

Salesforce

Exam Questions MuleSoft-Integration-Architect-I

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam



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

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 3

A retail company is implementing a MuleSoft API to get inventory details from two vendors by invoking each vendor's online applications. Due to network issues, the invocations to the vendor applications are timing out intermittently, but the requests are successful after re-invoking each vendor application. What is the most performant way of implementing the API to invoke each vendor application and to retry invocations that generate timeout errors?

- A. Use a For-Each scope to invoke the two vendor applications in series, one after the other. Place the For-Each scope inside an Until-Successful scope to retry requests that raise timeout errors.
- B. Use a Choice scope to invoke each vendor application on a separate route
- C. Place the Choice scope inside an Until-Successful scope to retry requests that raise timeout errors.
- D. Use a Scatter-Gather scope to invoke each vendor application on a separate route
- E. Use an Until-Successful scope in each route to retry requests that raise timeout errors.
- F. Use a Round-Robin scope to invoke each vendor application on a separate route
- G. Use a Try-Catch scope in each route to retry requests that raise timeout errors.

Answer: C

NEW QUESTION 4

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 5

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 6

An API has been updated in Anypoint Exchange by its API producer from version 3.1.1 to 3.2.0 following accepted semantic versioning practices and the changes have been communicated via the API's public portal. The API endpoint does NOT change in the new version. How should the developer of an API client respond to this change?

- A. The update should be identified as a project risk and full regression testing of the functionality that uses this API should be run.
- B. The API producer should be contacted to understand the change to existing functionality.
- C. The API producer should be requested to run the old version in parallel with the new one.
- D. The API client code ONLY needs to be changed if it needs to take advantage of new features.

Answer: D

NEW QUESTION 7

Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improved 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

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

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

Answer: A

NEW QUESTION 9

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

An organization uses Mule runtimes which are managed by Anypoint Platform - Private Cloud Edition. What MuleSoft component is responsible for feeding analytics data to non- MuleSoft analytics platforms?

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

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 14

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL <https://api.acme.com>

As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub

What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

- A. Deploy the Mule applications to CloudhubUpdate the CNAME record for an api.acme.com in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map URLs to their corresponding Mule applications
- B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB)Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications
- C. Deploy the Mule applications to CloudhubCreate CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancerApply mapping rules in the SLB to map URLs to their corresponding Mule applications
- D. Deploy the Mule applications to CloudhubUpdate the CNAME record for api.acme.com in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB)Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

Answer: A

NEW QUESTION 16

A high-volume eCommerce retailer receives thousands of orders per hour and requires notification of its order management, warehouse, and billing system for subsequent processing within 15 minutes of order submission through its website.

Which integration technology, when used for its typical and intended purpose, meets the retailer's requirements for this use case?

- A. Managed File Transfer (MFT)
- B. Publish/Subscriber Messaging Bus (Pub/Sub)
- C. Enterprise Data Warehouse (EDW)
- D. Extract Transform Load (ETL)

Answer: B

NEW QUESTION 18

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

- A. Metadata API
- B. REST API
- C. SOAP API
- D. Bulk API

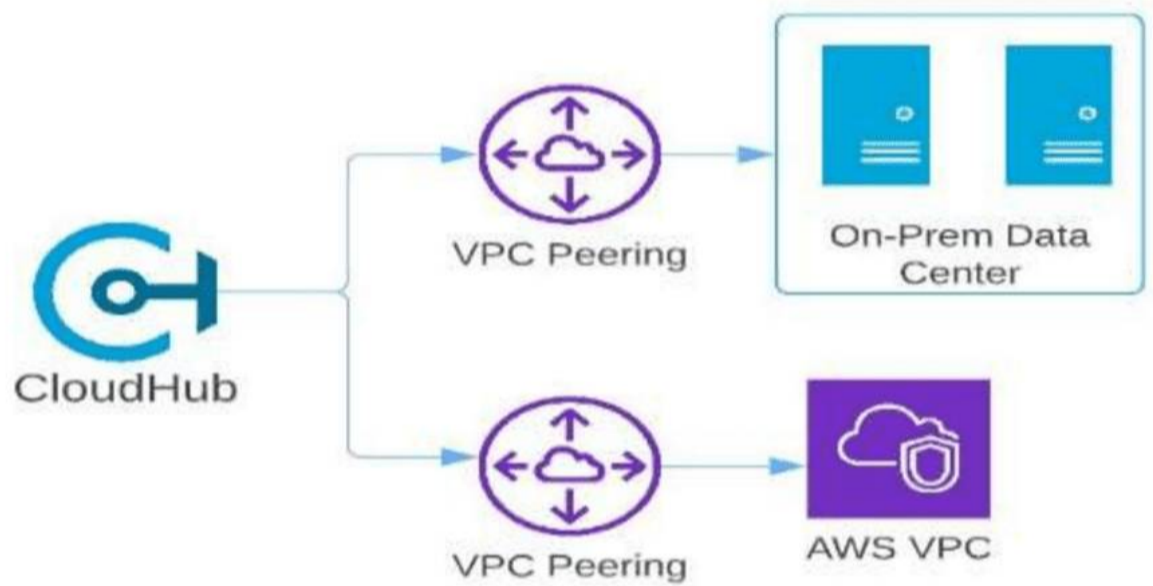
Answer: A

NEW QUESTION 20

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.

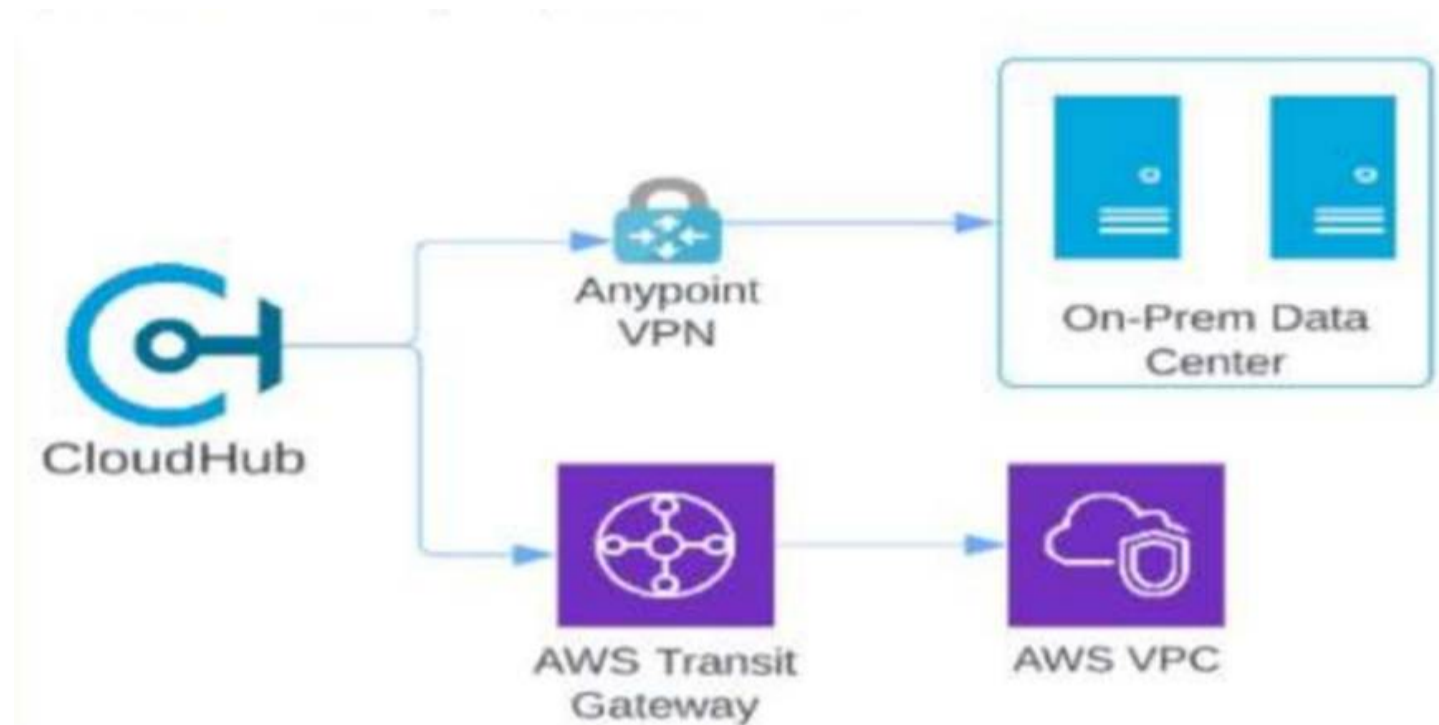


- B) In the AWS account, attach the CloudHub 2.0 private space to an AWS transit gateway that routes from the CloudHub 2.0 private space to the on-premises data center

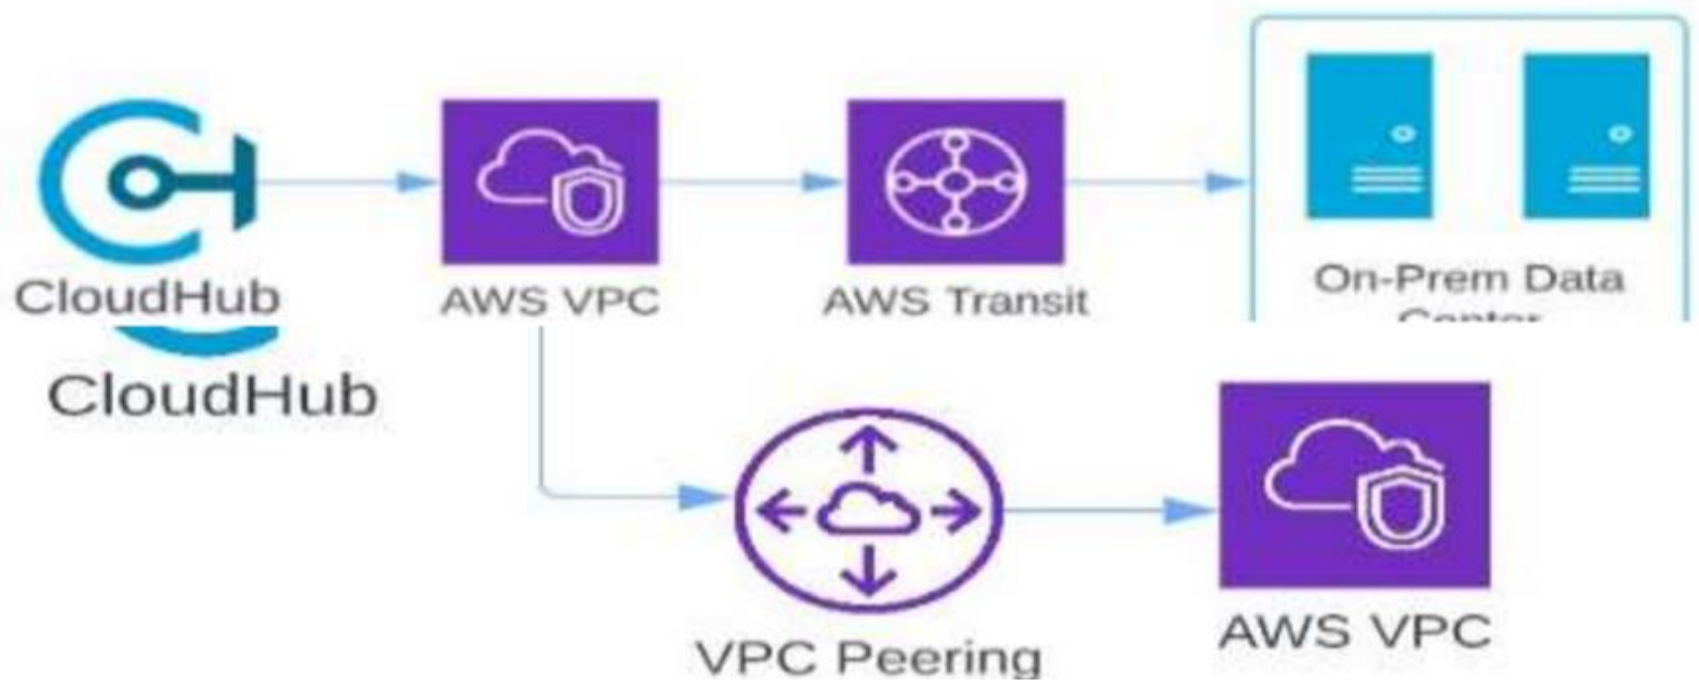
- B. In Runtime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the AWS VPC.



- C) In Runtime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the on-premises data center
- 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 24

A mule application uses an HTTP request operation to involve an external API. The external API follows the HTTP specification for proper status code usage. What is possible cause when a 3xx status code is returned to the HTTP Request operation from the external API?

- A. The request was not accepted by the external API
- B. The request was Redirected to a different URL by the external API
- C. The request was NOT RECEIVED by the external API
- D. The request was ACCEPTED by the external API

Answer: B

NEW QUESTION 26

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

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

Answer: B

NEW QUESTION 29

A developer needs to discover which API specifications have been created within the organization before starting a new project. Which Anypoint Platform component can the developer use to find and try out the currently released API specifications?

