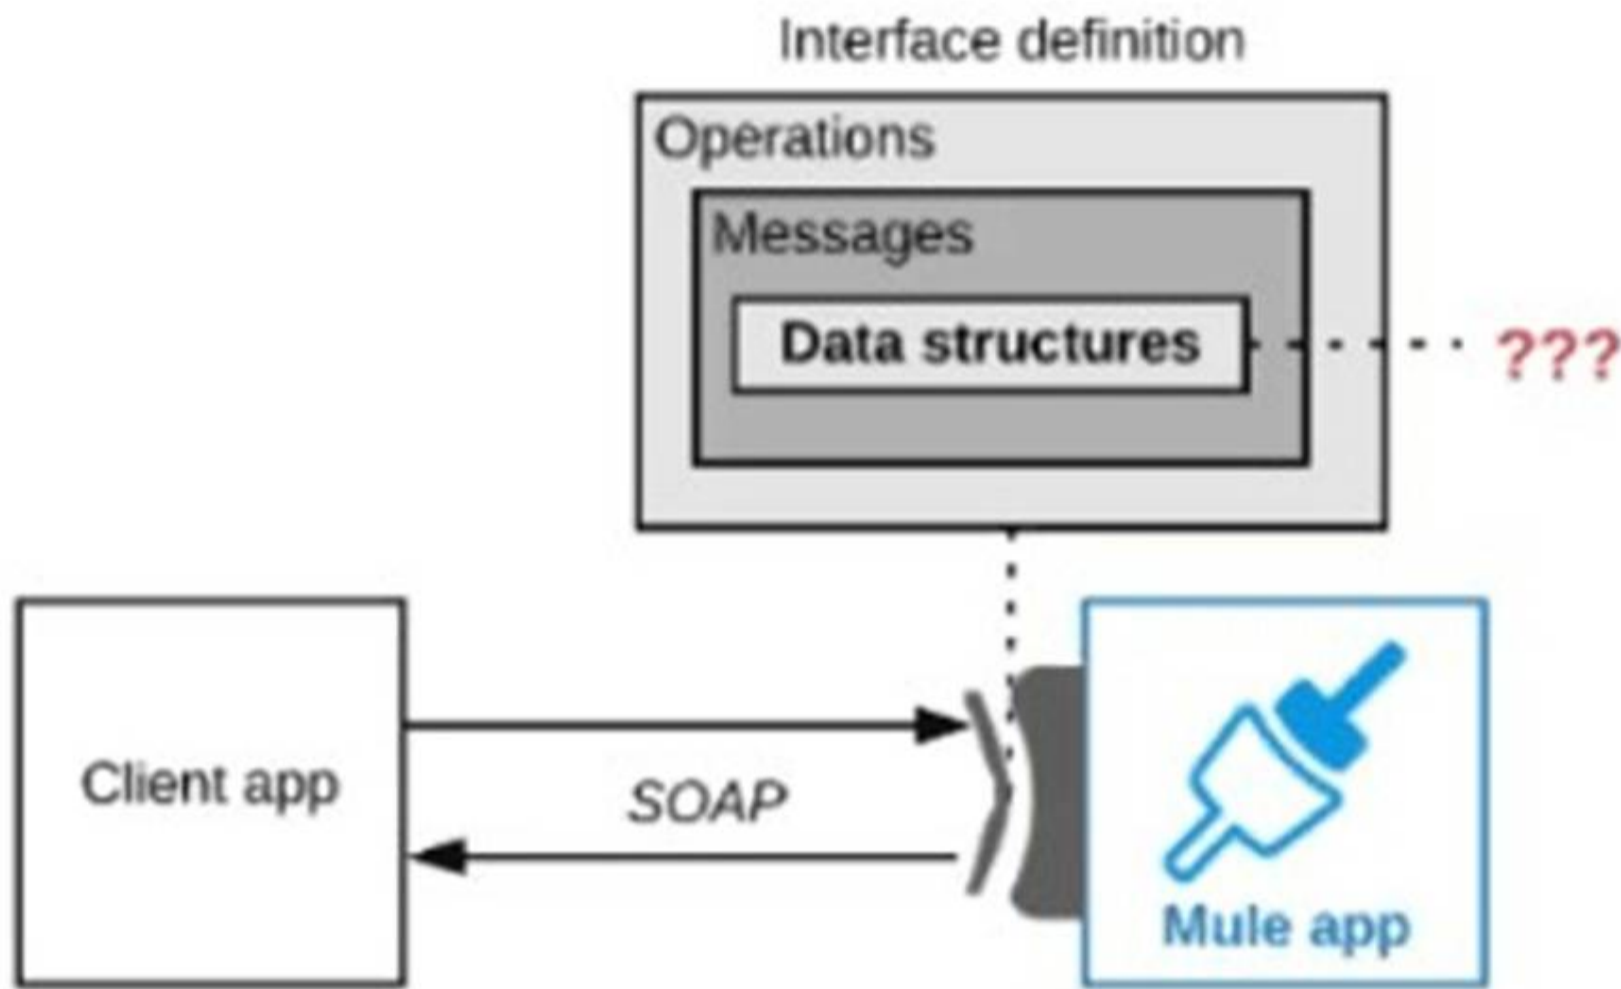
Answer: A

NEW QUESTION 318

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

Answer: B

NEW QUESTION 323

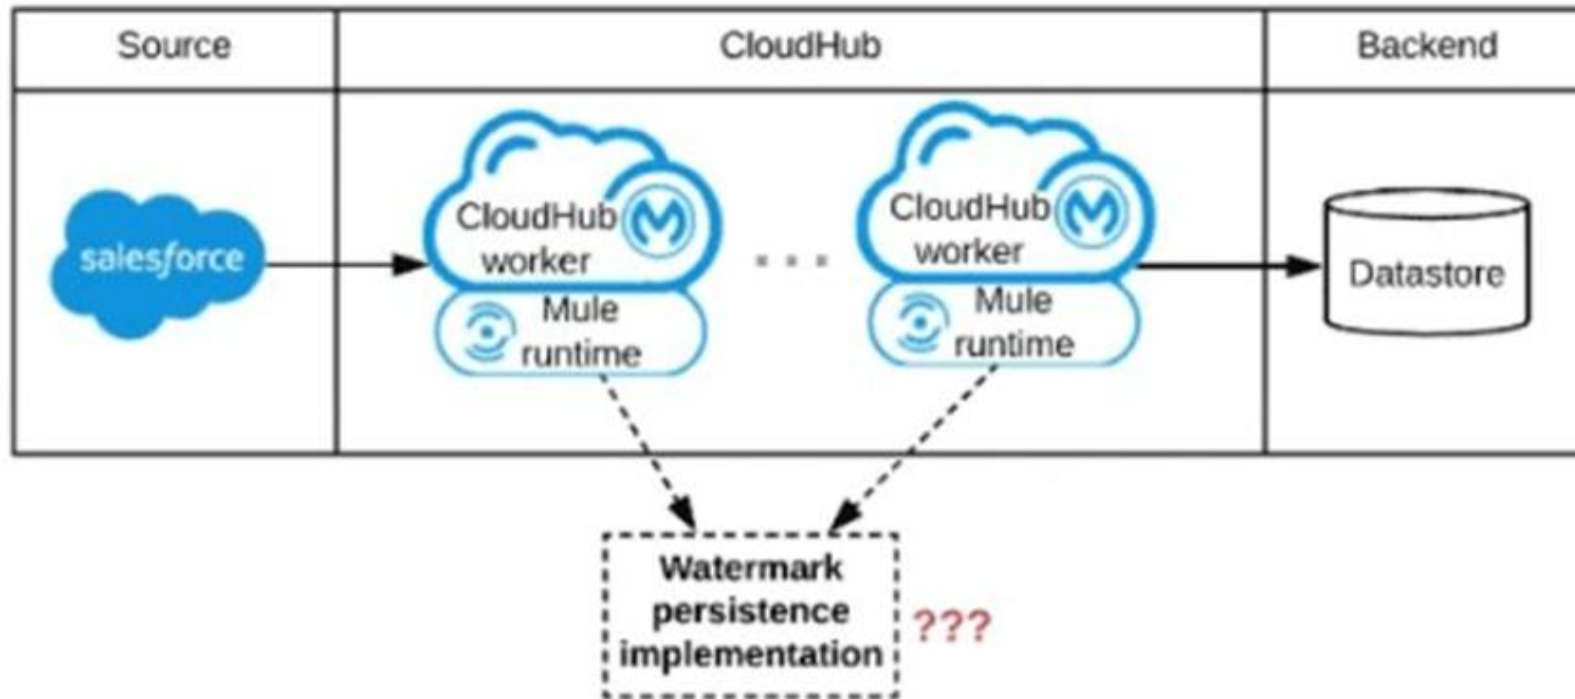
An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments.
 What must the testing team do before they can start integration testing the API in the Staging environment?

