MuleSoft

Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1

NEW QUESTION 1
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 cache Use an object store caching strategy and set expiration interval to 1-hour

Answer: A

NEW QUESTION 2
A company is modernizing its legal systems lo accelerate access lo applications and data while supporting the adoption of new technologies. The key to achieving this business goal is unlocking the companies' key systems and dala including microservices miming under Docker and kubernetes containers using apis.
Considering the current aggressive backlog and project delivery requirements the company wants to take a strategic approach in the first phase of its transformation projects by quickly deploying API's in mule runtime that are able lo scale, connect to on premises systems and migrate as needed.
Which runtime deployment option supports company's goals?

A. Customer hosted self provisioned runtimes
B. Cloudhub runtimes
C. Runtime fabric on self managed Kubernetes
D. Runtime fabric on Vmware metal

Answer: C

NEW QUESTION 3
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

Explanation:
* Clustering is a group of servers or mule runtime which acts as a single unit.
* Mulesoft Enterprise Edition supports scalable clustering to provide high availability for the Mulesoft application.
* In simple terms, virtual servers composed of multiple nodes and they communicate and share information through a distributed shared memory grid.
* By default, Mulesoft ensures the High availability of applications if clustering implemented.
* Let's consider the scenario one of the nodes in cluster crashed or goes down and under maintenance. In such cases, Mulesoft will ensure that requests are processed by other nodes in the cluster. Mulesoft clustering also ensures that the request is load balanced between all the nodes in a cluster.
* Clustering is only supported by on-premise Mule runtime and it is not supported in Cloudhub.
Correct answer is External monitoring tools or log aggregators must be configured to recognize the new nodes
* Rest of the options are automatically taken care of when a new node is added in cluster.

NEW QUESTION 4
An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to custom hosted Mule runtimes (on-premises) while others execute in the MuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems.
How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access the backend systems?

A. Create a Mule domain project that maintains the credentials as Mule domain-shared resources Deploy the Mule applications to the Mule domain, so the credentials are available to the Mule applications
B. Store the credentials in properties files in a shared folder within the organization's data center Have the Mule applications load properties files from this shared location at startup
C. Segregate the credentials for each backend system into environment-specific properties files Package these properties files in each Mule application, from where they are loaded at startup
D. Configure or create a credentials service that returns the credentials for each backend system, and that is accessible from customer-hosted and MuleSoft-hosted Mule runtimes Have the Mule applications toad the properties at startup by invoking that credentials service

Answer: D

Explanation:
* "Create a Mule domain project that maintains the credentials as Mule domain-shared resources" is wrong as domain project is not supported in Cloudhub * We should Avoid Creating duplicates in each Mule application but below two options cause duplication of credentials - Store the credentials in properties files in a shared folder within the organization's data center. Have the Mule applications load properties files from this shared location at startup - Segregate the credentials for each backend system into environment-specific properties files. Package these properties files in each Mule application, from where they are loaded at startup So these are also wrong choices * Credentials service is the best approach in this scenario. Mule domain projects are not supported on CloudHub. Also its is not recommended to have multiple copies of configuration values as this makes difficult to maintain Use the Mule Credentials Vault to encrypt data in a .properties file. (In the context of this document, we refer to the .properties file simply as the properties file.) The properties file in Mule stores data as key-value pairs which may contain information such as usernames, first and last names, and credit card numbers. A Mule application may access this data as it processes messages, for example, to acquire login credentials for an external Web service. However, though this sensitive, private data must be stored in a properties file for Mule to

access, it must also be protected against unauthorized – and potentially malicious – use by anyone with access to the Mule application

NEW QUESTION 5
A trading company handles millions of requests a day. Due to nature of its business, it requires excellent performance and reliability within its application.
For this purpose, company uses a number of event-based API's hosted on various mule clusters that communicate across a shared message queue sitting within its network.
Which method should be used to meet the company's requirement for its system?

A. XA transactions and XA connected components
B. JMS transactions
C. JMS manual acknowledgements with a reliability pattern
D. VM queues with reliability pattern

Answer: C

NEW QUESTION 6
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

Explanation:
Scatter-Gather sends a request message to multiple targets concurrently. It collects the responses from all routes, and aggregates them into a single message.

NEW QUESTION 7
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 intented 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

Explanation:
When an operation in a Mule application fails to connect to an external server, the default behavior is for the operation to fail immediately and return a connectivity error. You can modify this default behavior by configuring a reconnection strategy for the operation. You can configure a reconnection strategy for an operation either by modifying the operation properties or by modifying the configuration of the global element for the operation. The following are the available reconnection strategies and their behaviors: None Is the default behavior, which immediately returns a connectivity error if the attempt to connect is unsuccessful Standard (reconnect) Sets the number of reconnection attempts and the interval at which to execute them before returning a connectivity error Forever (reconnect-forever) Attempts to reconnect continually at a given interval

NEW QUESTION 8
An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.
What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

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

Answer: B

NEW QUESTION 9
A Mule application is being designed To receive nightly a CSV file containing millions of records from an external vendor over SFTP, The records from the file need to be validated, transformed. And then written to a database. Records can be inserted into the database in any order.
In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

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

Answer: C

Explanation:
Correct answer is Use a Batch job scope to bulk insert records into the database
* Batch Job is most efficient way to manage millions of records. A few points to note here are as follows :

Reliability: If you want reliabilty while processing the records, i.e should the processing survive a runtime crash or other unhappy scenarios, and when restarted process all the remaining records, if yes then go for batch as it uses persistent queues.