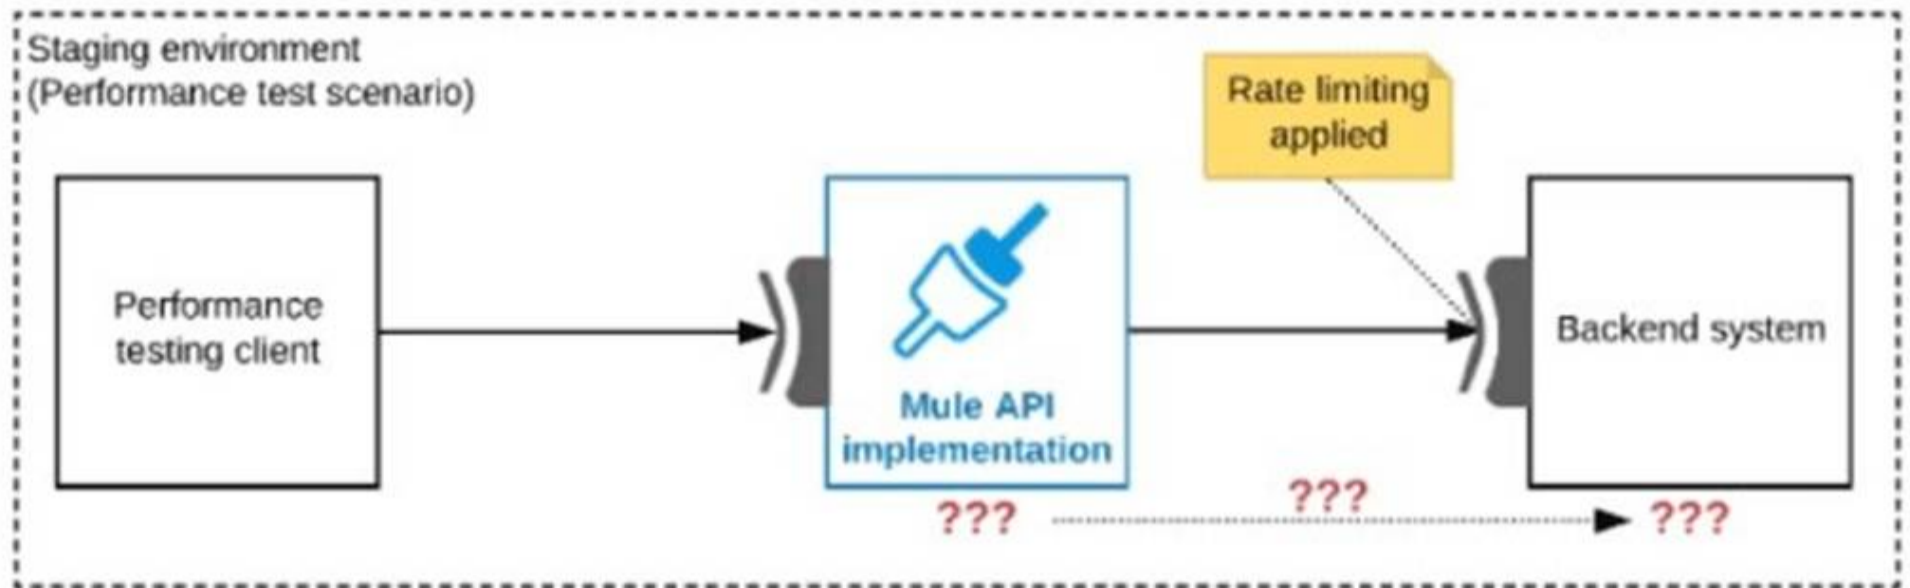
- A. Anypoint Exchange
- B. Runtime Manager
- C. API Manager

D. Object Store

Answer: A

NEW QUESTION 33

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.

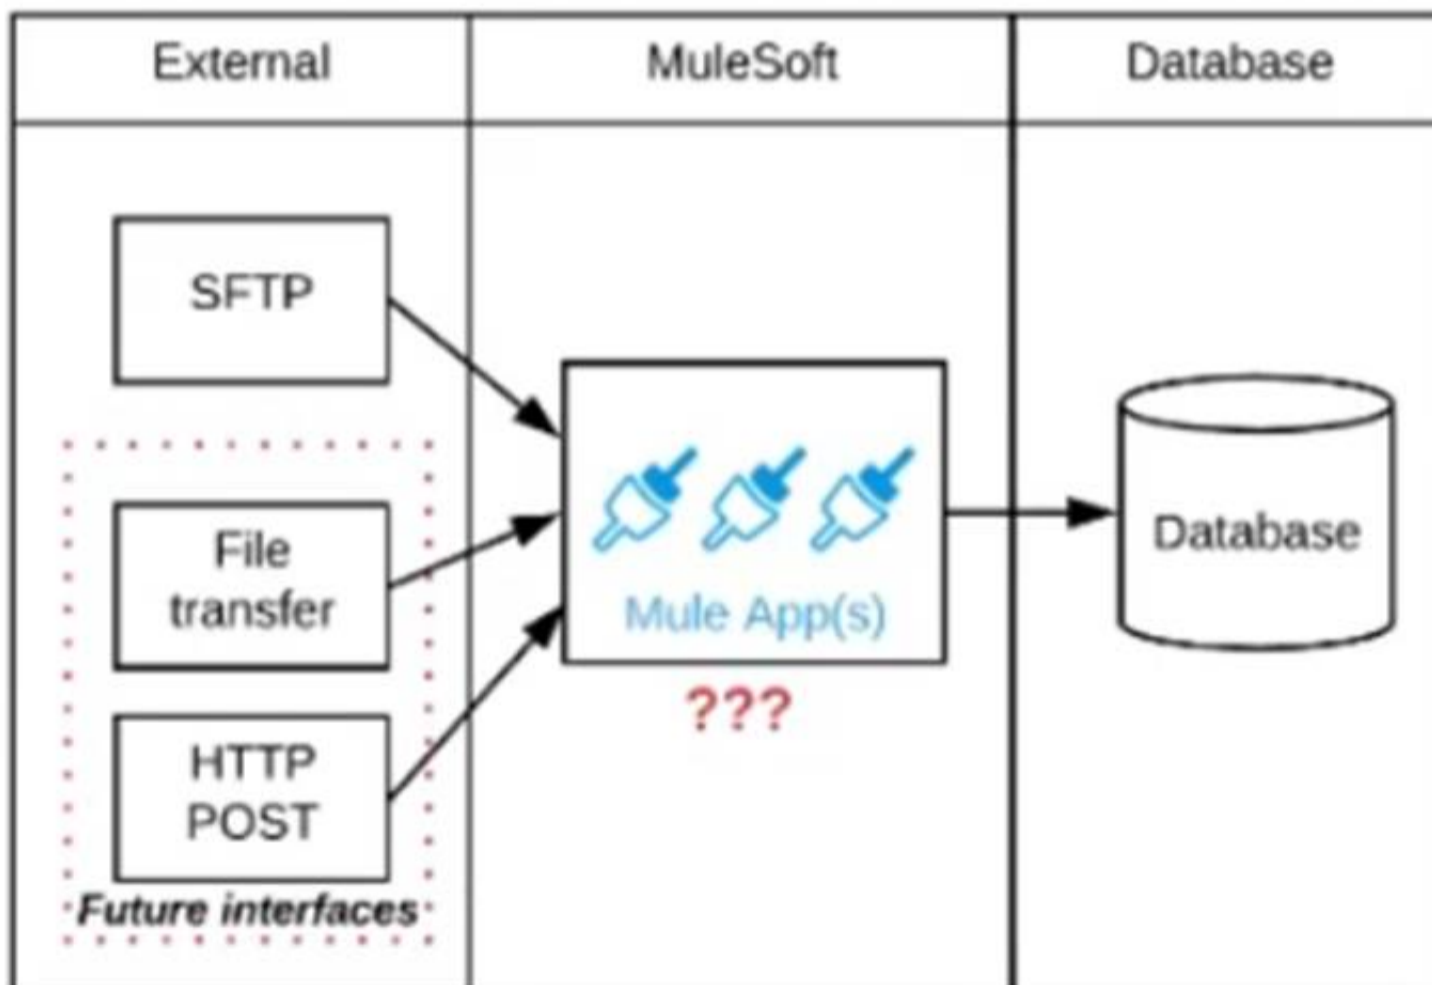
Rate limiting of the backend system applies to all non-production environments. The production environment, however, does NOT have any rate limiting. What is the most effective approach to conduct performance tests of the API implementation in a staging (non-production) environment?

- 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 35

Refer to the exhibit.



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

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

- A. Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- B. Create a Process API to receive the file and process it using a MuleSoft Batch Job while delegating the data save process to a System API
- C. Create an API that receives the file and invokes a Process API with the data contained in the file, then have the Process API process the data using a MuleSoft

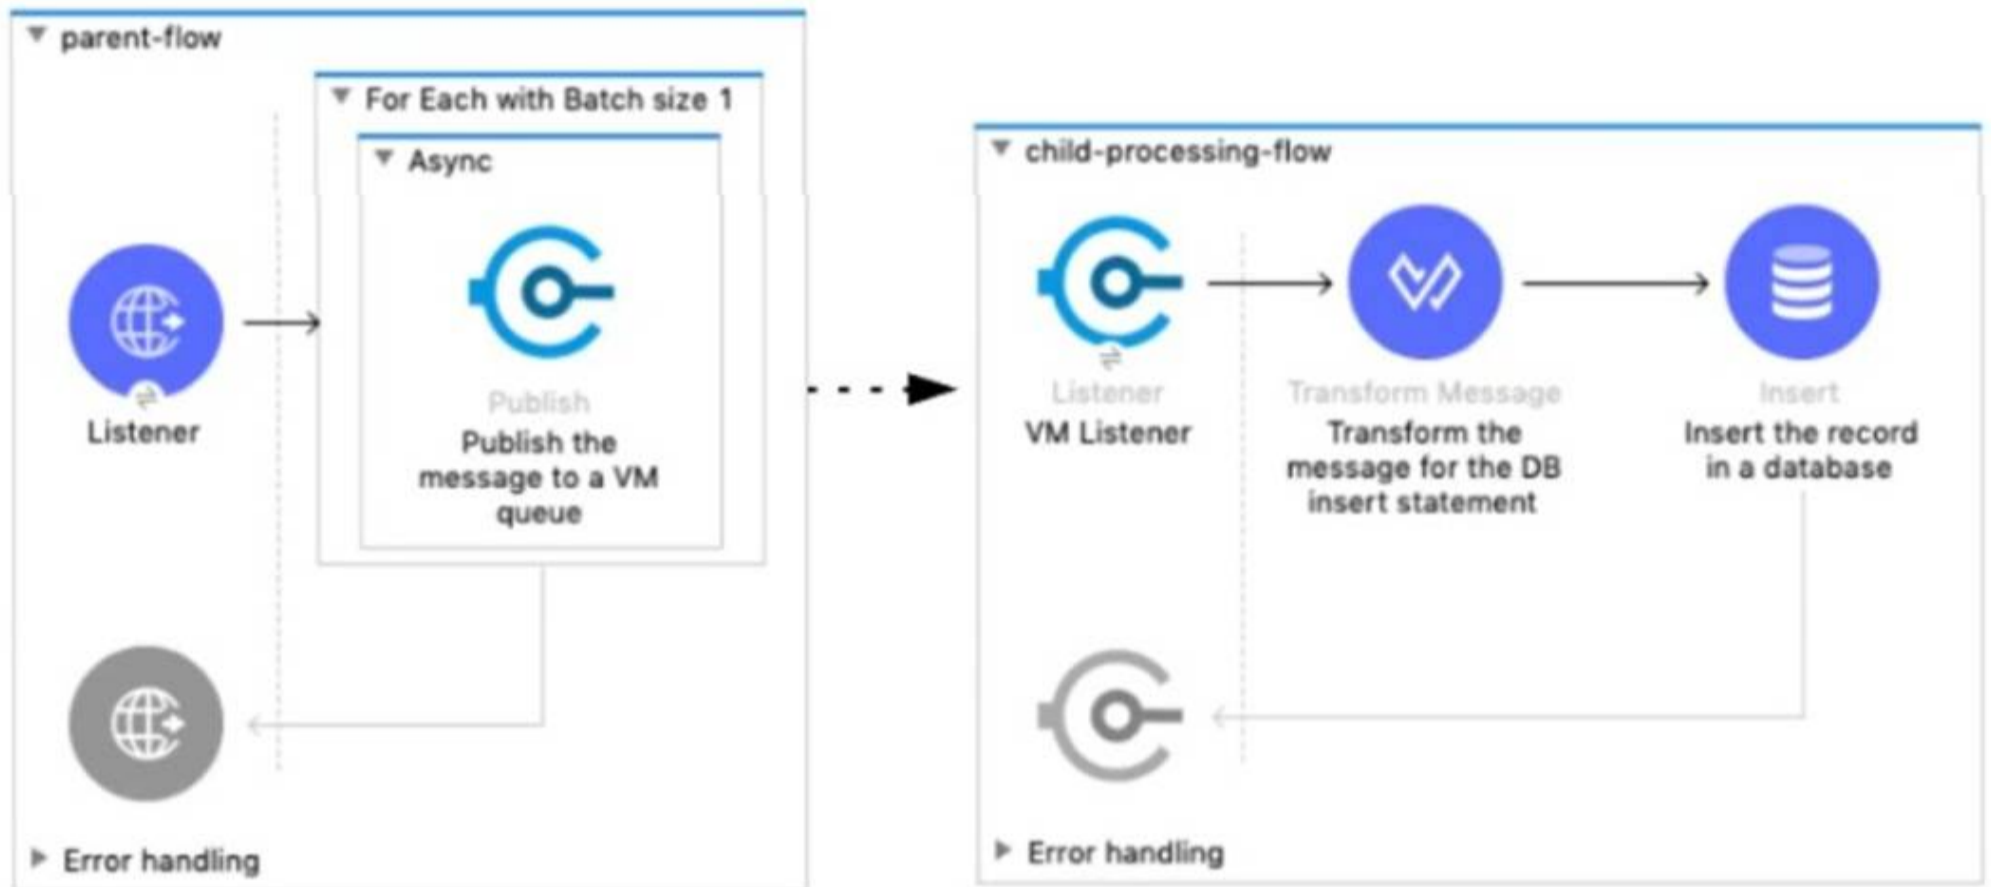
Batch Job and other System APIs as needed

D. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing

Answer: C

NEW QUESTION 36

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 40

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 44

As an enterprise architect, what are the two reasons for which you would use a canonical data model in the new integration project using Mulesoft Anypoint platform (choose two answers)

- A. To have consistent data structure aligned in processes
- B. To isolate areas within a bounded context
- C. To incorporate industry standard data formats
- D. There are multiple canonical definitions of each data type
- E. Because the model isolates the back and systems and support mule applications from change

Answer: AE

NEW QUESTION 47

Mule application muleA deployed in cloudhub uses Object Store v2 to share data across instances. As a part of new requirement , application muleB which is deployed in same region wants to access this Object Store.

Which of the following option you would suggest which will have minimum latency in this scenario?

- A. Object Store REST API
- B. Object Store connector
- C. Both of the above option will have same latency
- D. Object Store of one mule application cannot be accessed by other mule application.

Answer: A

NEW QUESTION 49

An organization is designing multiple new applications to run on CloudHub in a single Anypoint VPC and that must share data using a common persistent Anypoint object store V2 (OSv2).

Which design gives these mule applications access to the same object store instance?

- A. AVM connector configured to directly access the persistence queue of the persistent object store
- B. An Anypoint MQ connector configured to directly access the persistent object store
- C. Object store V2 can be shared across cloudhub applications with the configured osv2 connector
- D. The object store V2 rest API configured to access the persistent object store

Answer: C

NEW QUESTION 52

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 55

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 58

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 61

An organization is designing a mule application to support an all or nothing transaction between several database operations and some other connectors so that they all roll back if there is a problem with any of the connectors

Besides the database connector , what other connector can be used in the transaction.

- A. VM
- B. Anypoint MQ
- C. SFTP
- D. ObjectStore

Answer: A

NEW QUESTION 66

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 71

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 76

When the mule application using VM is deployed to a customer-hosted cluster or multiple cloudhub workers, how are messages consumed by the Mule engine?

- A. in non-deterministic way
- B. by starting an XA transaction for each new message
- C. in a deterministic way
- D. the primary only in order to avoid duplicate processing

Answer: A

NEW QUESTION 78

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 81

An organization has defined a common object model in Java to mediate the communication between different Mule applications in a consistent way. A Mule application is being built to use this common object model to process responses from a SOAP API and a REST API and then write the processed results to an order management system.

The developers want Anypoint Studio to utilize these common objects to assist in creating mappings for various transformation steps in the Mule application. What is the most idiomatic (used for its intended purpose) and performant way to utilize these common objects to map between the inbound and outbound systems in the Mule application?

- A. Use JAXB (XML) and Jackson (JSON) data bindings
- B. Use the WSS module
- C. Use the Java module
- D. Use the Transform Message component

Answer: A

NEW QUESTION 84

A platform architect includes both an API gateway and a service mesh in the architect of a distributed application for communication management. Which type of communication management does a service mesh typically perform in this architecture?

- A. Between application services and the firewall
- B. Between the application and external API clients
- C. Between services within the application
- D. Between the application and external API implementations.

Answer: C

NEW QUESTION 88

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 90

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 column. In the same flow, 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 application. Write the Mule application to publish each Mule event as a message to an Anypoint MQ exchange. Write 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 91

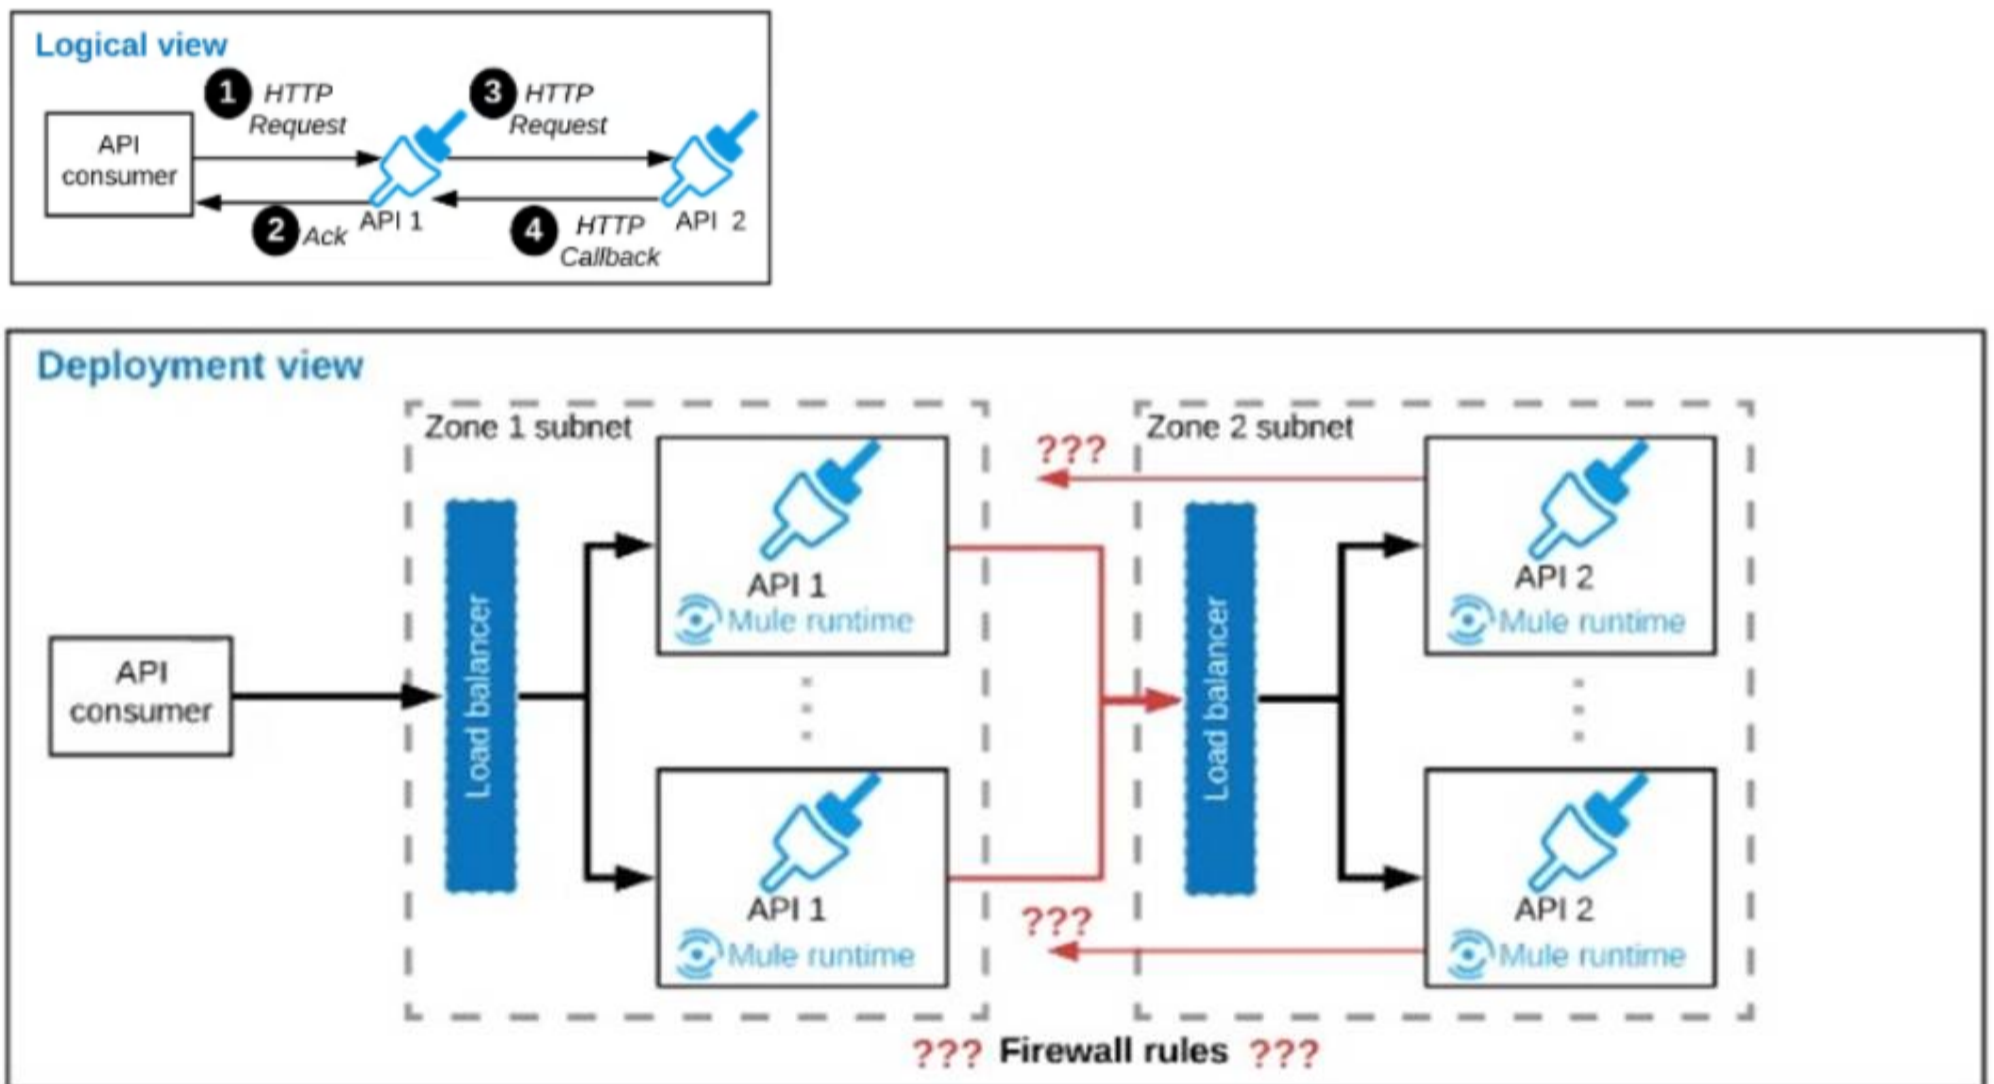
How are the API implementation, API client, and API consumer combined to invoke and process an API?

- A. The API consumer creates an API implementation, which receives API invocations from an API such that they are processed for an API client
- B. The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation
- C. An API client creates an API consumer, which receives API invocation from an API such that they are processed for an API implementation
- D. The API client creates an API consumer which sends API invocations to an API such that they are processed by API implementation

Answer: C

NEW QUESTION 95

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 APIs
- D. To allow communication between load balancers used by each API

Answer: B

NEW QUESTION 96

What is the MuleSoft-recommended best practice to share the connector and configuration information among the APIs?

- A. Build a Mule domain project, add the Database connector and configuration to it, and reference this one domain project from each System API
- B. Build a separate Mule domain project for each API, and configure each of them to use a file on a shared file store to load the configuration information dynamically
- C. Build another System API that connects to the database, and refactor all the other APIs to make requests through the new System API to access the database

D. Create an API proxy for each System API and share the Database connector configuration with all the API proxies via an automated policy

Answer: A

NEW QUESTION 97

A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.

The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.

What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

- A. Retrieve the OpenAPI Specification file(s) from API Manager
- B. Retrieve the WSDL file(s) from the deployed Mule application
- C. Retrieve the RAML file(s) from the deployed Mule application
- D. Retrieve the XML file(s) from Runtime Manager

Answer: D

NEW QUESTION 101

A DevOps team has adequate observability of individual system behavior and performance, but it struggles to track the entire lifecycle of each request across different microservices.