- 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 327

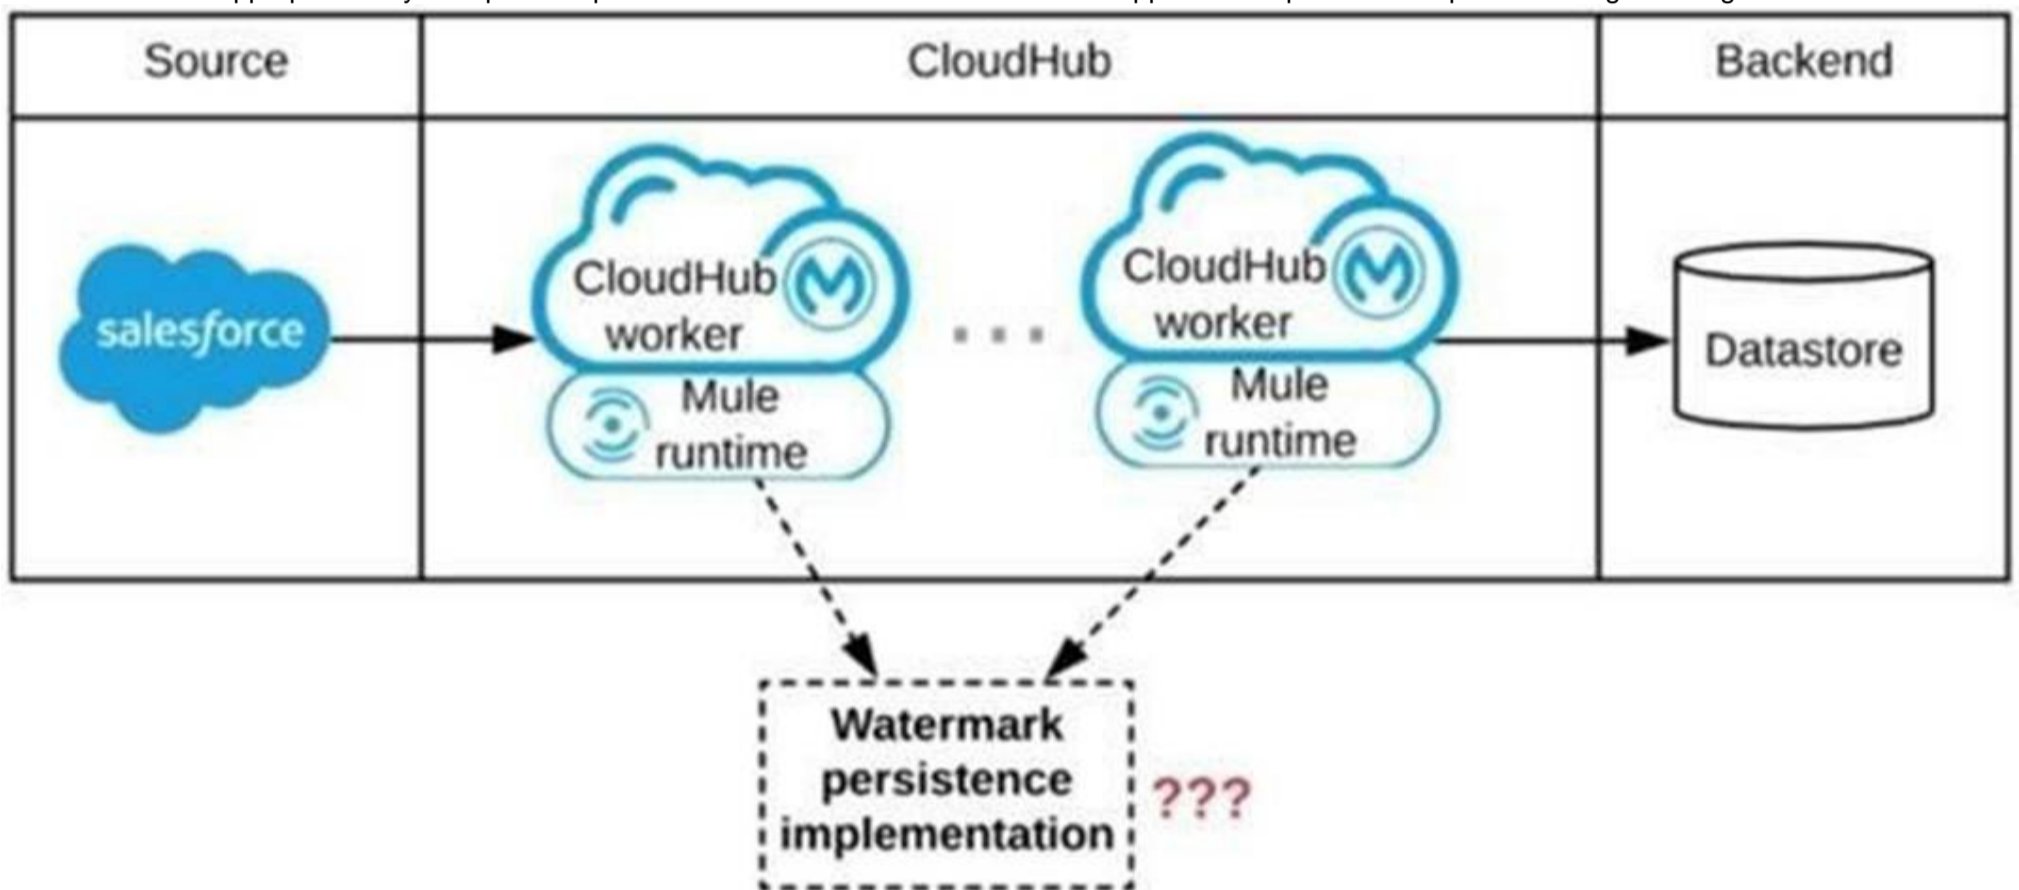
Refer to the exhibit.