Error Handling: In Parallel for each an error in a particular route will stop processing the remaining records in that route and in such case you'd need to handle it using on error continue, batch process does not stop during such error instead you can have a step for failures and have a dedicated handling in it.

Memory footprint: Since question said that there are millions of records to process, parallel for each will aggregate all the processed records at the end and can possibly cause Out Of Memory.

Batch job instead provides a BatchResult in the on complete phase where you can get the count of failures and success. For huge file processing if order is not a concern definitely go ahead with Batch Job

NEW QUESTION 10

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: B

NEW QUESTION 10

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: C

NEW QUESTION 13

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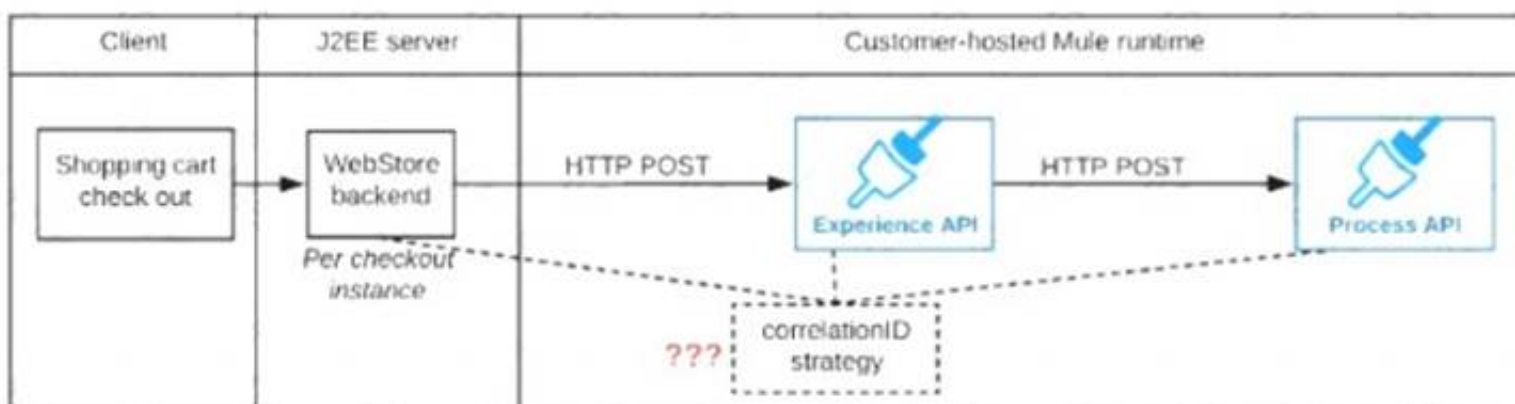
A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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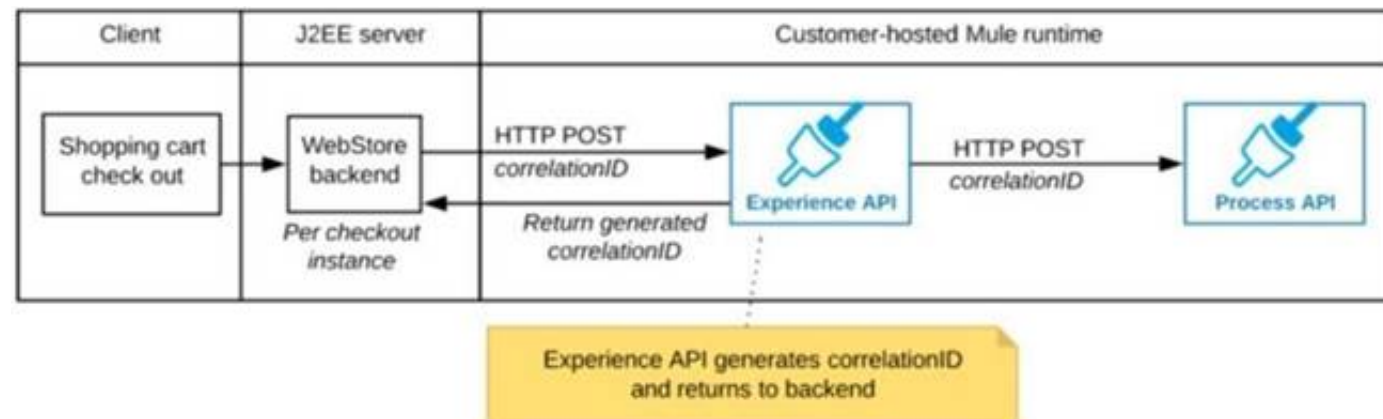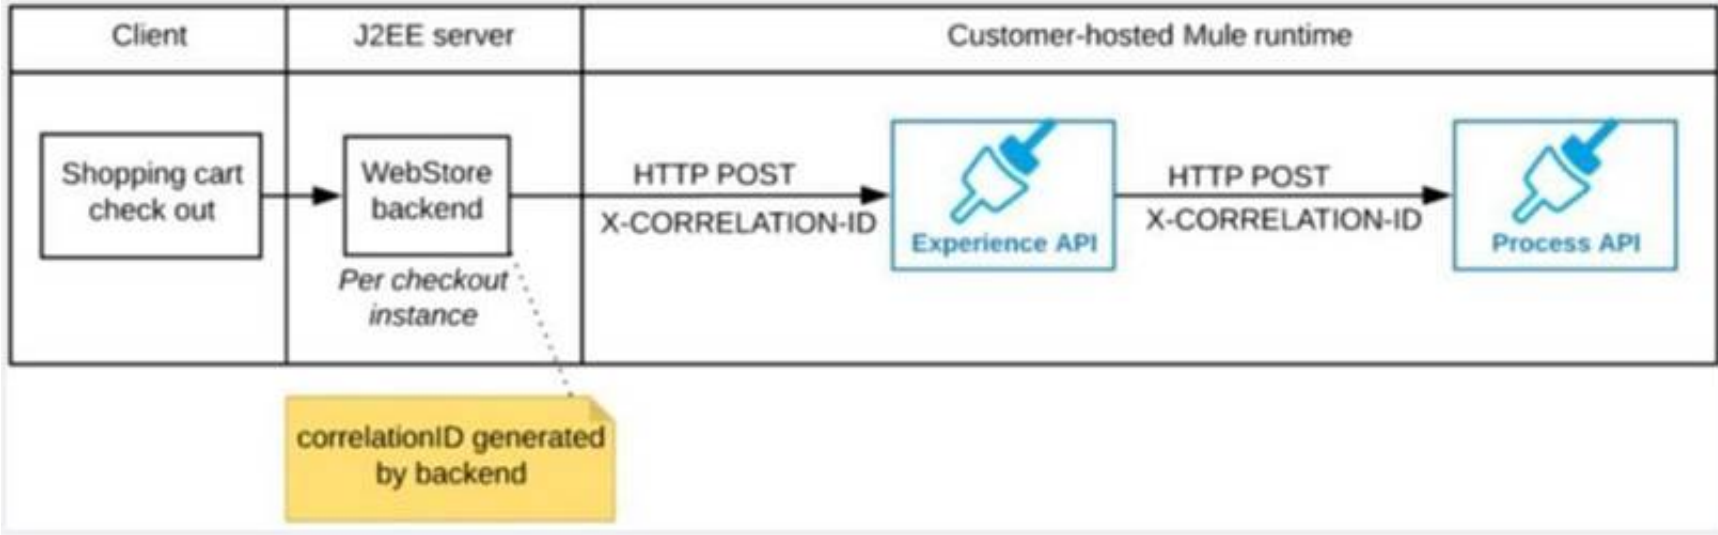