Which additional observability approach should this team consider adopting?

- A. Analytics
- B. Metrics
- C. Tracing
- D. Data mining

Answer: C

NEW QUESTION 104

A bank is implementing a REST API in a Mule application to receive an array of accounts from an online banking platform user interface (UI), retrieve account balances for those accounts from a backend Finance system, and then return the account balances so they can be displayed in the online banking platform UI. As part of the processing, the MuleSoft API also needs to insert the retrieved account data into an AuditDatabase for auditing purposes. The auditing process should not add latency to the account balance retrieval response back to the online banking platform UI.

The retrieveBalances flow in the Mule application is designed to use an operation in a connector to the Finance system (the Finance operation) that can only look up one account record at a time, and a operation from a different connector to the Audit system (the Audit operation) that can only insert one account record at a time.

To best meet the performance-related requirements, what scope or scopes should be used and how should they be used to incorporate the Finance operation and Audit operation into the retrieveBalances flow?

- A. Wrap the Finance operation in a Parallel For-Each scope
- B. Wrap the Audit operation in a Async scope.
- C. Wrap the Finance operation in a Until-Successful scope
- D. Wrap the Audit operation in a Try-Catch scope.
- E. Wrap both connector operations in a Async scope.
- F. Wrap both connector operations in a For-Each scope.

Answer: A

NEW QUESTION 109

An organization needs to procure an enterprise software system to increase cross-selling opportunities and better track prospect data.

Which category of enterprise software has these core capabilities, when used for its typical and intended purpose?

- A. Supply Chain Management (SCM)
- B. IT Service Management (ITSM)
- C. Business-to-Business (A2B)
- D. Customer Relationship Management (CRM)

Answer: D

NEW QUESTION 110

What best describes the Fully Qualified Domain Names (FQDNs), also known as DNS entries, created when a Mule application is deployed to the CloudHub Shared Worker Cloud?

- A. A fixed number of FQDNs are created, IRRESPECTIVE of the environment and VPC design
- B. The FQDNs are determined by the application name chosen, IRRESPECTIVE of the region
- C. The FQDNs are determined by the application name, but can be modified by an administrator after deployment
- D. The FQDNs are determined by both the application name and the region

Answer: D

NEW QUESTION 112

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 115

An organization has implemented the cluster with two customer hosted Mule runtimes is hosting an application.

This application has a flow with a JMS listener configured to consume messages from a queue destination. As an integration architect can you advise which JMS listener configuration must be used to receive messages in all the nodes of the cluster?

- A. Use the parameter `primaryNodeOnly= "false"` on the JMS listener
- B. Use the parameter `primaryNodeOnly= "false"` on the JMS listener with a shared subscription
- C. Use the parameter `primaryNodeOnly= "true"` on the JMS listener with a non-shared subscription
- D. Use the parameter `primaryNodeOnly= "true"` on the JMS listener

Answer: B

NEW QUESTION 116

An organization has deployed both Mule and non-Mule API implementations to integrate its customer and order management systems. All the APIs are available to REST clients on the public internet.

The organization wants to monitor these APIs by running health checks: for example, to determine if an API can properly accept and process requests. The organization does not have subscriptions to any external monitoring tools and also does not want to extend its IT footprint.

What Anypoint Platform feature provides the most idiomatic (used for its intended purpose) way to monitor the availability of both the Mule and the non-Mule API implementations?

- A. API Functional Monitoring
- B. Runtime Manager
- C. API Manager
- D. Anypoint Visualizer

Answer: D

NEW QUESTION 120

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 122

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 RuntimeManager Define 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 Manager Define the splunk appender in EACH Mule application's log4j2.xml file
- C. Disable Cloudhub logging in Runtime Manager Define the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager 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 124

A Mule application is deployed to a cluster of two(2) customer-hosted Mule runtimes. Currently the node name Alice is the primary node and node named bob is the secondary node. The mule application has a flow that polls a directory on a file system for new files.

The primary node Alice fails for an hour and then restarted.

After the Alice node completely restarts, from what node are the files polled, and what node is now the primary node for the cluster?

- A. Files are polled from Alice node Alice is now the primary node
- B. Files are polled from Bob node Alice is now the primary node
- C. Files are polled from Alice node Bob is the now the primary node
- D. Files are polled from Bob node Bob is now the primary node

Answer: D

NEW QUESTION 127

An IT integration delivery team begins a project by gathering all of the requirements, and proceeds to execute the remaining project activities as sequential, non-repeating phases.

Which IT project delivery methodology is this team following?

- A. Kanban
- B. Scrum
- C. Waterfall
- D. Agile

Answer: C

NEW QUESTION 132

A system administrator needs to determine when permissions were last changed for an Anypoint Platform user.

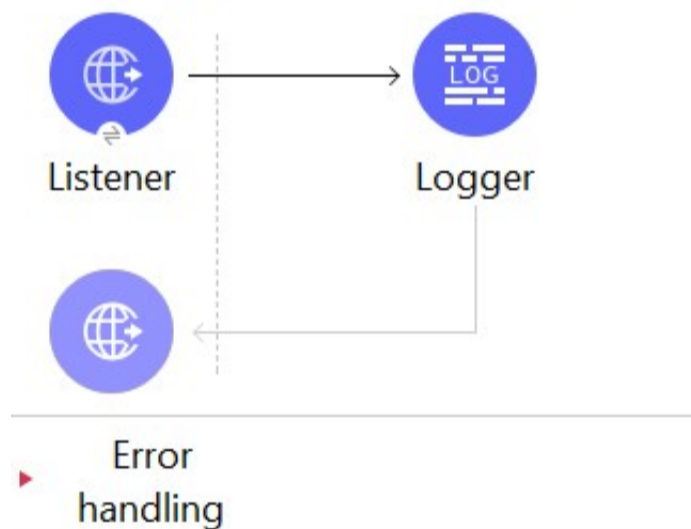
Which Anypoint Platform component should the administrator use to obtain this information?

- A. Audit Logging
- B. Anypoint Monitoring
- C. Anypoint Studio
- D. Mule Stack Traces

Answer: A

NEW QUESTION 137

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 141

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 eventFor each Batch Step scope, all 25 records within a block are processed in parallelAll 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 eventFor each Batch Step scope, all 25 records within a block are processed sequentially, one at a timeAll 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 eventFor each Batch Step scope, all 25 records within a block are processed sequentially, one record at a timeAll 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 parallelEach Batch Step scope is invoked with a batch of 25 records in the payload of the received Mule eventFor each Batch Step scope, all 4000 records are processed in parallelIndividual 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 146

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 149

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 152

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 156

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

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

Answer: C

NEW QUESTION 157

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 160

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 165

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 167

What is not true about Mule Domain Project?

- A. This allows Mule applications to share resources
- B. Expose multiple services within the Mule domain on the same port
- C. Only available Anypoint Runtime Fabric
- D. Send events (messages) to other Mule applications using VM queues

Answer: C

NEW QUESTION 170

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 mule application is required to periodically process large data set from a back-end database to Salesforce CRM using batch job scope configured properly process the higher rate of records.

The application is deployed to two cloudhub workers with no persistence queues enabled. What is the consequence if the worker crashes during records processing?

- A. Remaining records will be processed by a new replacement worker
- B. Remaining records be processed by second worker
- C. Remaining records will be left and processed
- D. All the records will be processed from scratch by the second worker leading to duplicate processing

Answer: D

NEW QUESTION 173

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 175

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 176

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 181

A Mule application uses the Database connector.

What condition can the Mule application automatically adjust to or recover from without needing to restart or redeploy the Mule application?

- A. One of the stored procedures being called by the Mule application has been renamed
- B. The database server was unavailable for four hours due to a major outage but is now fully operational again
- C. The credentials for accessing the database have been updated and the previous credentials are no longer valid
- D. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade

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 190

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 193

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(??healthy")
- D. status: healthy

Answer: B

NEW QUESTION 196

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 scope that participates in the local transaction.

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

- A. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
- B. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back

- C. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
- D. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back

Answer: A

NEW QUESTION 201

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 202

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

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

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

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

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

Answer: C

NEW QUESTION 206

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 transaction3) Acknowledge the JMS message
- B. 1) Read the JMS message (NOT in an XA transaction)2) Perform EACH DB insert in a SEPARATE DB transaction3) Acknowledge the JMS message
- C. 1) Read the JMS message in an XA transaction2) 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 210

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 consumed these files?

- A. ATLS context Key Store requiring the private key and certificate for the company PGP public key of partnerPGP 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 partnerTLS 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 213

According to MuleSoft, a synchronous invocation of a RESTful API using HTTP to get an individual customer record from a single system is an example of which system integration interaction pattern?

- A. Request-Reply
- B. Multicast
- C. Batch
- D. One-way

Answer: A

NEW QUESTION 215

An organization is not meeting its growth and innovation objectives because IT cannot deliver projects fast enough to keep up with the pace of change required by the business.

According to MuleSoft's IT delivery and operating model, which step should the organization take to solve this problem?

- A. Modify IT governance and security controls so that line of business developers can have direct access to the organization's systems of record
- B. Switch from a design-first to a code-first approach for IT development
- C. Adopt a new approach that decouples core IT projects from the innovation that happens within each line of business
- D. Hire more IT developers, architects, and project managers to increase IT delivery

Answer: C

NEW QUESTION 219

A customer wants to use the mapped diagnostic context (MDC) and logging variables to enrich its logging and improve tracking by providing more context in the logs.

The customer also wants to improve the throughput and lower the latency of message processing.

As an Mulesoft integration architect can you advise, what should the customer implement to meet these requirements?

- A. Use synchronous logging and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
- B. Use async logger at the level greater than INFO and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
- C. Use async logger at the level equal to DEBUG or TRACE and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
- D. Use synchronous logging at the INFO DEBUG or Trace level and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables

Answer: B

NEW QUESTION 220

An organization is struggling frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects.

How can these best practices be achieved with the applications having the least amount of responsibility?

- A. Create a Mule plugin project with all the dependencies and add it as a dependency in each application's POM.xml file
- B. Create a mule domain project with all the dependencies define in its POM.xml file and add each application to the domain Project
- C. Add all dependencies in each application's POM.xml file
- D. Create a parent POM of all the required dependencies and reference each in each application's POM.xml file

Answer: D

NEW QUESTION 223

An integration team uses Anypoint Platform and follows MuleSoft's recommended approach to full lifecycle API development.

Which step should the team's API designer take before the API developers implement the API Specification?

- A. Generate test cases using MUnit so the API developers can observe the results of running the API
- B. Use the scaffolding capability of Anypoint Studio to create an API portal based on the API specification
- C. Publish the API specification to Exchange and solicit feedback from the API's consumers
- D. Use API Manager to version the API specification

Answer: C

NEW QUESTION 228

An organization is in the process of building automated deployments using a CI/CD process. As a part of automated deployments, it wants to apply policies to API Instances.

What tool can the organization use to promote and deploy API Manager policies?

- A. Anypoint CLI
- B. MUnit Maven plugin
- C. Mule Maven plugin
- D. Runtime Manager agent

Answer: A

NEW QUESTION 232

An organization's governance process requires project teams to get formal approval from all key stakeholders for all new Integration design specifications. An integration Mule application is being designed that interacts with various backend systems. The Mule application will be created using Anypoint Design Center or Anypoint Studio and will then be deployed to a customer-hosted runtime.

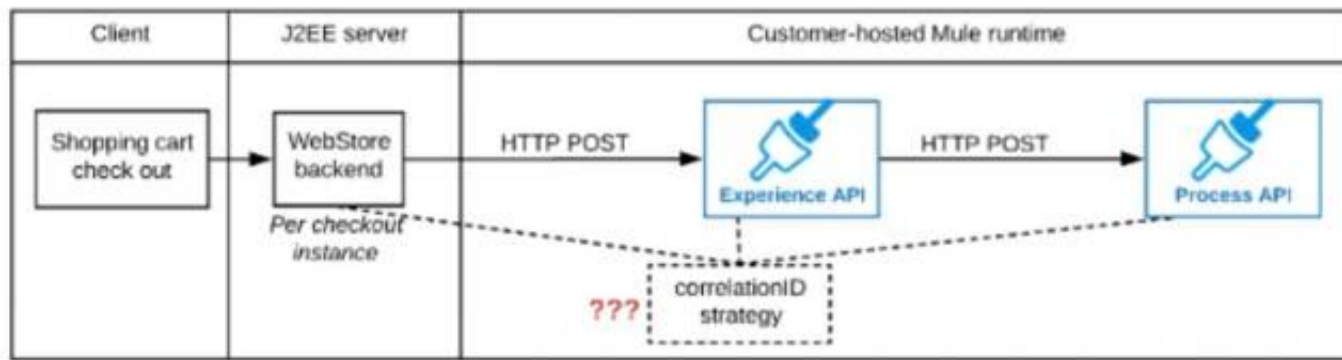
What key elements should be included in the integration design specification when requesting approval for this Mule application?