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 328

An application load balancer routes requests to a RESTful web API secured by Anypoint Flex Gateway. Which protocol is involved in the communication between the load balancer and the Gateway?

- A. SFTP
- B. HTTPS
- C. LDAP
- D. SMTP

Answer: B

NEW QUESTION 333

An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.

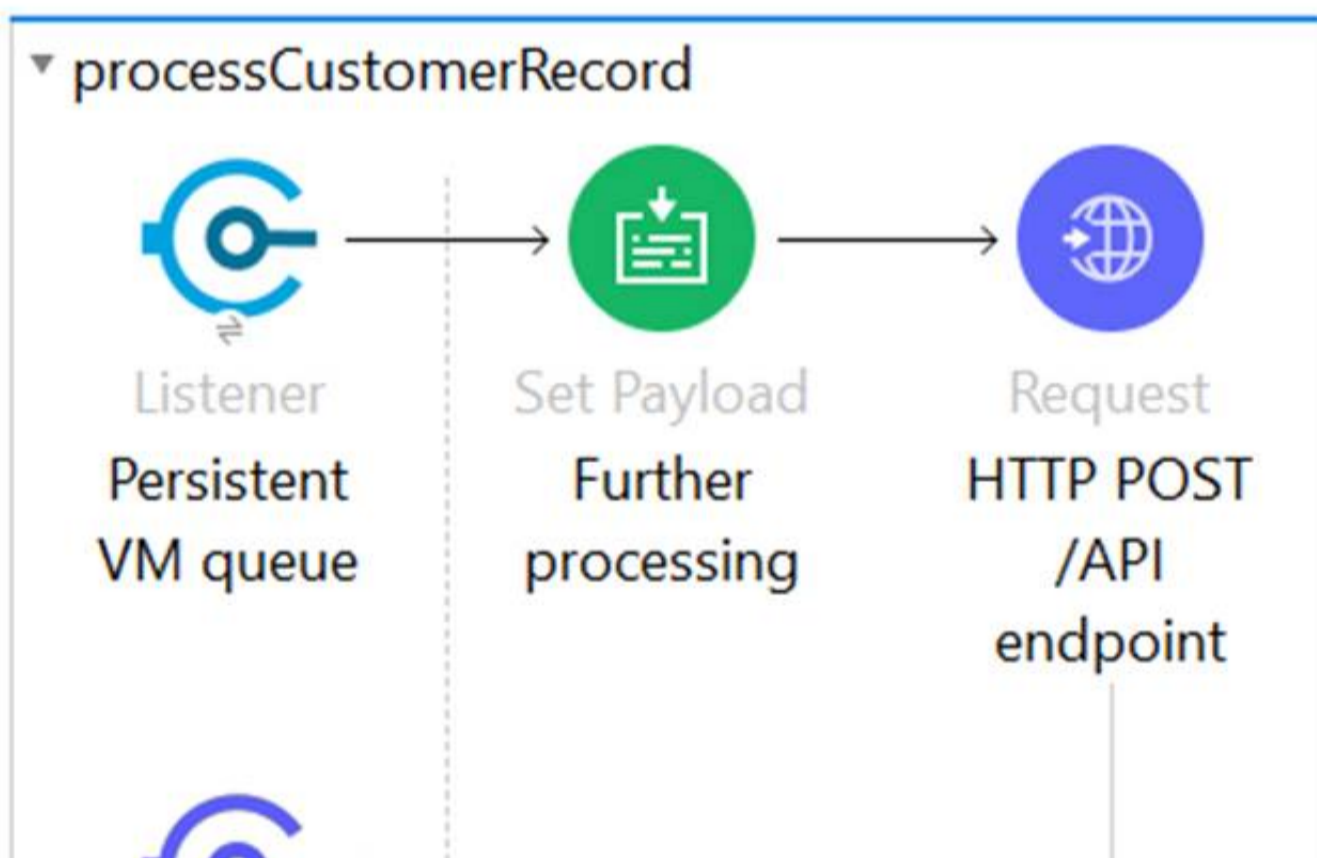
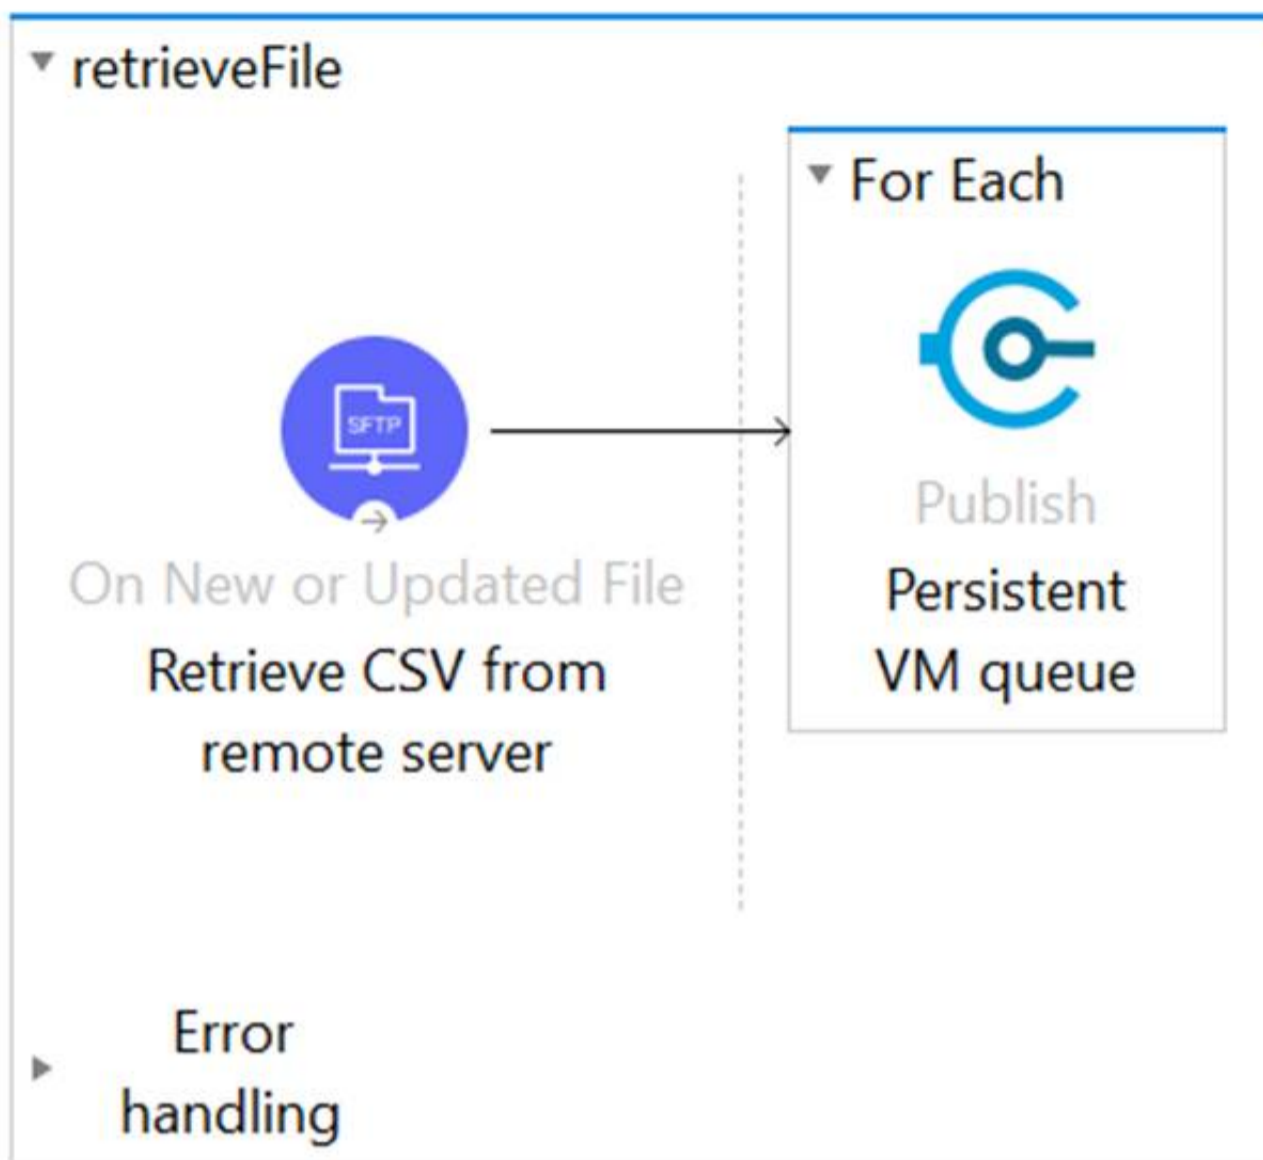
What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

