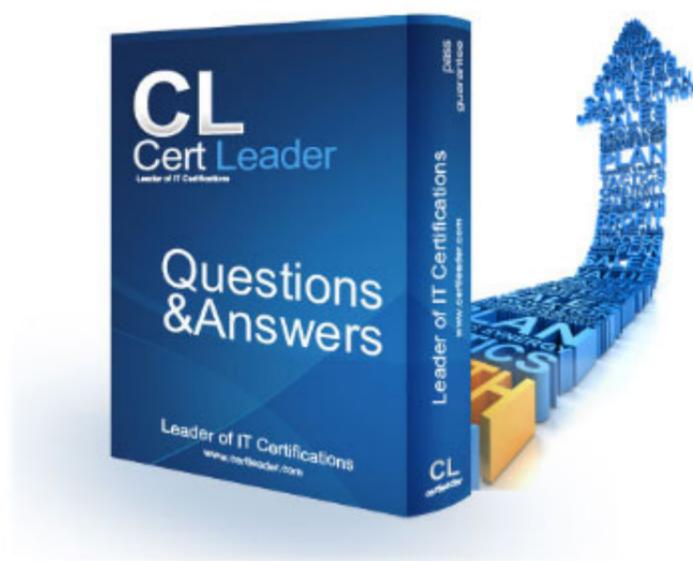


MuleSoft-Integration-Architect-I Dumps

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam

<https://www.certleader.com/MuleSoft-Integration-Architect-I-dumps.html>



NEW QUESTION 1

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 2

Insurance organization is planning to deploy Mule application in MuleSoft Hosted runtime plane. As a part of requirement, application should be scalable, highly available. It also has regulatory requirement which demands logs to be retained for at least 2 years. As an Integration Architect what step you will recommend in order to achieve this?

- A. It is not possible to store logs for 2 years in CloudHub deployment
- B. External log management system is required.
- C. When deploying an application to CloudHub, logs retention period should be selected as 2 years
- D. When deploying an application to CloudHub, worker size should be sufficient to store 2 years data
- E. Logging strategy should be configured accordingly in log4j file deployed with the application.

Answer: A

NEW QUESTION 3

When designing an upstream API and its implementation, the development team has been advised to not set timeouts when invoking downstream API. Because the downstream API has no SLA that can be relied upon. This is the only downstream API dependency of that upstream API. Assume the downstream API runs uninterrupted without crashing. What is the impact of this advice?

- A. The invocation of the downstream API will run to completion without timing out.
- B. An SLA for the upstream API CANNOT be provided.
- C. A default timeout of 500 ms will automatically be applied by the Mule runtime in which the upstream API implementation executes.
- D. A load-dependent timeout of less than 1000 ms will be applied by the Mule runtime in which the downstream API implementation executes.

Answer: B

NEW QUESTION 4

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

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

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

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

Answer: D

NEW QUESTION 5

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

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

Answer: A

NEW QUESTION 6

A Mule application name Pub uses a persistence object store. The Pub Mule application is deployed to Cloudhub and it is 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 7

A Mule application is synchronizing customer data between two different database systems. What is the main benefit of using XA transaction over local transactions to synchronize these two database system?

- A. Reduce latency
- B. Increase throughput
- C. Simplifies communication
- D. Ensure consistency