- A. SLAs and non-functional requirements to access the backend systems
- B. Snapshots of the Mule application's flows, including their error handling
- C. A list of current and future consumers of the Mule application and their contact details
- D. The credentials to access the backend systems and contact details for the administrator of each system

Answer: A

NEW QUESTION 237

Refer to the exhibit.

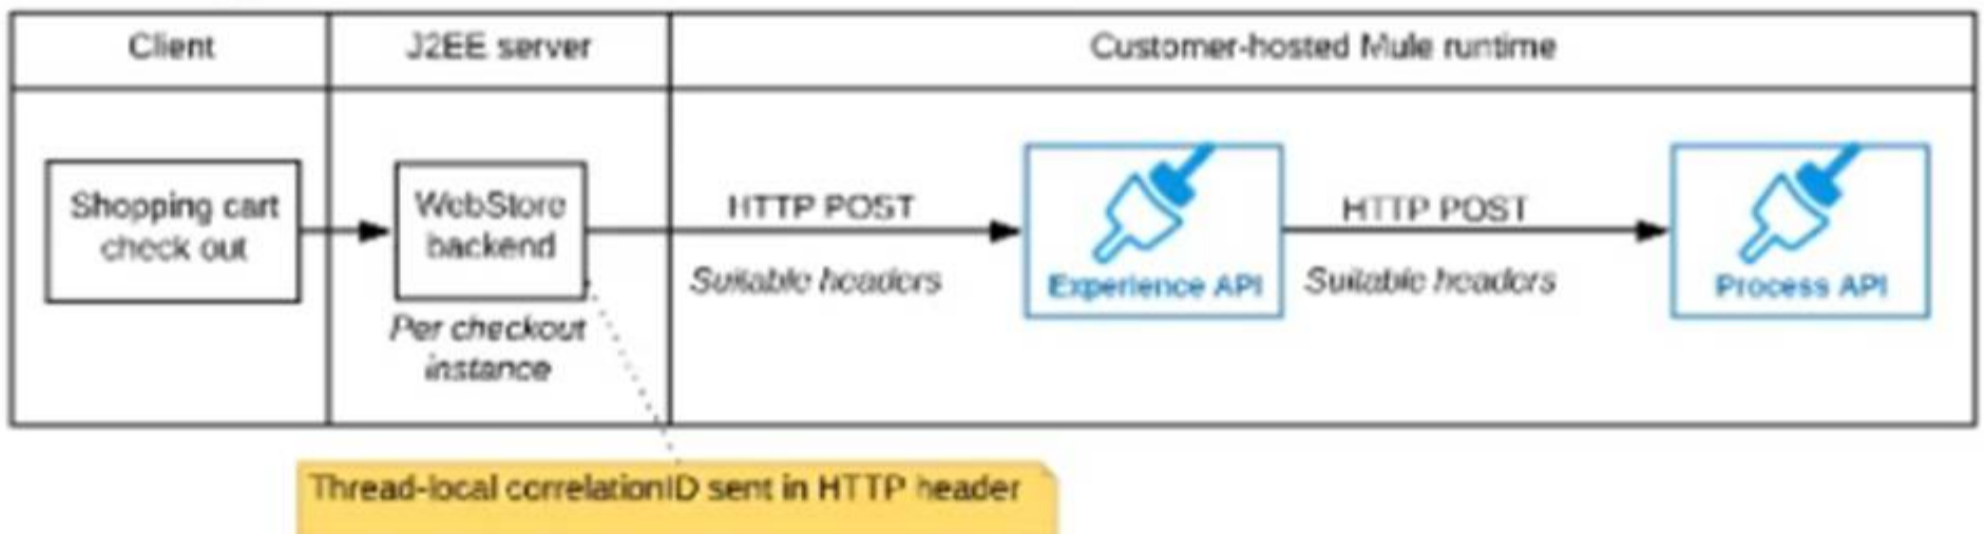


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

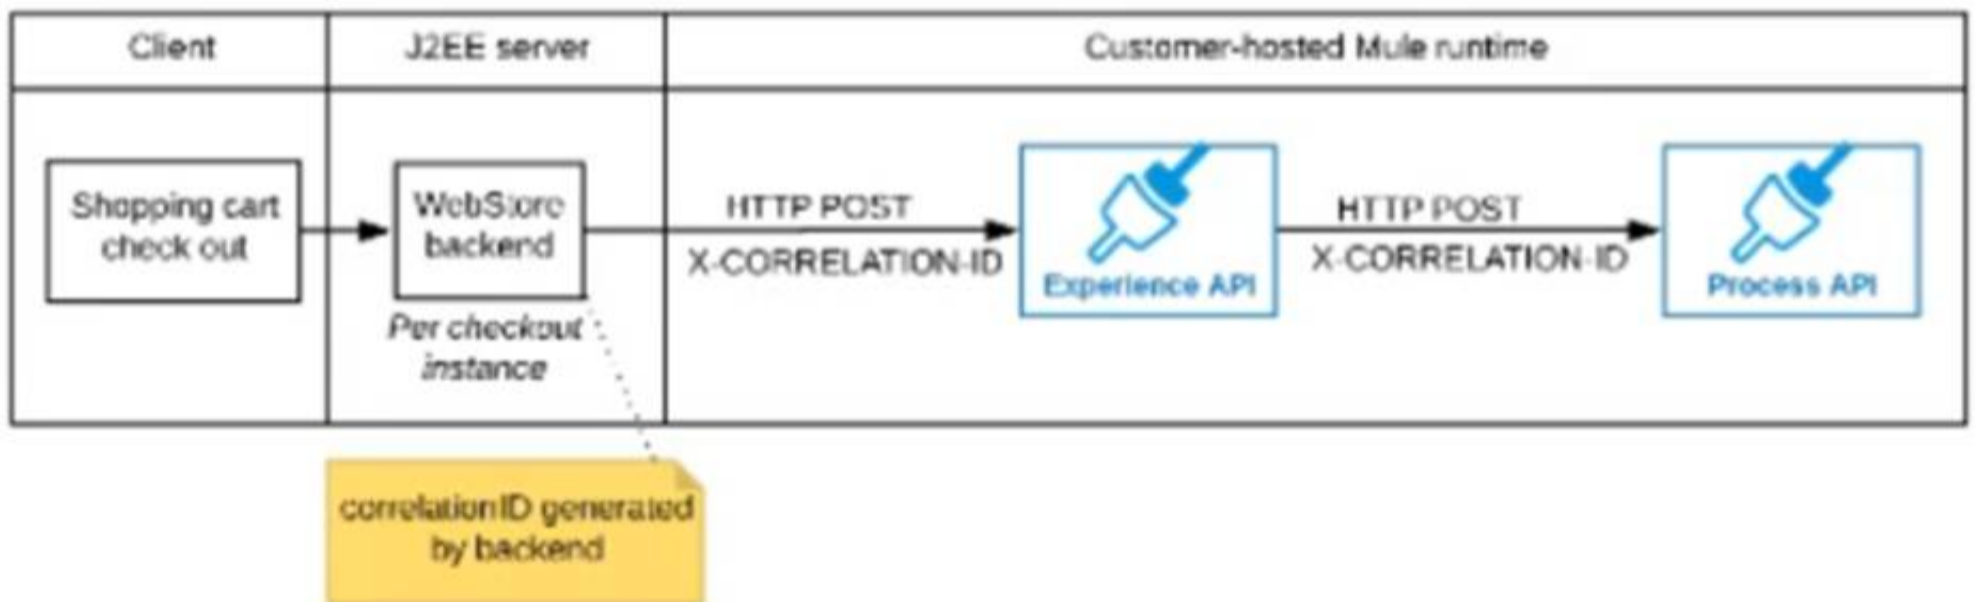
What is the most efficient way (using the least amount of custom coding or configuration) for the web store backend and the implementations of the Experience API and Process API to participate in end-to-end correlation of the API invocations for each checkout instance?

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



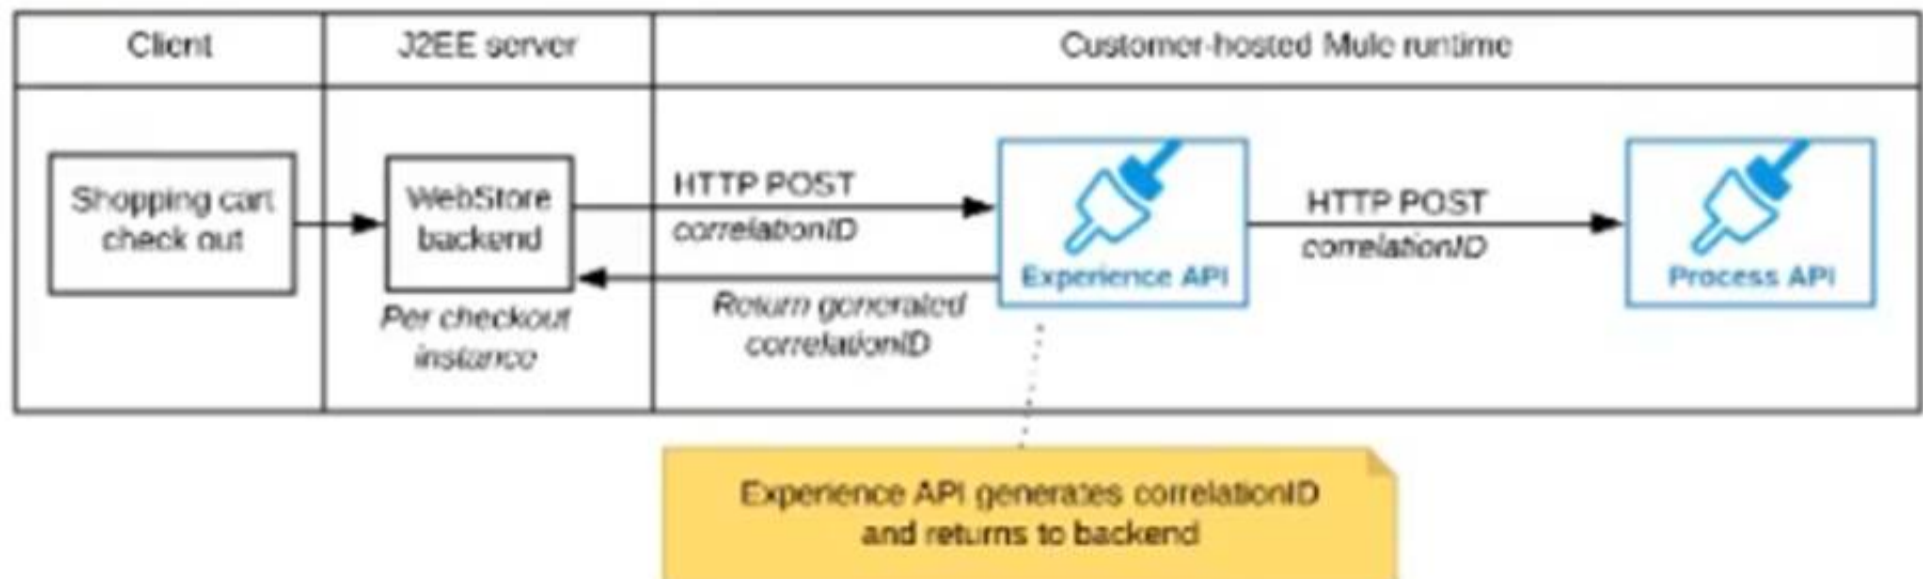
B) The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header In each API invocation belonging to that checkout