B)

The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-It 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
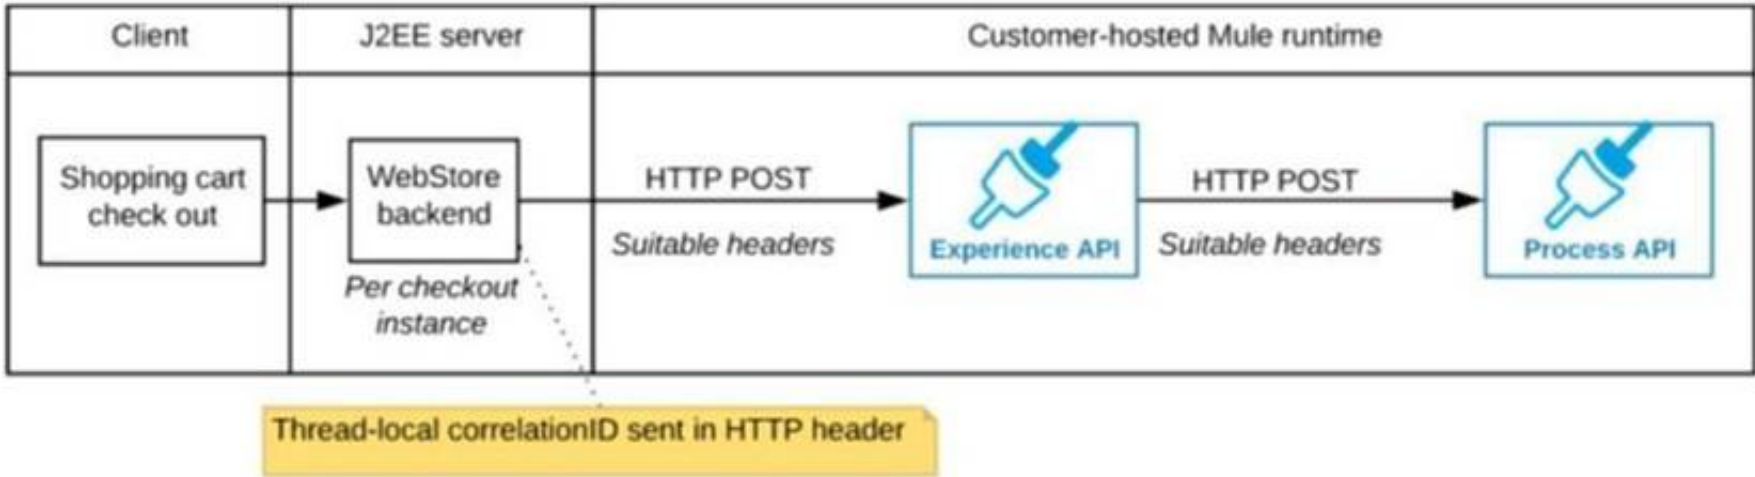
D)