Answer: D

NEW QUESTION 8

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

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

Answer: A

NEW QUESTION 9

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 10

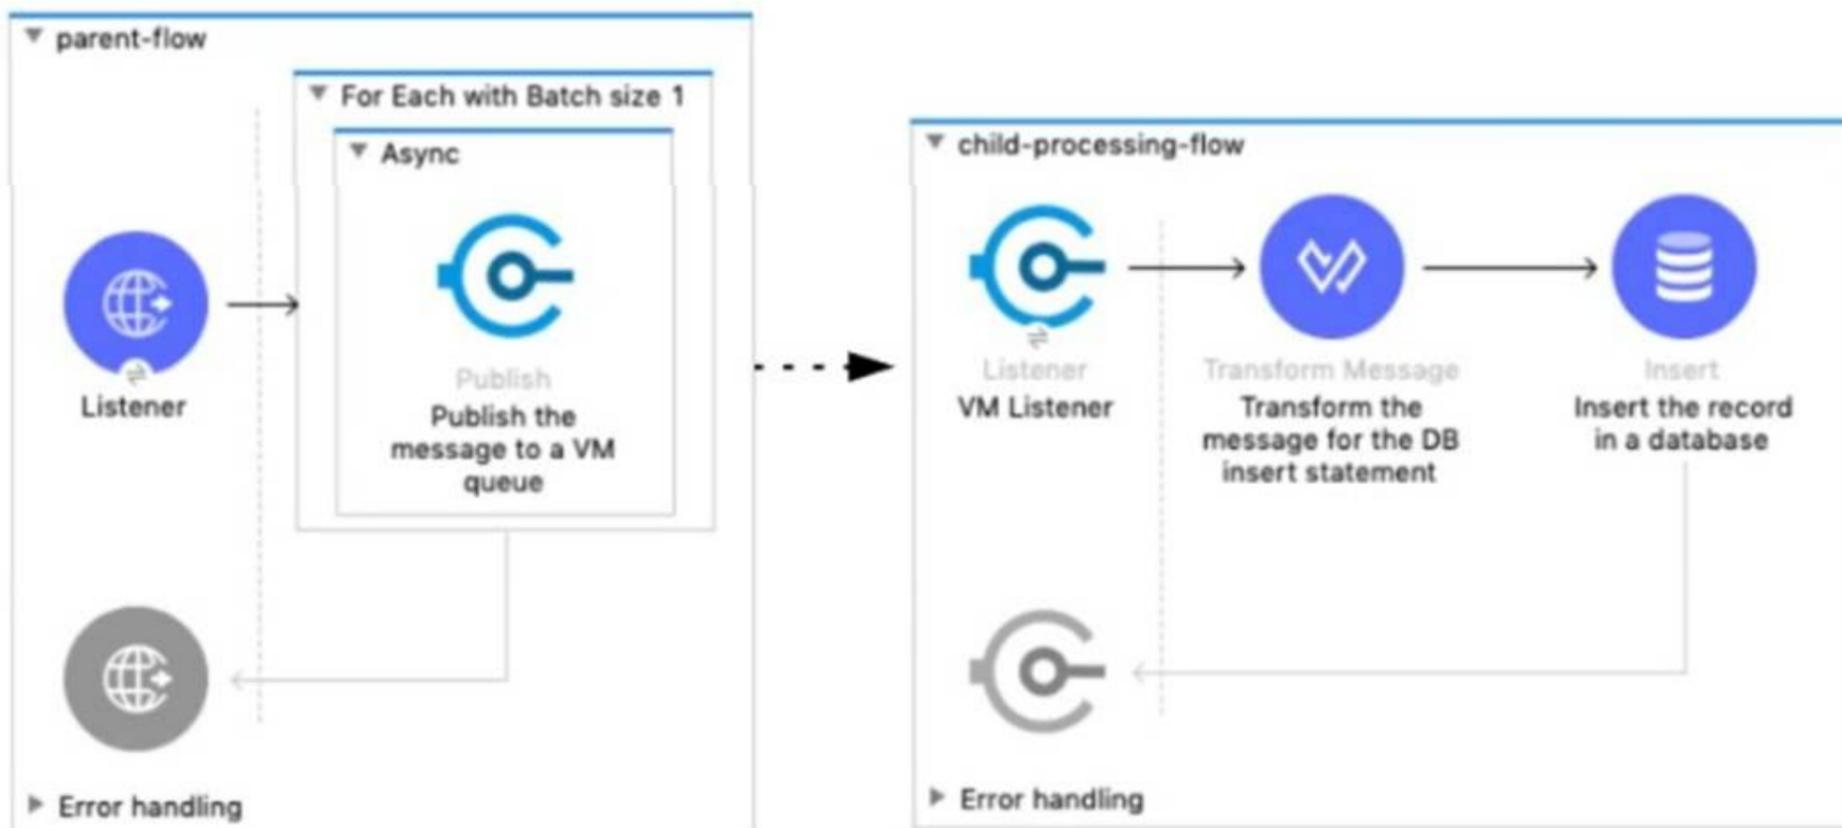
When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

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

Answer: B

NEW QUESTION 10

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 14

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 19

Which Anypoint Platform component should a MuleSoft developer use to create an API specification prior to building the API implementation?

- A. MUnit
- B. API Designer
- C. API Manager
- D. Runtime Manager

Answer: B

NEW QUESTION 21

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 24

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 25

A company is planning to migrate its deployment environment from on-premises cluster to a Runtime Fabric (RTF) cluster. It also has a requirement to enable Mule applications deployed to a Mule runtime instance to store and share data across application replicas and restarts.

How can these requirements be met?

- A. Anypoint object store V2 to share data between replicas in the RTF cluster
- B. Install the object store pod on one of the cluster nodes
- C. Configure Persistence Gateway in any of the servers using Mule Object Store
- D. Configure Persistent Gateway at the RTF

Answer: A

NEW QUESTION 28

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 32

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 35

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

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

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

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

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

Answer: C

NEW QUESTION 39

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

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

Answer: C

NEW QUESTION 43

An organization plans to migrate all its Mule applications to Runtime Fabric (RTF). Currently, all Mule applications have been deployed to CloudHub using automated CI/CD scripts.

What steps should be taken to properly migrate the applications from CloudHub to RTF, while keeping the same automated CI/CD deployment strategy?