No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID



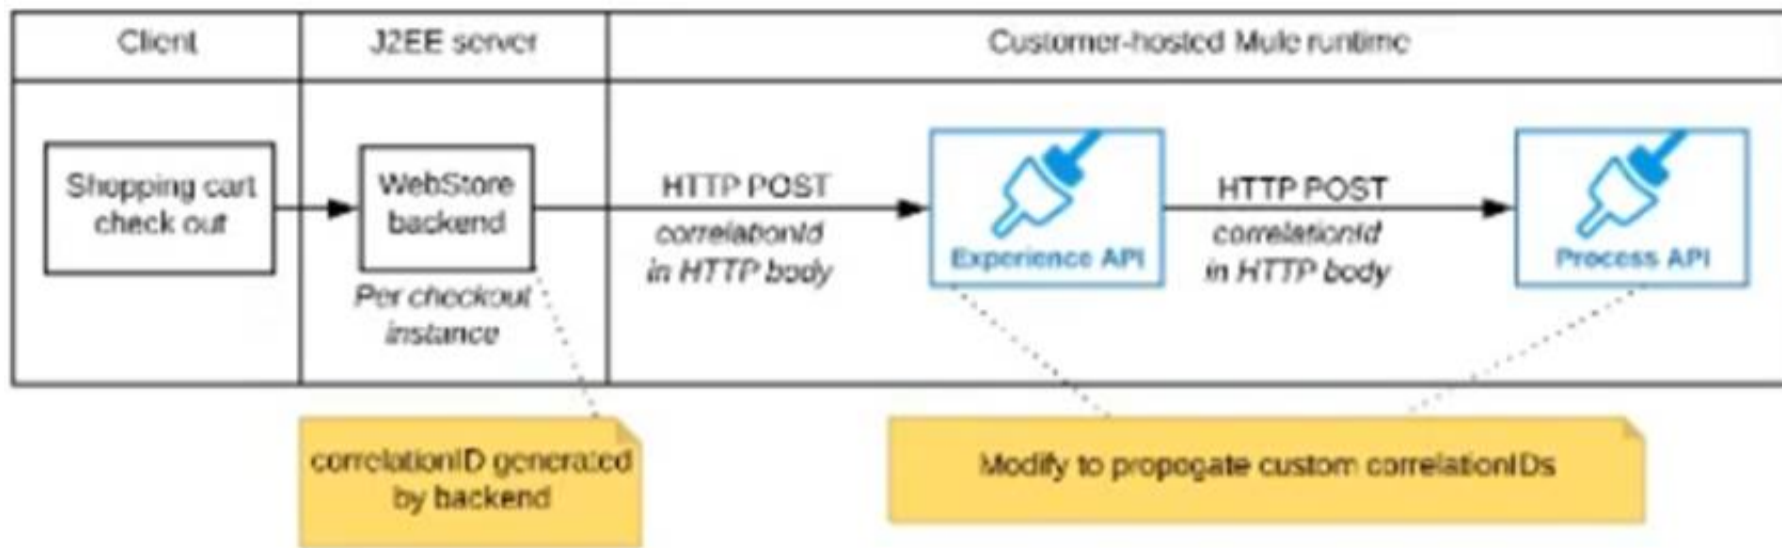
C) The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

The Experience API implementation must be coded to also propagate the correlation 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 241

A corporation has deployed multiple mule applications implementing various public and private API's to different cloudhub workers. These API's arc Critical applications that must be highly available and in line with the reliability SLA as defined by stakeholders.

How can API availability (liveliness or readiness) be monitored so that Ops team receives outage notifications?

- A. Enable monitoring of individual applications from Anypoint monitoring
- B. Configure alerts with failure conditions in runtime manager
- C. Configure alerts failure conditions in API manager
- D. Use any point functional monitoring test API's functional behavior

Answer: A

NEW QUESTION 243

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 247

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 250

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 255

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 260

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 263

An auto mobile company want to share inventory updates with dealers D1 and D2 asynchronously and concurrently via queues Q1 and Q2. Dealer D1 must consume the message from the queue Q1 and dealer D2 to must consume a message from the queue Q2.

Dealer D1 has implemented a retry mechanism to reprocess the transaction in case of any errors while processing the inventers updates. Dealer D2 has not implemented any retry mechanism.

How should the dealers acknowledge the message to avoid message loss and minimize impact on the current implementation?

- A. Dealer D1 must use auto acknowledgement and dealer D2 can use manual acknowledgement and acknowledge the message after successful processing
- B. Dealer D1 can use auto acknowledgement and dealer D2 can use IMMEDIATE acknowledgement and acknowledge the message of successful processing
- C. Dealer D1 and dealer D2 must use AUTO acknowledgement and acknowledge the message after successful processing
- D. Dealer D1 can use AUTO acknowledgement and dealer D2 must use manual acknowledgement and acknowledge the message after successful processing

Answer: D

NEW QUESTION 267

Which Exchange asset type represents configuration modules that extend the functionality of an API and enforce capabilities such as security?

- A. Rulesets
- B. Policies
- C. REST APIs
- D. Connectors

Answer: B

NEW QUESTION 268

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 269

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 Use persistent storage when creating RESP

Answer: D

NEW QUESTION 271

An ABC Farms project team is planning to build a new API that is required to work with data from different domains across the organization.

The organization has a policy that all project teams should leverage existing investments by reusing existing APIs and related resources and documentation that other project teams have already developed and deployed.

To support reuse, where on Anypoint Platform should the project team go to discover and read existing APIs, discover related resources and documentation, and interact with mocked versions of those APIs?

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

Answer: D

NEW QUESTION 274

As a part of project , existing java implementation is being migrated to Mulesoft. Business is very tight on the budget and wish to complete the project in most economical way possible.

Canonical object model using java is already a part of existing implementation. Same object model is required by mule application for a business use case. What is the best way to achieve this?

- A. Make use of Java module
- B. Create similar model for Mule applications
- C. Create a custom application to read Java code and make it available for Mule application
- D. Use Anypoint exchange

Answer: A

NEW QUESTION 277

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 279

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 stable 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 said new employees data to cacheUse an object store caching strategy and set expiration interval to 1-hour

Answer: A

NEW QUESTION 281

As a part of project requirement, client will send a stream of data to mule application. Payload size can vary between 10mb to 5GB. Mule application is required to transform the data and send across multiple sftp servers. Due to the cost cuttings in the organization, mule application can only be allocated one worker with size of 0.2 vCore.

As an integration architect , which streaming strategy you would suggest to handle this scenario?

- A. In-memory non repeatable stream
- B. File based non-repeatable stream
- C. In-memory repeatable stream
- D. File based repeatable storage

Answer: D

NEW QUESTION 282

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 285

To implement predictive maintenance on its machinery equipment, ACME Tractors has installed thousands of IoT sensors that will send data for each machinery asset as sequences of JMS messages, in near real-time, to a JMS queue named SENSOR_DATA on a JMS server. The Mule application contains a JMS Listener operation configured to receive incoming messages from the JMS servers SENSOR_DATA JMS queue. The Mule application persists each received JMS message, then sends a transformed version of the corresponding Mule event to the machinery equipment back-end systems.

The Mule application will be deployed to a multi-node, customer-hosted Mule runtime cluster. Under normal conditions, each JMS message should be processed exactly once.

How should the JMS Listener be configured to maximize performance and concurrent message processing of the JMS queue?

- A. Set numberOfConsumers = 1 Set primaryNodeOnly = false
- B. Set numberOfConsumers = 1 Set primaryNodeOnly = true
- C. Set numberOfConsumers to a value greater than one Set primaryNodeOnly = true
- D. Set numberOfConsumers to a value greater than one Set primaryNodeOnly = false

Answer: D

NEW QUESTION 288

49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58.

<https://api.acme.com/partnefs/v1>. The API implementation is estimated to require deployment to 5 CloudHub workers.

The retailer has obtained a public X.509 certificate for the name apl.acme.com, signed by a reputable CA, to be used as the server certificate.

Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

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

Answer: C

NEW QUESTION 290

A stock broking company makes use of CloudHub VPC to deploy Mule applications. Mule application needs to connect to a database application in the customers on-premises corporate data center and also to a Kafka cluster running in AWS VPC.

How is access enabled for the API to connect to the database application and Kafka cluster securely?

- A. Set up a transit gateway to the customers on-premises corporate datacenter to AWS VPC
- B. Setup AnyPoint VPN to the customer's on-premise corporate data center and VPC peering with AWS VPC
- C. Setup VPC peering with AWS VPC and the customers devices corporate data center
- D. Setup VPC peering with the customers onto my service corporate data center and Anypoint VPN to AWS VPC

Answer: B

NEW QUESTION 295

When a Mule application using VM queues is deployed to a customer-hosted cluster or multiple CloudHub v1.0 workers/replicas, how are messages consumed across the nodes?

- A. Sequentially, from a dedicated Anypoint MQ queue
- B. Sequentially, only from the primary node
- C. In a non-deterministic way
- D. Round-robin, within an XA transaction

Answer: C

NEW QUESTION 297

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 sequentiallyAny 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 parallelAny 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 sequentiallyAny 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 parallelAny error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations

Answer: A

NEW QUESTION 301

What Anypoint Connectors support transactions?

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

Answer: A

NEW QUESTION 306

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

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

Answer: C

NEW QUESTION 310

During a planning session with the executive leadership, the development team director presents plans for a new API to expose the data in the company's order database. An earlier effort to build an API on top of this data failed, so the director is recommending a design-first approach.

Which characteristics of a design-first approach will help make this API successful?

- A. Building MUnit tests so administrators can confirm code coverage percentage during deployment
- B. Publishing the fully implemented API to Exchange so all developers can reuse the API
- C. Adding global policies to the API so all developers automatically secure the implementation before coding anything
- D. Developing a specification so consumers can test before the implementation is built

Answer: D

NEW QUESTION 312

The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure. Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).

Given the company's current setup and requirements, what is the most idiomatic (used for its intended purpose) way to send Mule application logs to the external LMT?

- A. In RTF-VM, install and configure the external LMT's log-forwarding agent
- B. In RTF-VM, edit the pod configuration to automatically install and configure an Anypoint Monitoring agent
- C. In each Mule application, configure custom Log4j settings
- D. In RTF-V
- E. configure the out-of-the-box external log forwarder

Answer: A

NEW QUESTION 313

An organization uses one specific CloudHub (AWS) region for all CloudHub deployments. How are CloudHub workers assigned to availability zones (AZs) when the organization's Mule applications are deployed to CloudHub in that region?

- A. Workers belonging to a given environment are assigned to the same AZ within that region.
- B. AZs are selected as part of the Mule application's deployment configuration.
- C. Workers are randomly distributed across available AZs within that region.
- D. An AZ is randomly selected for a Mule application, and all the Mule application's CloudHub workers are assigned to that one AZ

Answer: C

NEW QUESTION 316

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 317

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 319

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

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

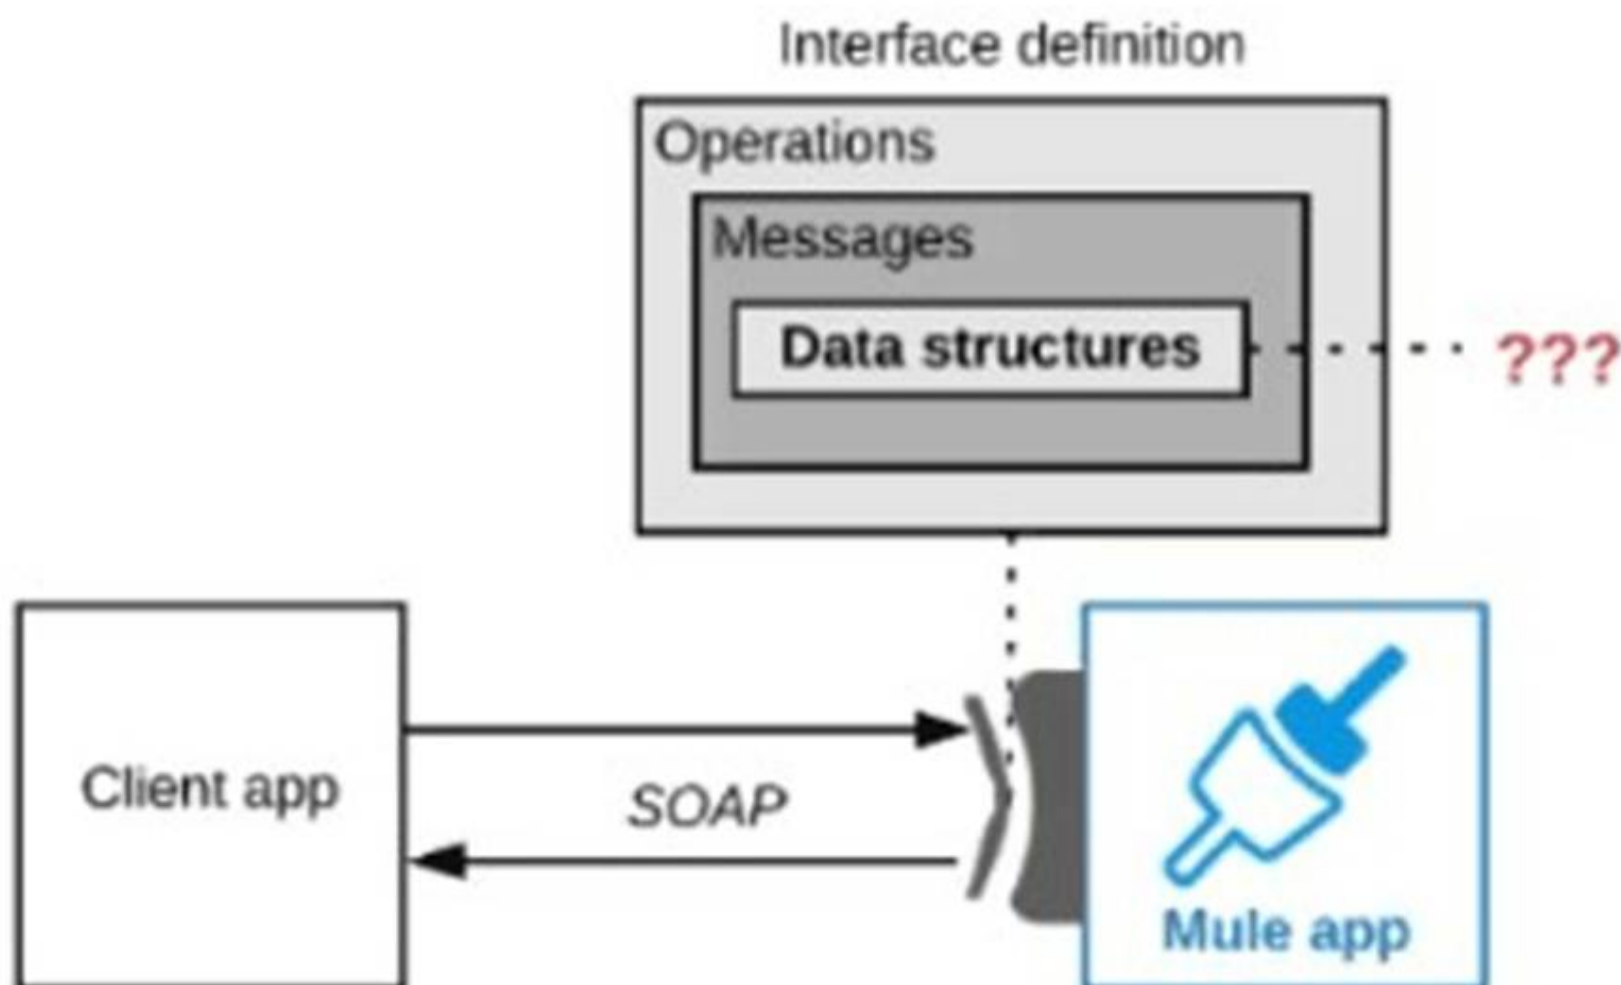
Answer: C

NEW QUESTION 321

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 326

One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make. Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments. The production environment however does not have any rate limiting. What is the cost-effective approach to conduct performance test of the API implementation in the non-production staging environment?