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
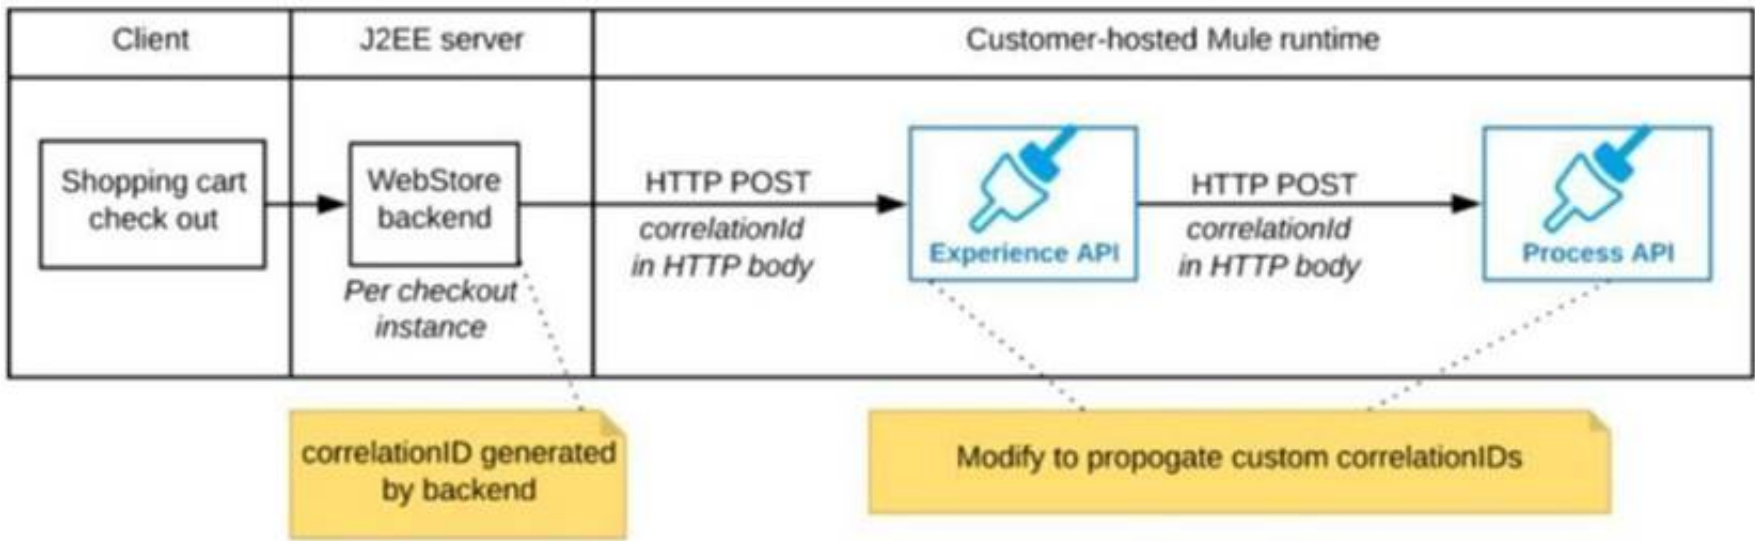

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:
: By design, Correlation Ids cannot be changed within a flow in Mule 4 applications and can be set only at source. This ID is part of the Event Context and is generated as soon as the message is received by the application. When a HTTP Request is received, the request is inspected for "X-Correlation-Id" header. If "X-Correlation-Id" header is present, HTTP connector uses this as the Correlation Id. If "X-Correlation-Id" header is NOT present, a Correlation Id is randomly generated. For
Incoming HTTP Requests: In order to set a custom Correlation Id, the client invoking the HTTP request must set "X-Correlation-Id" header. This will ensure that the Mule Flow uses this Correlation Id. For Outgoing HTTP Requests: You can also propagate the existing Correlation Id to downstream APIs. By default, all outgoing HTTP Requests send "X-Correlation-Id" header. However, you can choose to set a different value to "X-Correlation-Id" header or set "Send Correlation Id" to NEVER.
Mulesoft Reference:
https://help.mulesoft.com/s/article/How-to-Set-Custom-Correlation-Id-for-Flows-with-HTTP-Endpoint-in-Mule
Graphical user interface, application, Word Description automatically generated
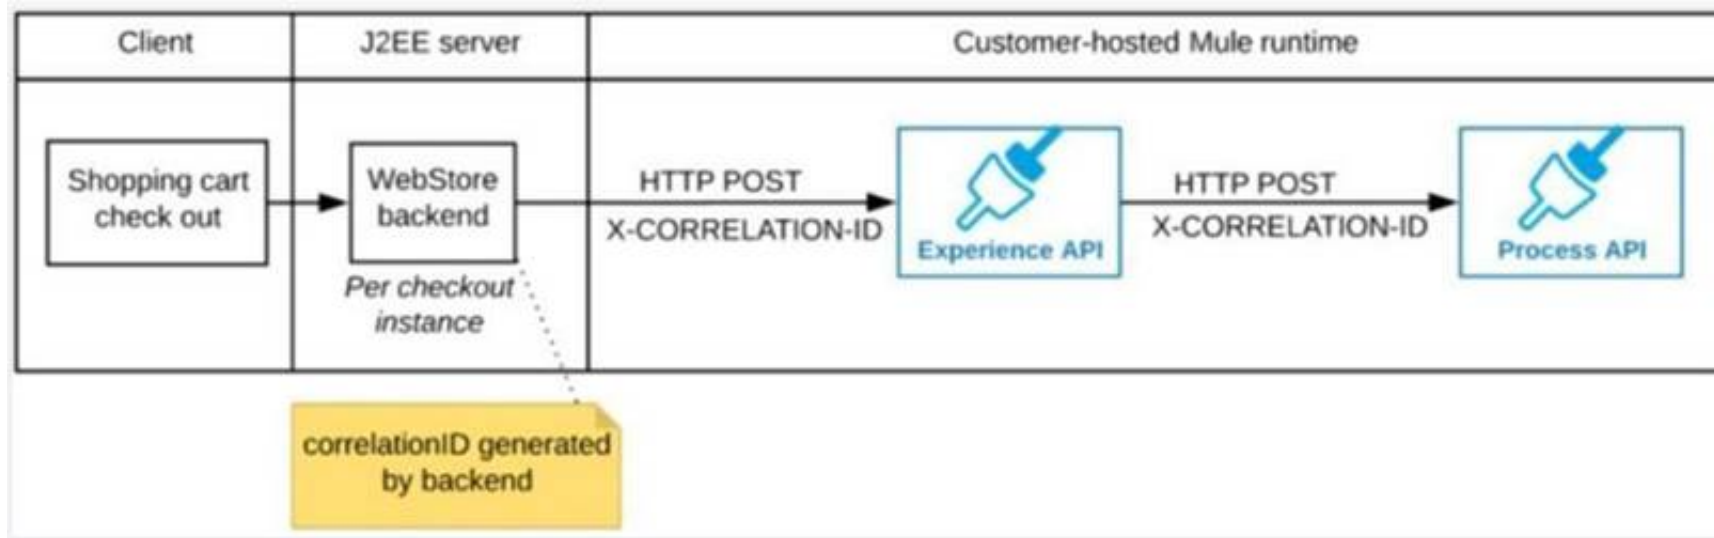
NEW QUESTION 18
Customer has deployed mule applications to different customer hosted mule run times. Mule applications are managed from Anypoint platform.
What needs to be configured to monitor these Mule applications from Anypoint monitoring and what sends monitoring data to Anypoint monitoring?