- A. A runtimefabric dependency should be added as a mule-plugin to the pom.xml file in all the Mule applications.
- B. runtimeFabric command-line parameter should be added to the CI/CD deployment scripts.
- C. A runtimeFabricDeployment profile should be added to Mule configuration properties YAML files in all the Mule applications. CI/CD scripts must be modified to use the new configuration properties.
- D. runtimeFabricDeployment profile should be added to the pom.xml file in all the Mule application
- E. CI/CD scripts must be modified to use the new RTF profile.
- F. - The pom.xml and Mule configuration YAML files can remain unchanged in each Mule application. A --runtimeFabric command-line parameter should be added to the CI/CD deployment scripts

Answer: D

NEW QUESTION 48

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 53

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 57

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 58

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 59

What are two reasons why a typical MuleSoft customer favors a MuleSoft-hosted Anypoint Platform runtime plane over a customer-hosted runtime for its Mule application deployments? (Choose two.)

- A. Reduced application latency
- B. Increased application isolation
- C. Reduced time-to-market for the first application
- D. Increased application throughput
- E. Reduced IT operations effort

Answer: CE

NEW QUESTION 62

As part of a growth strategy, a supplier signs a trading agreement with a large customer. The customer sends purchase orders to the supplier according to the ANSI X12 EDI standard, and the supplier creates the orders in its ERP system using the information in the EDI document.

The agreement also requires that the supplier provide a new RESTful API to process request from the customer for current product inventory level from the supplier's ERP system.

Which two fundamental integration use cases does the supplier need to deliver to provide an end-to-end solution for this business scenario? (Choose two.)

- A. Synchronized data transfer
- B. Sharing data with external partners
- C. User interface integration
- D. Streaming data ingestion
- E. Data mashups

Answer: AB

NEW QUESTION 64

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 now the primary node
- D. Files are polled from Bob node Bob is now the primary node

Answer: D

NEW QUESTION 65

An organization has deployed runtime fabric on an eight node cluster with performance profile. An API uses an in-memory object store for maintaining some of its state data. What will be the impact to the state data if server crashes?

- A. State data is preserved
- B. State data is rolled back to a previously saved version
- C. State data is lost
- D. State data is preserved as long as more than one node is unaffected by the crash

Answer: C

NEW QUESTION 66

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

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

Answer: C

NEW QUESTION 70

What is an advantage of using OAuth 2.0 client credentials and access tokens over only API keys for API authentication?

- A. If the access token is compromised, the client credentials do not have to be reissued.
- B. If the access token is compromised, it can be exchanged for an API key.
- C. If the client ID is compromised, it can be exchanged for an API key
- D. If the client secret is compromised, the client credentials do not have to be reissued.

Answer: A

NEW QUESTION 74

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 76

A large life sciences customer plans to use the Mule Tracing module with the Mapped Diagnostic Context (MDC) logging operations to enrich logging in its Mule application and to improve tracking by providing more context in the Mule application logs. The customer also wants to improve throughput and lower the message processing latency in its Mule application flows.

After installing the Mule Tracing module in the Mule application, how should logging be performed in flows in Mule applications, and what should be changed in the log4j2.xml files?

- A. In the flows, add Mule Tracing module Set logging variable operations before any Core Logger components. In log4j2.xml files, change the appender's pattern layout to use %MDC and then assign the appender to a Logger or Root element.
- B. In the flows, add Mule Tracing module Set logging variable operations before any Core Logger components. In log4j2.xml files, change the appender's pattern

layout to use the %MDC placeholder and then assign the appender to an AsyncLogger element.
C. In the flows, add Mule Tracing module Set logging variable operations before any Core Logger components. In log4j2.xml files, change the appender's pattern layout to use %asyncLogger placeholder and then assign the appender to an AsyncLogger element.
D. In the flows, wrap Logger components in Async scope
E. In log4j2.xml files, change the appender's pattern layout to use the %asyncLoggerplaceholder and then assign the appender to a Logger or Root element.

Answer: A

NEW QUESTION 78

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 81

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 86

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

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

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

Answer: C

NEW QUESTION 87

An organization's IT team follows an API-led connectivity approach and must use Anypoint Platform to implement a System API that securely accesses customer data. The organization uses Salesforce as the system of record for all customer data, and its most important objective is to reduce the overall development time to release the System API.