- A. Client management configured in access management
- B. Identity management configured in access management
- C. Enterprise Security module coded in Mule applications
- D. External access configured in API Manager

Answer: A

NEW QUESTION 335

Refer to the exhibit.



This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow's VM Listener receives messages from the same VM queue and then processes each message separately.

How are messages routed to the cloudhub workers as messages are received by the VM Listener?

- A. Each message is routed to ONE of the Cloudhub workers in a DETERMINISTIC round robin fashion thereby EXACTLY BALANCING messages among the cloudhub workers
- B. Each messages routes to ONE of the available Clouhub workers in a NON- DETERMINISTIC non round-robin fashion thereby APPROXIMATELY BALANCING messages among the cloudhub workers
- C. Each message is routed to the SAME Cloudhub worker that retrieved the file, thereby BINDING ALL messages to ONLY that ONE Cloudhub worker
- D. Each message is duplicated to ALL of the Cloudhub workers, thereby SHARING EACH message with ALL the Cloudhub workers.

Answer: B

NEW QUESTION 340

The implementation of a Process API must change. What is a valid approach that minimizes the impact of this change on API clients?

- A. Implement required changes to the Process API implementation so that whenever possible, the Process API's RAML definition remains unchanged
- B. Update the RAML definition of the current Process API and notify API client developers by sending them links to the updated RAML definition
- C. Postpone changes until API consumers acknowledge they are ready to migrate to a new Process API or API version
- D. Implement the Process API changes in a new API implementation, and have the old API implementation return an HTTP status code 301 - Moved Permanently to inform API clients they should be calling the new API implementation

Answer: A

NEW QUESTION 344

A company is planning to extend its Mule APIs to the Europe region. Currently all new applications are deployed to Cloudhub in the US region following this naming convention {API name}-{environment}. for example, Orders-SAPI-dev, Orders-SAPI-prod etc. Considering there is no network restriction to block communications between API's, what strategy should be implemented in order to apply the same new API's running in the EU region of CloudHub as well to minimize latency between API's and target users and systems in Europe?

- A. Set region property to Europe (eu-de) in API manager for all the mule application No need to change the naming convention
- B. Set region property to Europe (eu-de) in API manager for all the mule application Change the naming convention to {API name}-{environment}-{region} and communicate this change to the consuming applications and users
- C. Set region property to Europe (eu-de) in runtime manager for all the mule application No need to change the naming convention
- D. Set region property to Europe (eu-de) in runtime manager for all the mule application Change the naming convention to {API name}-{environment}-{region} and communicate this change to the consuming applications and users

Answer: D

NEW QUESTION 347

According to MuleSoft, what Action should an IT organization take regarding its technology assets in order to close the IT delivery.

- A. Make assets easily discoverable via a central repository
- B. Focus project delivery efforts on custom assets that meet the specific requirements of each individual line of business
- C. Create weekly meetings that all members of IT attend to present justification and request approval to use existing assets
- D. Hire additional staff to meet the demand for asset creation required for approved projects and timelines

Answer: A

NEW QUESTION 352

An organization is using Mulesoft cloudhub and develops API's in the latest version. As a part of requirements for one of the API's, third party API needs to be called. The security team has made it clear that calling any external API needs to have include listing
As an integration architect please suggest the best way to accomplish the design plan to support these requirements?

- A. Implement includelist IP on the cloudhub VPC firewall to allow the traffic
- B. Implement the validation of includelisted IP operation
- C. Implement the Any point filter processor to implement the include list IP
- D. Implement a proxy for the third party API and enforce the IPinclude list policy and call this proxy from the flow of the API

Answer: D

NEW QUESTION 357

Refer to the exhibit.

```
traits:  
  error-responses: traits/error-responses.raml  
  jwt-required:  
    headers:  
      x-jwt:  
        type: string  
        description: JWT token string
```

What is the type data format shown in the exhibit?

- A. JSON
- B. XML
- C. YAML
- D. CSV

Answer: C

NEW QUESTION 360

What is a defining characteristic of an integration-Platform-as-a-Service (iPaaS)?