A. Enable monitoring of individual applications from runtime manager application settingsRuntime manager agent sends monitoring data from the mule applications to Anypoint monitoring
B. Install runtime manager agent on each mule runtimeRuntime manager agent since monitoring data from the mule applications to Anypoint monitoring
C. Anypoint monitoring agent on each mule runtimeAnypoint monitoring agent sends monitoring data from the mule applications to Anypoint monitoring
D. By default, Anypoint monitoring agent will be installed on each Mule run timeAnypoint Monitoring agent automatically sends monitoring data from the Mule applications to Anypoint monitoring

Answer: C

NEW QUESTION 22
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

Explanation:
To use a dedicated load balancer in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.

NEW QUESTION 24
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

Explanation:
* Roles are business group specific. Configure identity management in the Anypoint Platform master organization. As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). * Roles and permissions can be set up at business group and organization level also. But Identity Management setup is only done at Organization level * Business groups are self-contained resource groups that contain Anypoint Platform resources such as applications and APIs. Business groups provide a way to separate and control access to Anypoint Platform resources because users have access only to the busine

NEW QUESTION 25
An Order microservice and a Fulfillment microservice are being designed to communicate with their dients 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 Mute application, the Fulfillment microservice.
The Fulfilment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilted message (a kind of event message). Each OrderFulfilted message can be consumed by any interested Mule application, and the Order microservice is one such Mute 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 queu
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 queu

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

Explanation:

* If you need to scale a JMS provider/ message broker, - add nodes to scale it horizontally or - add memory to scale it vertically * Cons of adding another JMS provider/ message broker: - adds cost. - adds complexity to use two JMS brokers - adds Operational overhead if we use two brokers, say, ActiveMQ and IBM MQ * So Two options that mention to use two brokers are not best choice. * It's mentioned that "The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message. Each OrderFulfilled message can be consumed by any interested Mule application." - When you publish a message on a topic, it goes to all the subscribers who are interested - so zero to many subscribers will receive a copy of the message. - When you send a message on a queue, it will be received by exactly one consumer. * As we need multiple consumers to consume the message below option is not valid choice: "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" * Order messages are only consumed by one Mule application, the Fulfillment microservice, so we will publish it on queue and OrderFulfilled message can be consumed by any interested Mule application so it need to be published on Topic using same broker. * Correct Answer Best choice in this scenario is: "Order messages are sent to a JMS queue. 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" Tried to depict scenario in diagram:

Diagram Description automatically generated
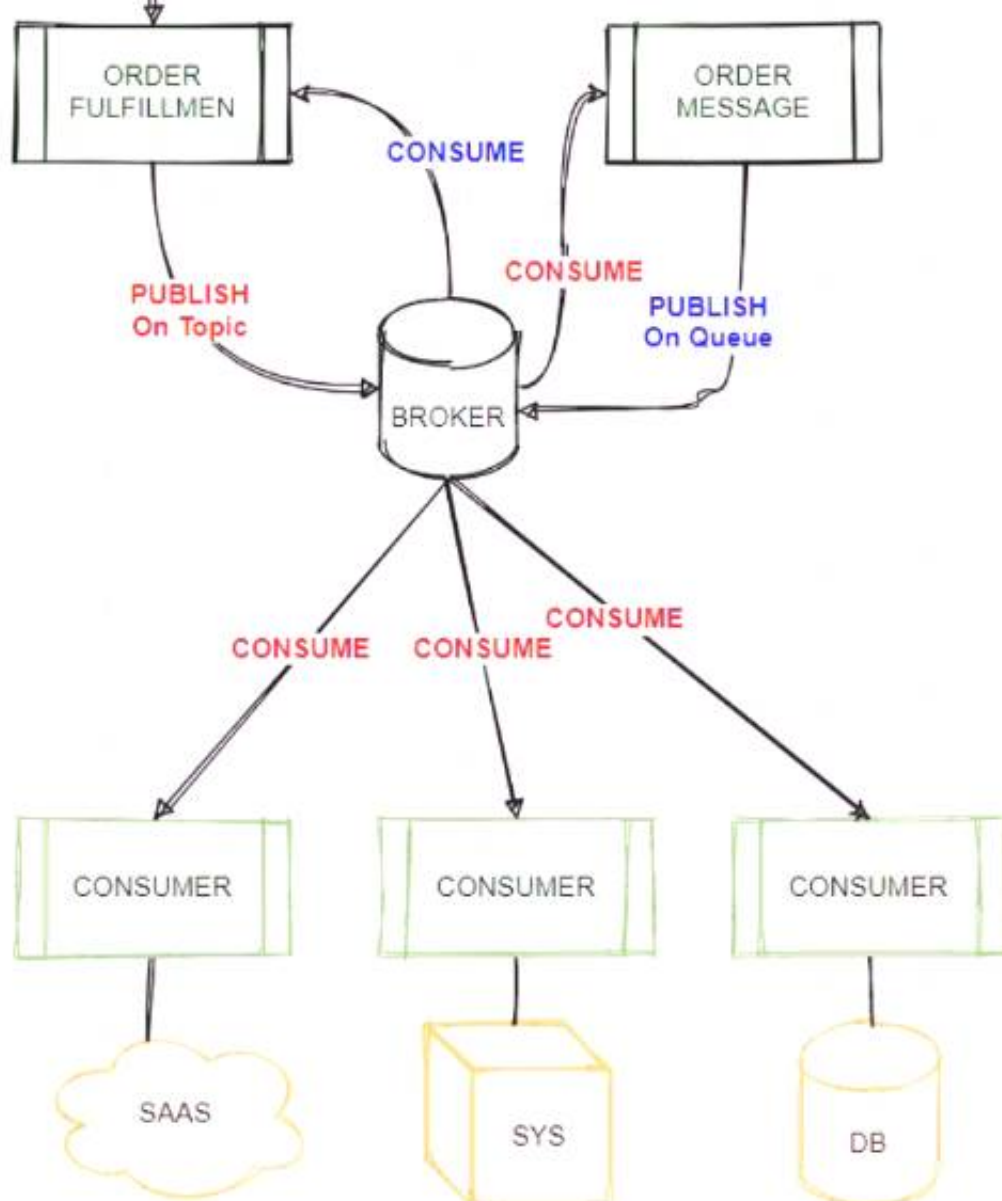
NEW QUESTION 29

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 HTTPSAMQP 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 33

Which of the below requirements prevent the usage of Anypoint MQ in a company's network? (Choose two answers)

A. single message payload can be up to 15 MB

B. payloads must be encrypted

C. the message broker must be hosted on premises

D. support for point-to-point messaging

E. ability for a third party outside the company's network to consume events from the queue

Answer: CD

NEW QUESTION 35
An organization has deployed runtime fabric on an eight note cluster with performance profile. An API uses and non persistent object store for maintaining some of its state data. What will be the impact to the stale 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 more is unaffected by the crash

Answer: D

NEW QUESTION 37
How does timeout attribute help inform design decisions while using JMS connector listening for incoming messages in an extended architecture (XA) transaction?

A. After the timeout is exceeded, stale JMS consumer threads are destroyed and new threads are created
B. The timeout specifies the time allowed to pass between receiving JMS messages on the same JMS connection and then after the timeout new JMS connection is established
C. The time allowed to pass between committing the transaction and the completion of the mule flow and then after the timeout flow processing triggers an error
D. The timeout defines the time that is allowed to pass without the transaction ending explicitly and after the timeout expires, the transaction rolls back

Answer: D

NEW QUESTION 40
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 44
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 46
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 50
Refer to the exhibit.
An organization is designing a Mule application to receive data from one external business partner. The two companies currently have no shared IT infrastructure and do not want to establish one. Instead, all communication should be over the public internet (with no VPN).
What Anypoint Connector can be used in the organization's Mule application to securely receive data from this external business partner?

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

Answer: C

Explanation:
* Object Store and VM Store is used for sharing data inter or intra mule applications in same setup. Can't be used with external Business Partner
* Also File connector will not be useful as the two companies currently have no shared IT infrastructure. It's specific for local use.
* Correct answer is SFTP connector. The SFTP Connector implements a secure file transport channel so that your Mule application can exchange files with external resources. SFTP uses the SSH security protocol to transfer messages. You can implement the SFTP endpoint as an inbound endpoint with a one-way exchange pattern, or as an outbound endpoint configured for either a one-way or request-response exchange pattern.

NEW QUESTION 52
A manufacturing company is planning to deploy Mule applications to its own Azure Kubernetes Service infrastructure.
The organization wants to make the Mule applications more available and robust by deploying each Mule application to an isolated Mule runtime in a Docker container while managing all the Mule applications from the MuleSoft-hosted control plane.
What is the most idiomatic (used for its intended purpose) choice of runtime plane to meet these organizational requirements?