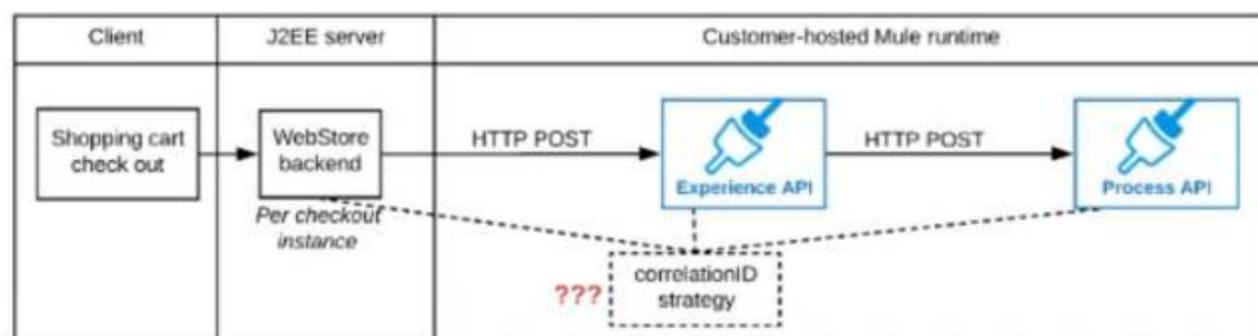
The team's integration architect has identified four different approaches to access the customer data from within the implementation of the System API by using different Anypoint Connectors that all meet the technical requirements of the project.

- A. Use the Anypoint Connector for Database to connect to a MySQL database to access a copy of the customer data
- B. Use the Anypoint Connector for HTTP to connect to the Salesforce APIs to directly access the customer data
- C. Use the Anypoint Connector for Salesforce to connect to the Salesforce APIs to directly access the customer data
- D. Use the Anypoint Connector for FTP to download a file containing a recent near-real time extract of the customer data

Answer: C

NEW QUESTION 92

Refer to the exhibit.



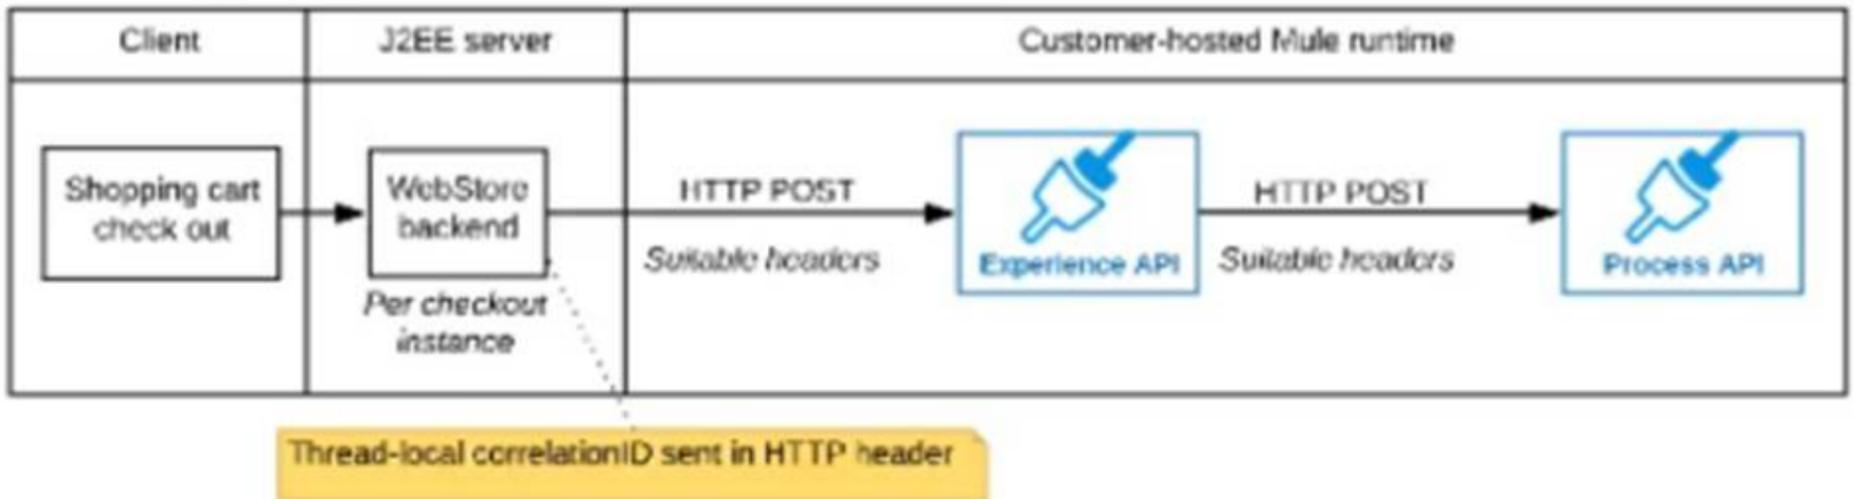
A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

End-to-end correlation of all HTTP requests and responses belonging to each individual checkout Instance is required. This is to be done through a common

correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the same correlation ID for all requests and responses belonging to the same checkout instance.