- A. A Cloud-based
- B. No-code
- C. Code-first
- D. On-premises

Answer: A

NEW QUESTION 363

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 Mute 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 Mute application using DataWeave, invoking existing transformation services when possible

Answer: D

NEW QUESTION 367

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, response size, response time
- D. Request size, number of requests, JDBC Select operation response time

Answer: C

NEW QUESTION 368

An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events

What are the two considerations that must be kept in mind while using test recorder (Choose two answers)

- A. Tests for flows cannot be created with Mule errors raised inside the flow or already existing in the incoming event
- B. Recorder supports smoking a message before or inside a ForEach processor
- C. The recorder support loops where the structure of the data been tested changes inside the iteration
- D. A recorded flow execution ends successfully but the result does not reach its destination because the application is killed
- E. Mocking values resulting from parallel processes are possible and will not affect the execution of the processes that follow in the test

Answer: AD

NEW QUESTION 369

An organization has implemented a continuous integration (CI) lifecycle that promotes Mule applications through code, build, and test stages. To standardize the organization's CI journey, a new dependency control approach is being designed to store artifacts that include information such as dependencies, versioning, and build promotions.

To implement these process improvements, the organization will now require developers to maintain all dependencies related to Mule application code in a shared location.

What is the most idiomatic (used for its intended purpose) type of system the organization should use in a shared location to standardize all dependencies related to Mule application code?

- A. A MuleSoft-managed repository at repository.mulesoft.org
- B. A binary artifact repository
- C. API Community Manager
- D. The Anypoint Object Store service at cloudhub.io

Answer: C

NEW QUESTION 374

What is a key difference between synchronous and asynchronous logging from Mule applications?

- A. Synchronous logging writes log messages in a single logging thread but does not block the Mule event being processed by the next event processor
- B. Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event
- C. Asynchronous logging produces more reliable audit trails with more accurate timestamps
- D. Synchronous logging within an ongoing transaction writes log messages in the same thread that processes the current Mule event

Answer: B

NEW QUESTION 379

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. 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
- B. The API implementation source code must be committed to a source control management system (such as GitHub)
- C. A RAML definition of the API must be created in API designer so it can then be published to Anypoint Exchange
- D. The API must be shared with the potential developers through an API portal so API consumers can interact with the API

Answer: A

NEW QUESTION 383

What is true about automating interactions with Anypoint Platform using tools such as Anypoint Platform REST API's, Anypoint CLI or the Mule Maven plugin?

- A. By default, the Anypoint CLI and Mule Maven plugin are not included in the Mule runtime
- B. Access to Anypoint Platform API's and Anypoint CLI can be controlled separately through the roles and permissions in Anypoint platform, so that specific users can get access to Anypoint CLI while others get access to the platform API's
- C. Anypoint Platform API's can only automate interactions with CloudHub while the Mule maven plugin is required for deployment to customer hosted Mule runtimes
- D. API policies can be applied to the Anypoint platform API's so that only certain LOS's has access to specific functions

Answer: A

NEW QUESTION 387

An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub.

For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub's default zero- downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers. This is expected to result in several Mule application deployments per hour.

- A. 10.0.0.0/21(2048 IPs)
- B. 10.0.0.0/22(1024IPs)
- C. 10.0.0.0/23(512 IPs)
- D. 10.0.0.0/24(256 IPs)

Answer: A

NEW QUESTION 392

A REST API is being designed to implement a Mule application.

What standard interface definition language can be used to define REST APIs?

- A. Web Service Definition Language(WSDL)
- B. OpenAPI Specification (OAS)
- C. YAML
- D. AsyncAPI Specification

Answer: B

NEW QUESTION 393

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: B

NEW QUESTION 397

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual MuleSoft-Integration-Architect-I Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the MuleSoft-Integration-Architect-I Product From:

<https://www.2passeasy.com/dumps/MuleSoft-Integration-Architect-I/>

Money Back Guarantee

MuleSoft-Integration-Architect-I Practice Exam Features:

- * MuleSoft-Integration-Architect-I Questions and Answers Updated Frequently
- * MuleSoft-Integration-Architect-I Practice Questions Verified by Expert Senior Certified Staff
- * MuleSoft-Integration-Architect-I Most Realistic Questions that Guarantee you a Pass on Your First Try
- * MuleSoft-Integration-Architect-I Practice Test Questions in Multiple Choice Formats and Updates for 1 Year