A. Anypoint Platform Private Cloud Edition
B. Anypoint Runtime Fabric
C. CloudHub
D. Anypoint Service Mesh

Answer: B

NEW QUESTION 57
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 system.Then conduct performance test to identify other bottlenecks in the system
C. Create a Mocking service that replicates the back-end system's production performance characteristicsThen configure the API implementation to use the mocking service and conduct the performance test
D. Conduct scaled-down performance tests in the staging environment against rate-limiting back-end syste
E. Then upscale performance results to full production scale

Answer: C

NEW QUESTION 62
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 availabiltty 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 67
An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events

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

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

Answer: AD

NEW QUESTION 70

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 75

A team would like to create a project skeleton that developers can use as a starting point when creating API Implementations with Anypoint Studio. This skeleton should help drive consistent use of best practices within the team.
What type of Anypoint Exchange artifact(s) should be added to Anypoint Exchange to publish the project skeleton?

A. A custom asset with the default API implementation
B. A RAML archetype and reusable trait definitions to be reused across API implementations
C. An example of an API implementation following best practices
D. a Mule application template with the key components and minimal integration logic

Answer: D

Explanation:
* Sharing Mule applications as templates is a great way to share your work with other people who are in your organization in Anypoint Platform. When they need to build a similar application they can create the mule application using the template project from Anypoint studio.
* Anypoint Templates are designed to make it easier and faster to go from a blank canvas to a production application. They're bit for bit Mule applications requiring only Anypoint Studio to build and design, and are deployable both on-premises and in the cloud.
* Anypoint Templates are based on five common data Integration patterns and can be customized and extended to fit your integration needs. So even if your use case involves different endpoints or connectors than those included in the template, they still offer a great starting point.
Some of the best practices while creating the template project: - Define the common error handler as part of template project, either using pom dependency or mule config file - Define common logger/audit framework as part of the template project - Define the env specific properties and secure properties file as per the requirement - Define global.xml for global configuration - Define the config file for connector configuration like Http,Salesforce,File,FTP etc - Create separate folders to create DWL,Properties,SSL certificates etc - Add the dependency and configure the pom.xml as per the business need - Configure the mule-artifact.json as per the business need

NEW QUESTION 80

Anypoint Exchange is required to maintain the source code of some of the assets committed to it, such as Connectors, Templates, and API specifications.
What is the best way to use an organization's source-code management (SCM) system in this context?

A. Organizations should continue to use an SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching, and merging
B. Organizations need to use Anypoint Exchange as the main SCM system to centralize versioning and avoid code duplication
C. Organizations can continue to use an SCM system of their choice for branching and merging, as long as they follow the branching and merging strategy enforced by Anypoint Exchange
D. Organizations need to point Anypoint Exchange to their SCM system so Anypoint Exchange can pull source code when requested by developers and provide it to Anypoint Studio

Answer: B

Explanation:
* Organization should continue to use SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching.
* Reason is that Anypoint exchange is not full fledged version repositories like GitHub.
* But at same time it is tightly coupled with Mule assets

NEW QUESTION 82

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 communncation
D. Ensure consistency

Answer: D

Explanation:
* XA transaction add tremendous latency so "Reduce Latency" is incorrect option XA transactions define "All or No" commit protocol.
* Each local XA resource manager supports the A.C.I.D properties (Atomicity, Consistency, Isolation, and Durability).

So correct choice is "Ensure consistency"