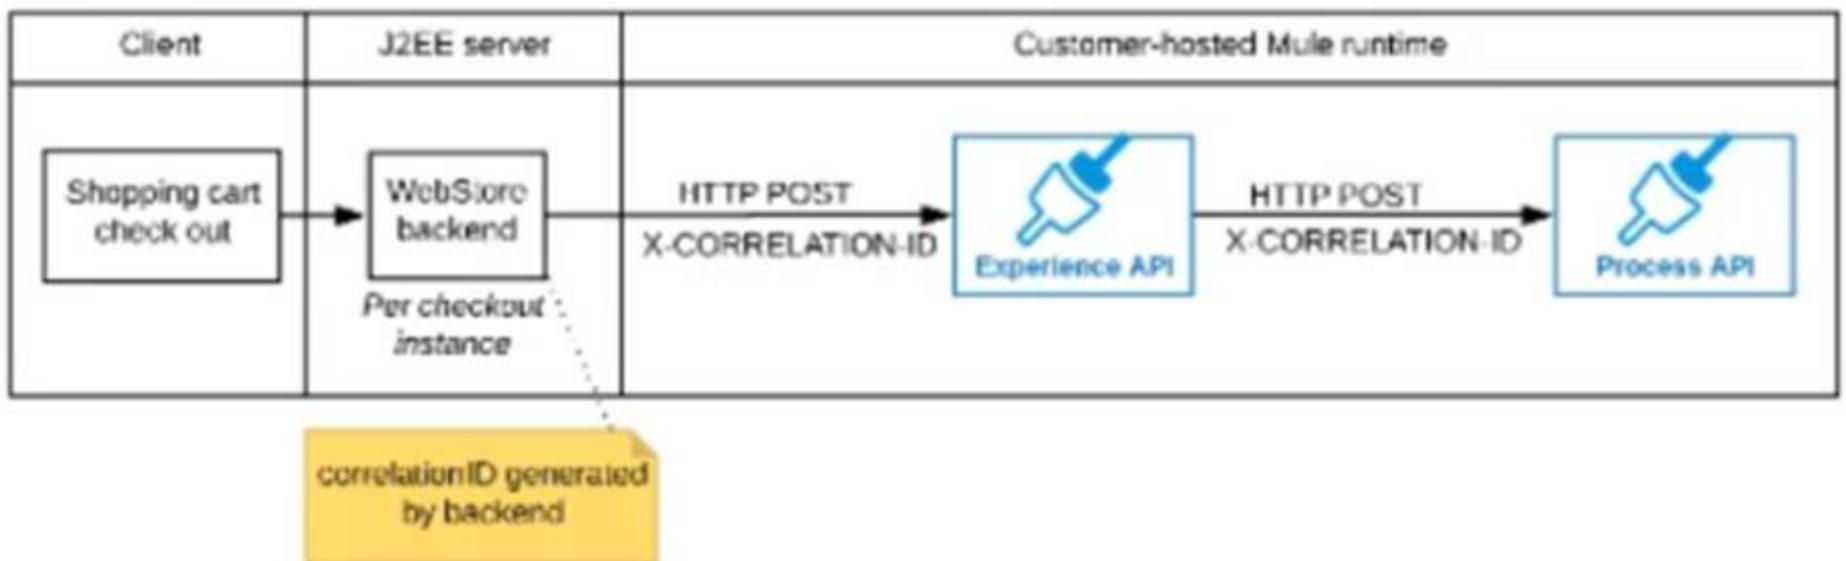
What is the most efficient way (using the least amount of custom coding or configuration) for the web store backend and the implementations of the Experience API and Process API to participate in end-to-end correlation of the API invocations for each checkout instance?