- A. Including logic within the API implementation that bypasses in locations of the back-end system in the staging environment and invoke a Mocking service that replicates typical back-end system responsesThen conduct performance test using this API implementation
- B. Use MUnit to simulate standard responses from the back-end syste
- C. Then conduct performance test to identify other bottlenecks in the system
- D. Create a Mocking service that replicates the back-end system'sproduction performance characteristicsThen configure the API implementation to use the mocking service and conduct the performance test
- E. Conduct scaled-down performance tests in the staging environment against rate-limiting back-end syste
- F. Then upscale performance results to full production scale

Answer: C

NEW QUESTION 330

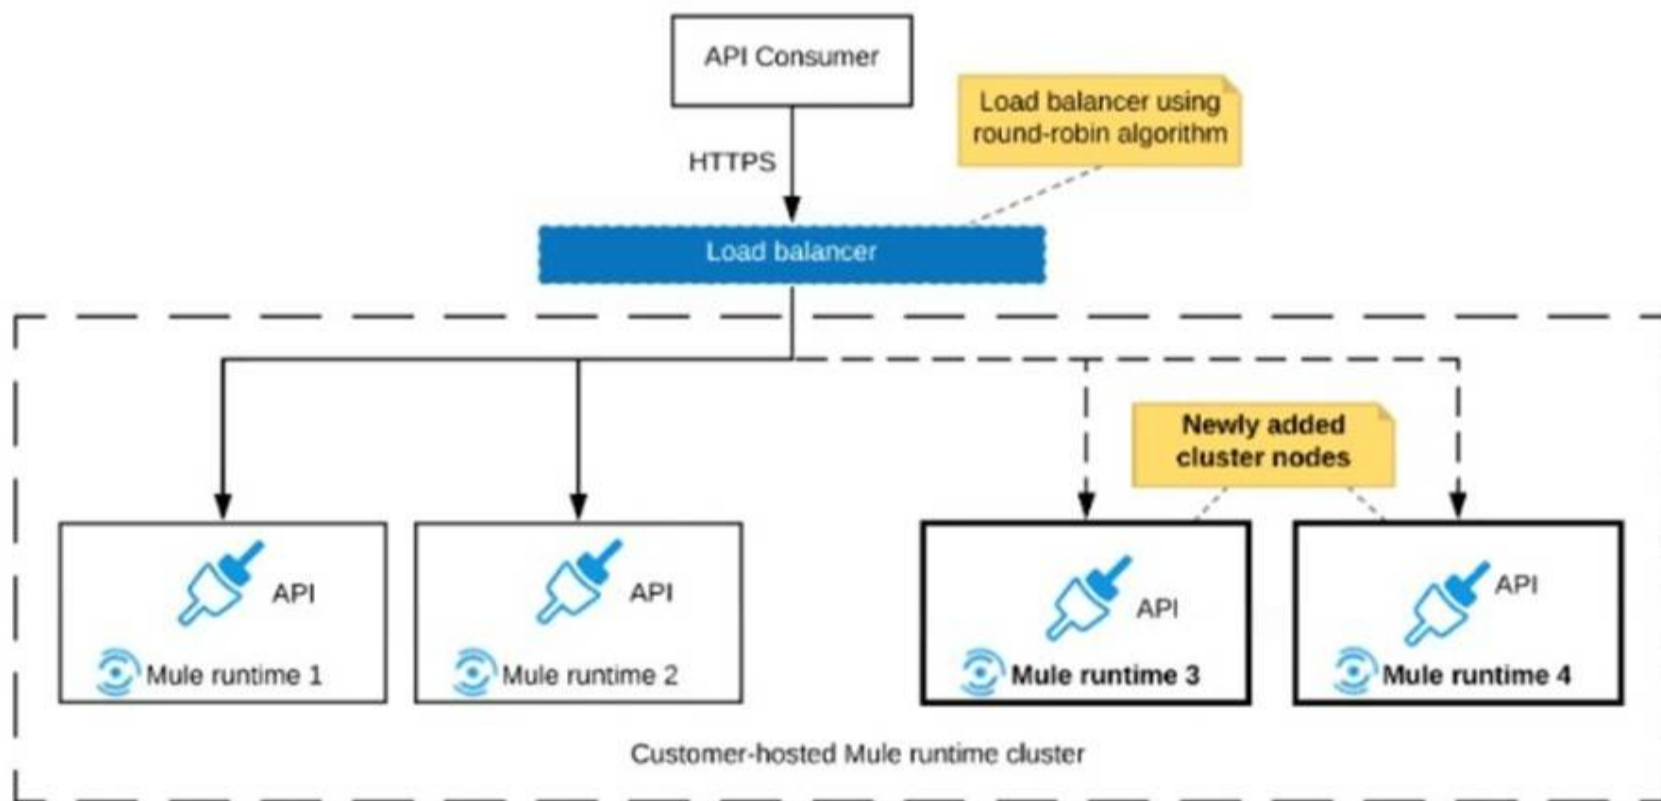
An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments. What must the testing team do before they can start integration testing the API in the Staging environment?

- 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 334

Refer to the exhibit.



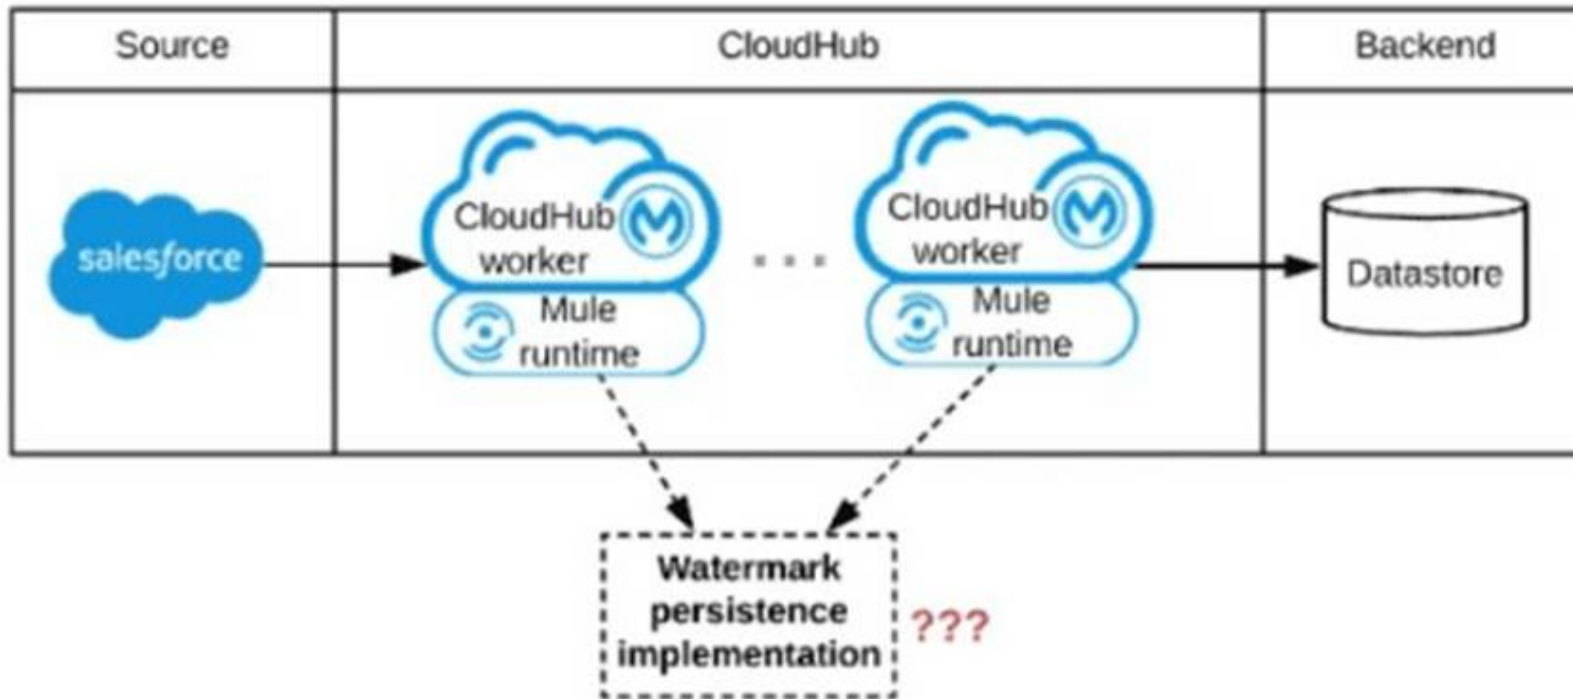
An organization uses a 2-node Mute runtime cluster to host one stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution. Two additional nodes have been added to the cluster and the load balancer has been configured to recognize the new nodes with no other change to the load balancer. What average performance change is guaranteed to happen, assuming all cluster nodes are fully operational?

- A. 50% reduction in the response time of the API
- B. 100% increase in the throughput of the API
- C. 50% reduction in the JVM heap memory consumed by each node
- D. 50% reduction in the number of requests being received by each node