NEW QUESTION 83
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 85
Refer to the exhibit.
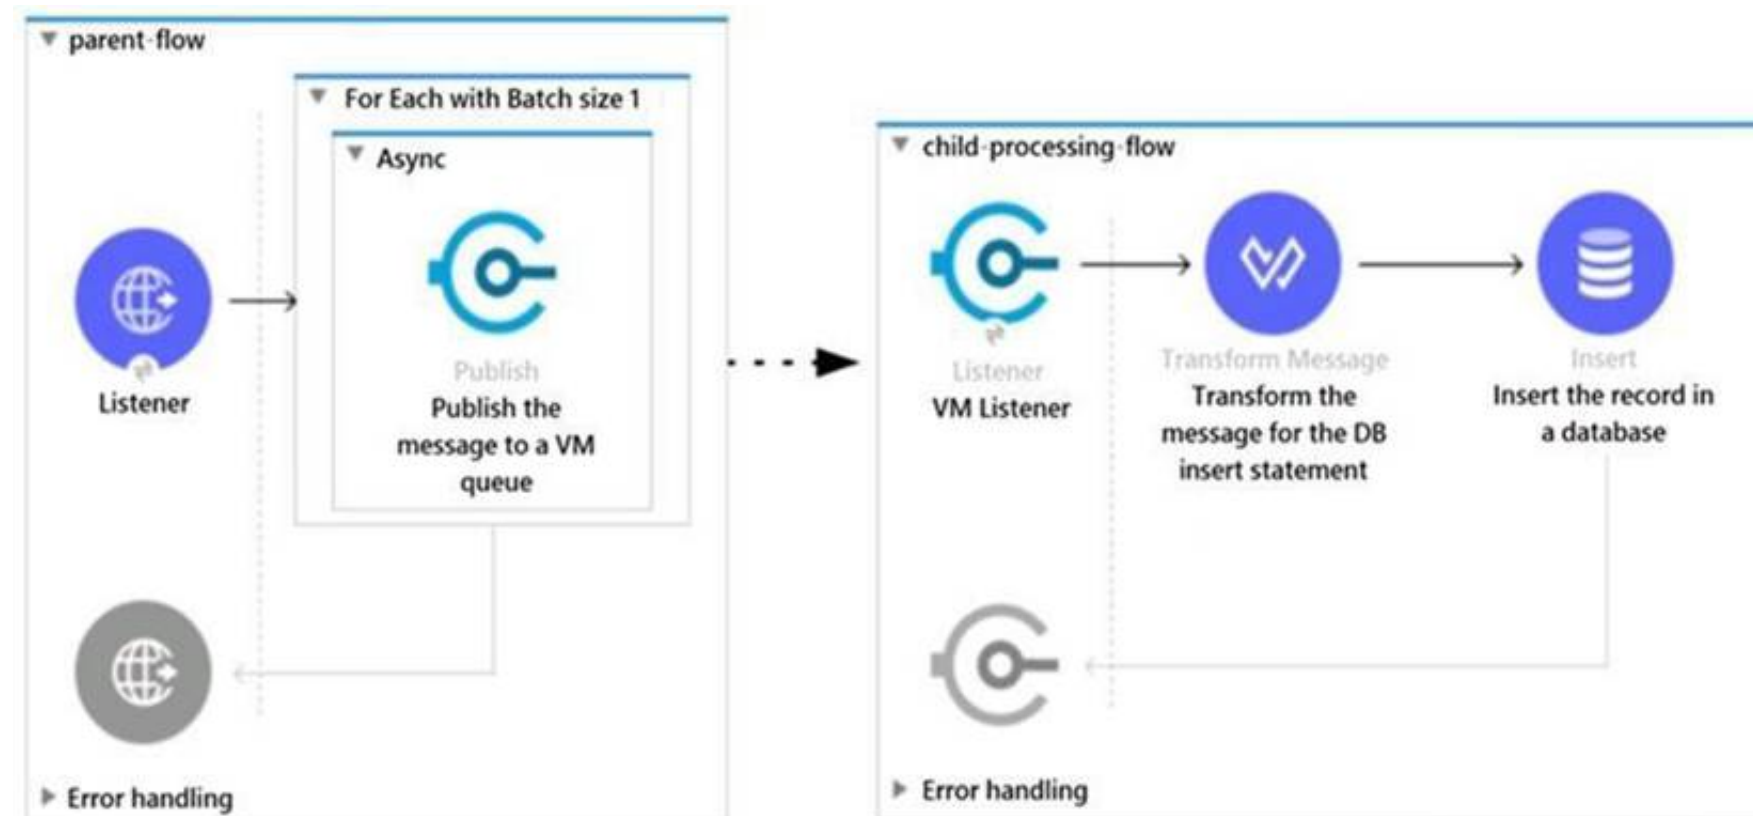
A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items, then sends each item one at a time inside an Async scope to a VM queue.
A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing 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 invokedThis 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

Explanation:
Correct answer is 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 In Cloudhub, each persistent VM queue is listened on by every CloudHub worker - But each message is read and processed at least once by only one CloudHub worker and the duplicate processing is possible - If the CloudHub worker fails , the message can be read by another worker to prevent loss of messages and this can lead to duplicate processing - By default , every CloudHub worker's VM Listener receives different messages from VM Queue Referenece: https://dzone.com/articles/deploying-mulesoft-application-on-1-worker-vs-mult

NEW QUESTION 89
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

Explanation:
Context of the question is about managing and governing mule applications deployed on Anypoint platform.
Anypoint API Manager (API Manager) is a component of Anypoint Platform that enables you to manage, govern, and secure APIs. It leverages the runtime capabilities of API Gateway and Anypoint Service Mesh, both of which enforce policies, collect and track analytics data, manage proxies, provide encryption and authentication, and manage applications.
Mule Ref Doc : https://docs.mulesoft.com/api-manager/2.x/getting-started-proxy

NEW QUESTION 90
An organization has strict unit test requirement that mandate every mule application must have an MUnit test suit with a test case defined for each flow and a minimum test coverage of 80%.
A developer is building Munit test suit for a newly developed mule application that sends API request to an
external rest API.
What is the effective approach for successfully executing the Munit tests of this new application while still achieving the required test coverage for the Munit tests?

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

Answer: C

NEW QUESTION 92
An integration Mule application is deployed to a customer-hosted multi-node Mule 4 runtime duster. The Mule application uses a Listener operation of a JMS connector to receive incoming messages from a JMS
queue.
How are the messages consumed by the Mule application?

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

Answer: C

Explanation:
Correct answer is 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
For applications running in clusters, you have to keep in mind the concept of primary node and how the connector will behave. When running in a cluster, the JMS listener default behavior will be to receive messages only in the primary node, no matter what kind of destination you are consuming from. In case of consuming messages from a Queue, you'll want to change this configuration to receive messages in all the nodes of the cluster, not just the primary.
This can be done with the primaryNodeOnly parameter:
<jms:listener config-ref="config" destination="${inputQueue}" primaryNodeOnly="false"/>

NEW QUESTION 95
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

Explanation:
The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation
This is based on below definitions API client • An application component • that accesses a service • by invoking an API of that service - by definition of the term API over HTTP API consumer • A business role, which is often assigned to an individual • that develops API clients, i.e., performs the activities necessary for enabling an API client to invoke APIs API implementation • An application component • that implements th functionality

NEW QUESTION 99
......

Relate Links

100% Pass Your MCIA-Level-1 Exam with Exambible Prep Materials

https://www.exambible.com/MCIA-Level-1-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/