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



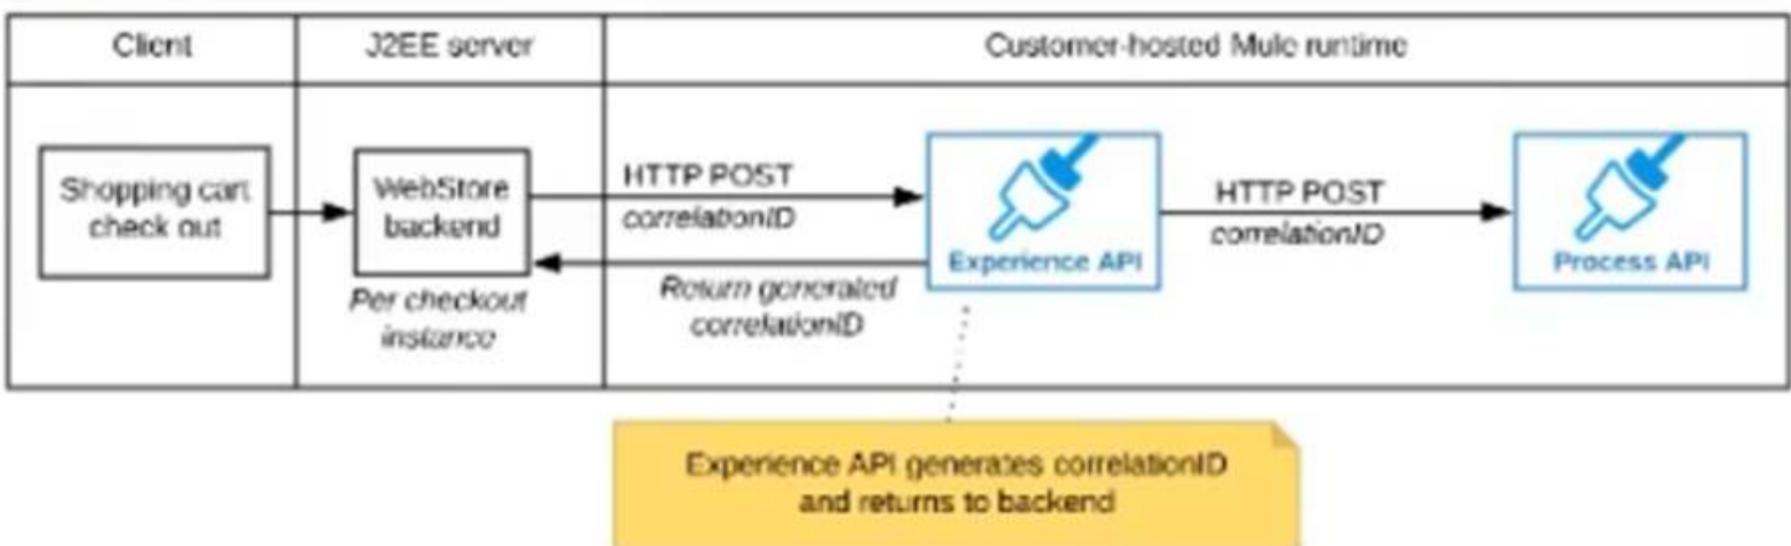
B) The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header In each API invocation belonging to that checkout

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



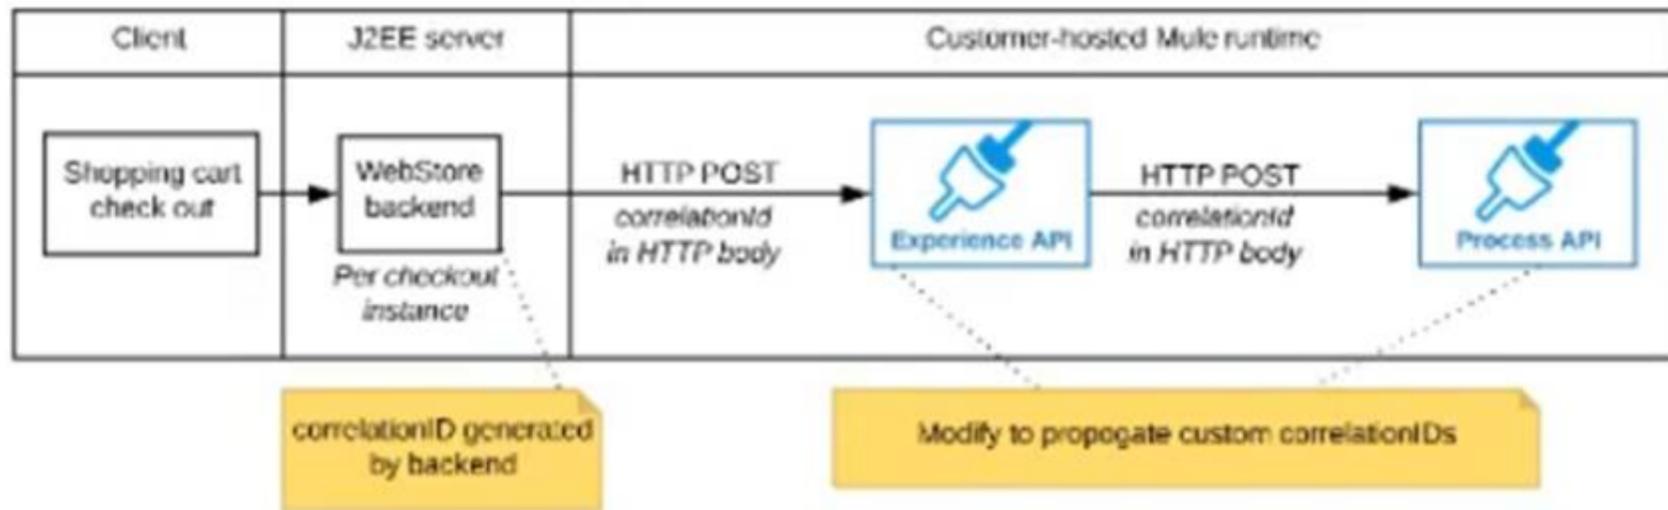
C) The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

The Experience API implementation must be coded to also propagate the correlation ID to the Process API in a suitable HTTP request header



D) The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API

The Experience API and Process API implementations must be coded to receive the custom correlation ID In the HTTP requests and propagate It in suitable HTTP request headers



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

Answer: B

NEW QUESTION 95

An API client is implemented as a Mule application that includes an HTTP Request operation using a default configuration. The HTTP Request operation invokes an external API that follows standard HTTP status code conventions, which causes the HTTP Request operation to return a 4xx status code. What is a possible cause of this status code response?