Answer: D

NEW QUESTION 336

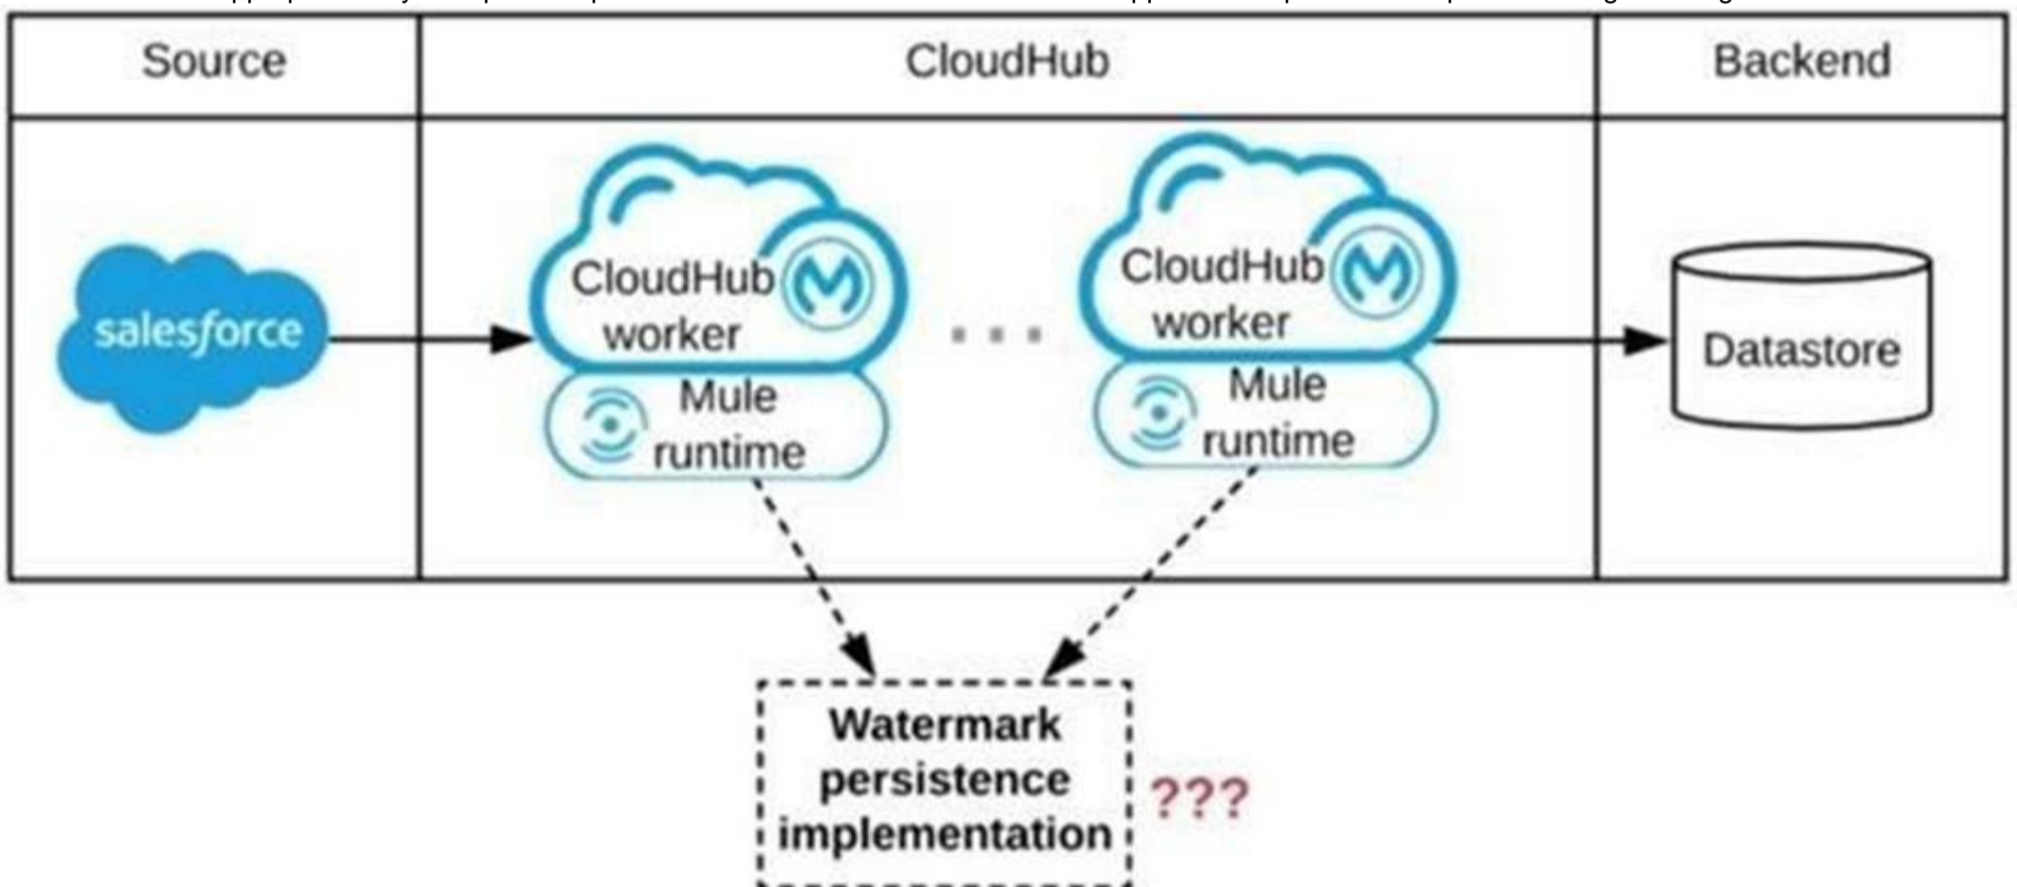
Refer to the exhibit.



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 339

What aspects of a CI/CD pipeline for Mule applications can be automated using MuleSoft- provided Maven plugins?

- A. Compile, package, unit test, validate unit test coverage, deploy
- B. Compile, package, unit test, deploy, integration test (Incorrect)
- C. Compile, package, unit test, deploy, create associated API instances in API Manager
- D. Import from API designer, compile, package, unit test, deploy, publish to Anypoint Exchange

Answer: A

NEW QUESTION 342

An integration Mute application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mute application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system. Orders that cannot be successfully

submitted due to rejections from the backend system will need to be processed manually (outside the backend system).

The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed.

The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages.

What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

- A. An On Error scope Non-persistent VM ActiveMQ Dead Letter Queue for manual processing

- B. An On Error scope MuleSoft Object Store ActiveMQ Dead Letter Queue for manual processing
- C. Until Successful component MuleSoft Object Store ActiveMQ is NOT needed or used
- D. Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing

Answer: D

NEW QUESTION 345

An organization designing a hybrid, load balanced, single cluster production environment. Due to performance service level agreement goals, it is looking into running the Mule applications in an active-active multi node cluster configuration. What should be considered when running its Mule applications in this type of environment?

- A. All event sources, regardless of time, can be configured as the target source by the primary node in the cluster
- B. An external load balancer is required to distribute incoming requests throughout the cluster nodes
- C. A Mule application deployed to multiple nodes runs in an isolation from the other nodes in the cluster
- D. Although the cluster environment is fully installed configured and running, it will not process any requests until an outage condition is detected by the primary node in the cluster.

Answer: B

NEW QUESTION 349

A developer is examining the responses from a RESTful web service that is compliant with the Hypertext Transfer Protocol (HTTP/1.1) as defined by the Internet Engineering Task Force (IETF).

In this HTTP/1.1-compliant web service, which class of HTTP response status codes should be specified to indicate when client requests are successfully received, understood, and accepted by the web service?

- A. 3xx
- B. 2xx
- C. 4xx
- D. 5xx

Answer: B

NEW QUESTION 351

An organization uses a four(4) node customer hosted Mule runtime cluster to host one(1) stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution. Each node in the cluster has been sized to be able to accept four(4) times the current number of requests.

Two(2) nodes in the cluster experience a power outage and are no longer available. The load balancer directs the outage and blocks the two unavailable nodes from receiving further HTTP requests.

What performance-related consequence is guaranteed to happen to average, assuming the remaining cluster nodes are fully operational?

- A. 100% increase in the average response time of the API
- B. 50% reduction in the throughput of the API
- C. 100% increase in the number of requests received by each remaining node
- D. 50% increase in the JVM heap memory consumed by each remaining node

Answer: A

NEW QUESTION 352

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 356

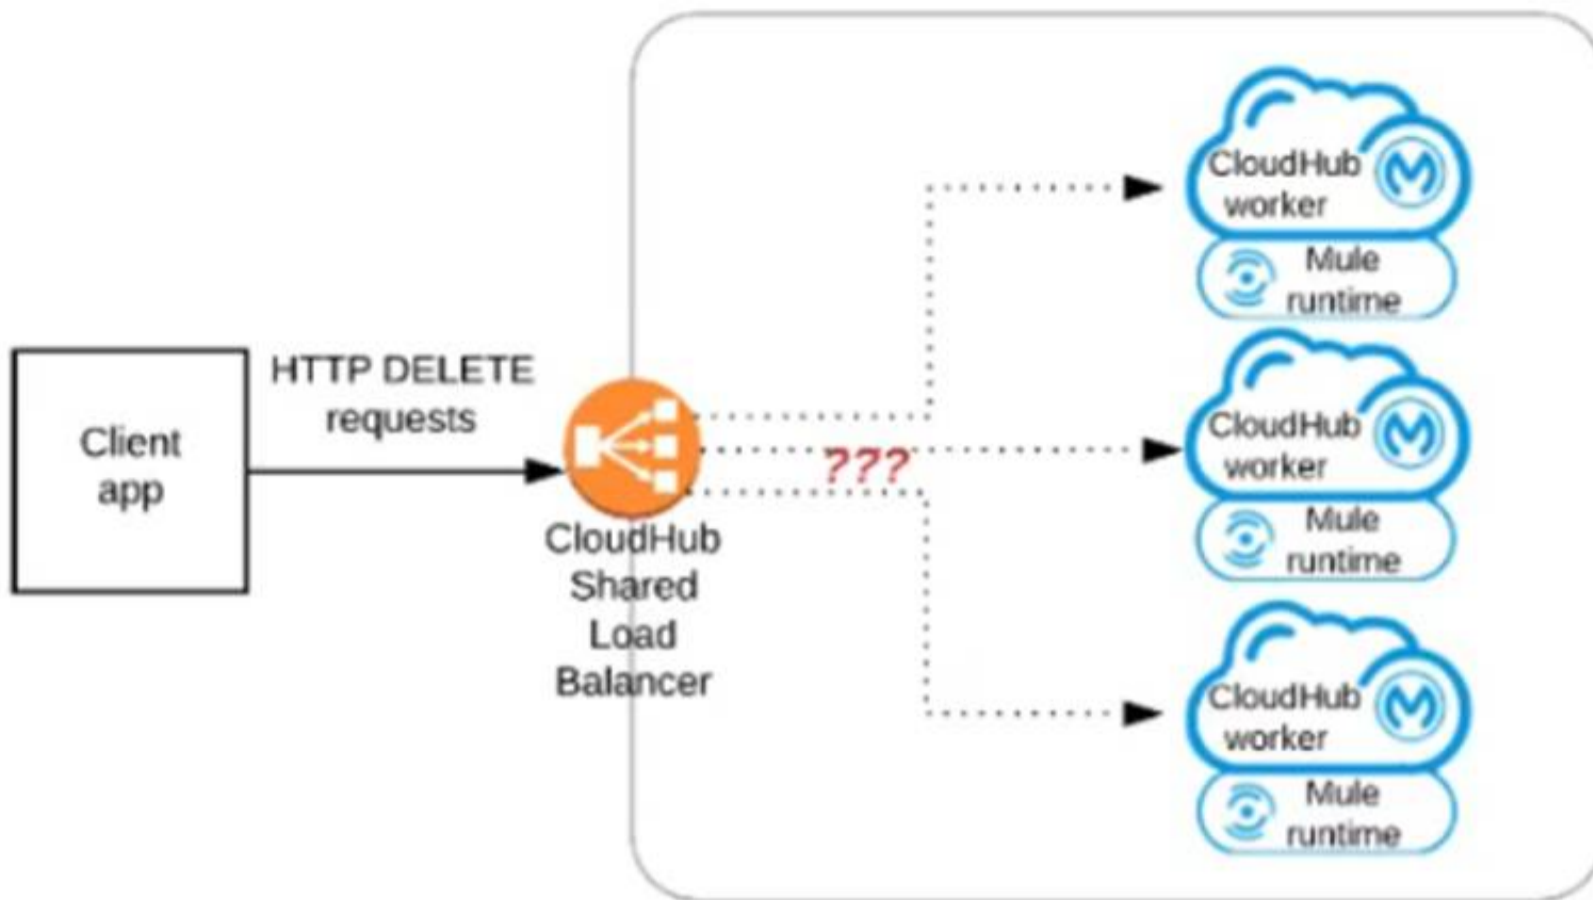
What is a defining characteristic of an integration-Platform-as-a-Service (iPaaS)?

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

Answer: A

NEW QUESTION 361

Refer to the exhibit.



A Mule application has an HTTP Listener that accepts HTTP DELETE requests. This Mule application is deployed to three CloudHub workers under the control of the CloudHub Shared Load Balancer.

A web client makes a sequence of requests to the Mule application's public URL.

How is this sequence of web client requests distributed among the HTTP Listeners running in the three CloudHub workers?

- A. Each request is routed to the PRIMARY CloudHub worker in the PRIMARY Availability Zone (AZ)
- B. Each request is routed to ONE ARBITRARY CloudHub worker in the PRIMARY Availability Zone (AZ)
- C. Each request is routed to ONE ARBITRARY CloudHub worker out of ALL three CloudHub workers
- D. Each request is routed (scattered) to ALL three CloudHub workers at the same time

Answer: C

NEW QUESTION 362

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 365

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 370

An organization is designing Mule application which connects to a legacy backend. It has been reported that backend services are not highly available and experience downtime quite often. As an integration architect which of the below approach you would propose to achieve high reliability goals?

- A. Alerts can be configured in Mule runtime so that backend team can be communicated when services are down
- B. Until Successful scope can be implemented while calling backend API's
- C. On Error Continue scope to be used to call in case of error again
- D. Create a batch job with all requests being sent to backend using that job as per the availability of backend API's

Answer: B

NEW QUESTION 374

A leading e-commerce giant will use Mulesoft API's on runtime fabric (RTF) to process customer orders. Some customer's sensitive information such as credit card information is also there as a part of a API payload.

What approach minimizes the risk of matching sensitive data to the original and can convert back to the original value whenever and wherever required?

- A. Apply masking to hide the sensitive information and then use API
- B. manager to detokenize the masking format to return the original value
- C. create a tokenization format and apply a tokenization policy to the API Gateway
- D. Used both masking and tokenization
- E. Apply a field level encryption policy in the API Gateway

Answer: C

NEW QUESTION 375

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 379

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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 FirstTry
- * MuleSoft-Integration-Architect-I Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The MuleSoft-Integration-Architect-I Practice Test Here](#)