- A. An error occurred inside the external API implementation when processing the HTTP request that was received from the outbound HTTP Request operation of the Mule application
- B. The external API reported that the API implementation has moved to a different external endpoint
- C. The HTTP response cannot be interpreted by the HTTP Request operation of the Mule application after it was received from the external API
- D. The external API reported an error with the HTTP request that was received from the outbound HTTP Request operation of the Mule application

Answer: D

NEW QUESTION 97

A stock trading company handles millions of trades a day and requires excellent performance and reliability within its stock trading system. The company operates a number of event-driven APIs implemented as Mule applications that are hosted on various customer-hosted Mule clusters and needs to enable message exchanges between the APIs within their internal network using shared message queues. What is an effective way to meet the cross-cluster messaging requirements of its event-driven APIs?

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

Answer: C

NEW QUESTION 99

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

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

Answer: D

NEW QUESTION 102

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

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

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

Answer: D

NEW QUESTION 103

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 106

In preparation for a digital transformation initiative, an organization is reviewing related IT integration projects that failed for various for reason. According to MuleSoft??s surveys of global IT leaders, what is a common cause of IT project failure that this organization may likely discover in its assessment?

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

Answer: D

NEW QUESTION 111

An organization plans to migrate its deployment environment from an onpremises cluster to a Runtime Fabric (RTF) cluster. The on-premises Mule applications are currently configured with persistent object stores.

There is a requirement to enable Mule applications deployed to the RTF cluster to store and share data across application replicas and through restarts of the entire RTF cluster, How can these reliability requirements be met?

- A. Replace persistent object stores with persistent VM queues in each Mule application deployment
- B. Install the Object Store pod on one of the cluster nodes
- C. Configure Anypoint Object Store v2 to share data between replicas in the RTF cluster
- D. Configure the Persistence Gateway in the RTF installation

Answer: C

NEW QUESTION 114

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

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

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

Answer: D

NEW QUESTION 119

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 124

Which key DevOps practice and associated Anypoint Platform component should a MuteSoft integration team adopt to improve delivery quality?

- A. A Continuous design with API Designer
- B. Automated testing with MUnit
- C. Passive monitoring with Anypoint Monitoring
- D. Manual testing with Anypoint Studio

Answer: B

NEW QUESTION 129

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

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

Answer: A

NEW QUESTION 130

An external REST client periodically sends an array of records in a single POST request to a Mule application API endpoint.

The Mule application must validate each record of the request against a JSON schema before sending it to a downstream system in the same order that it was received in the array

Record processing will take place inside a router or scope that calls a child flow. The child flow has its own error handling defined. Any validation or communication failures should not prevent further processing of the remaining records.

To best address these requirements what is the most idiomatic(used for it intended purpose) router or scope to used in the parent flow, and what type of error handler should be used in the child flow?

- A. First Successful router in the parent flow On Error Continue error handler in the child flow
- B. For Each scope in the parent flow On Error Continue error handler in the child flow
- C. Parallel For Each scope in the parent flow On Error Propagate error handler in the child flow
- D. Until Successful router in the parent flow On Error Propagate error handler in the child flow

Answer: B

NEW QUESTION 134

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 136

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 worker New 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 worker New 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 worker New 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 worker New web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.

Answer: C

NEW QUESTION 137

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

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

Answer: D

NEW QUESTION 140

A company is using Mulesoft to develop API's and deploy them to Cloudhub and on premises targets. Recently it has decided to enable Runtime Fabric deployment option as well and infrastructure is set up for this option.

What can be used to deploy Runtime Fabric?

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

Answer: A

NEW QUESTION 145

An insurance company has an existing API which is currently used by customers. API is deployed to customer hosted Mule runtime cluster. The load balancer that is used to access any APIs on the mule cluster is only configured to point to applications hosted on the server at port 443.

Mule application team of a company attempted to deploy a second API using port 443 but the application will not start and checking logs shows an error indicating the address is already in use.

Which steps must the organization take to resolve this error and allow customers to access both the API's?

- A. Change the base path of the HTTP listener configuration in the second API to a different one from the first API
- B. Set HTTP listener configuration in both API's to allow for connections from multiple ports
- C. Move the HTTP listener configurations from the API's and package them in a mule domain project using port 443
- D. Set the HTTP listener of the second API to use different port than the one used in the first API

Answer: C

NEW QUESTION 150

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 152

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

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

Answer: A

NEW QUESTION 156

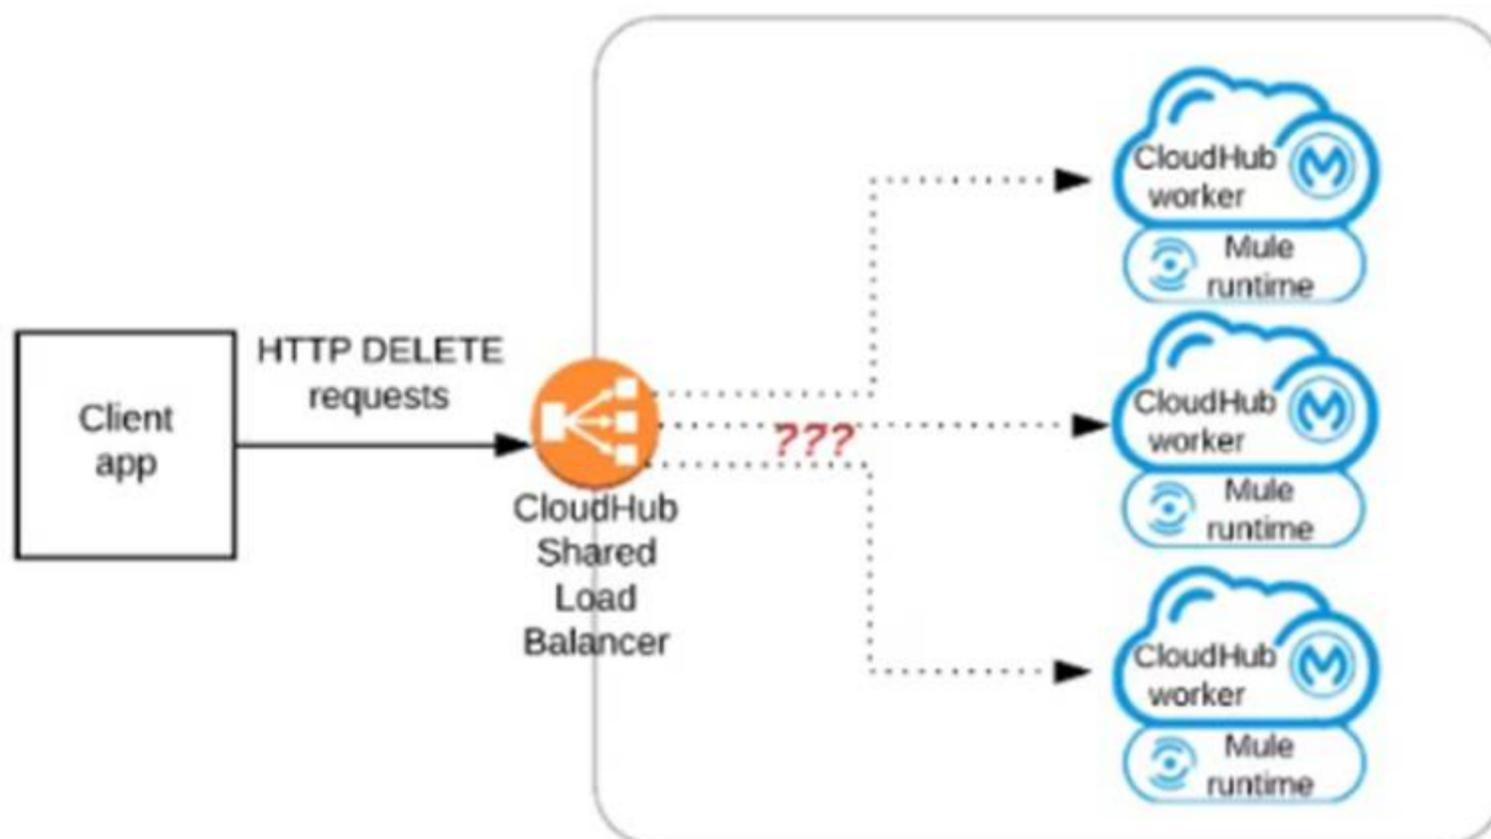
According to MuleSoft, what Action should an IT organization take regarding its technology assets in order to close the IT delivery.

- 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 161

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 162

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 165

What condition requires using a CloudHub Dedicated Load Balancer?

- A. When cross-region load balancing is required between separate deployments of the same Mule application
- B. When custom DNS names are required for API implementations deployed to customer- hosted Mule runtimes
- C. When API invocations across multiple CloudHub workers must be load balanced
- D. When server-side load-balanced TLS mutual authentication is required between API implementations and API clients

Answer: D

NEW QUESTION 166

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your MuleSoft-Integration-Architect-I Exam with Our Prep Materials Via below:

<https://www.certleader.com/MuleSoft-Integration-Architect-I-dumps.html>