

MCIA-Level-1 Dumps

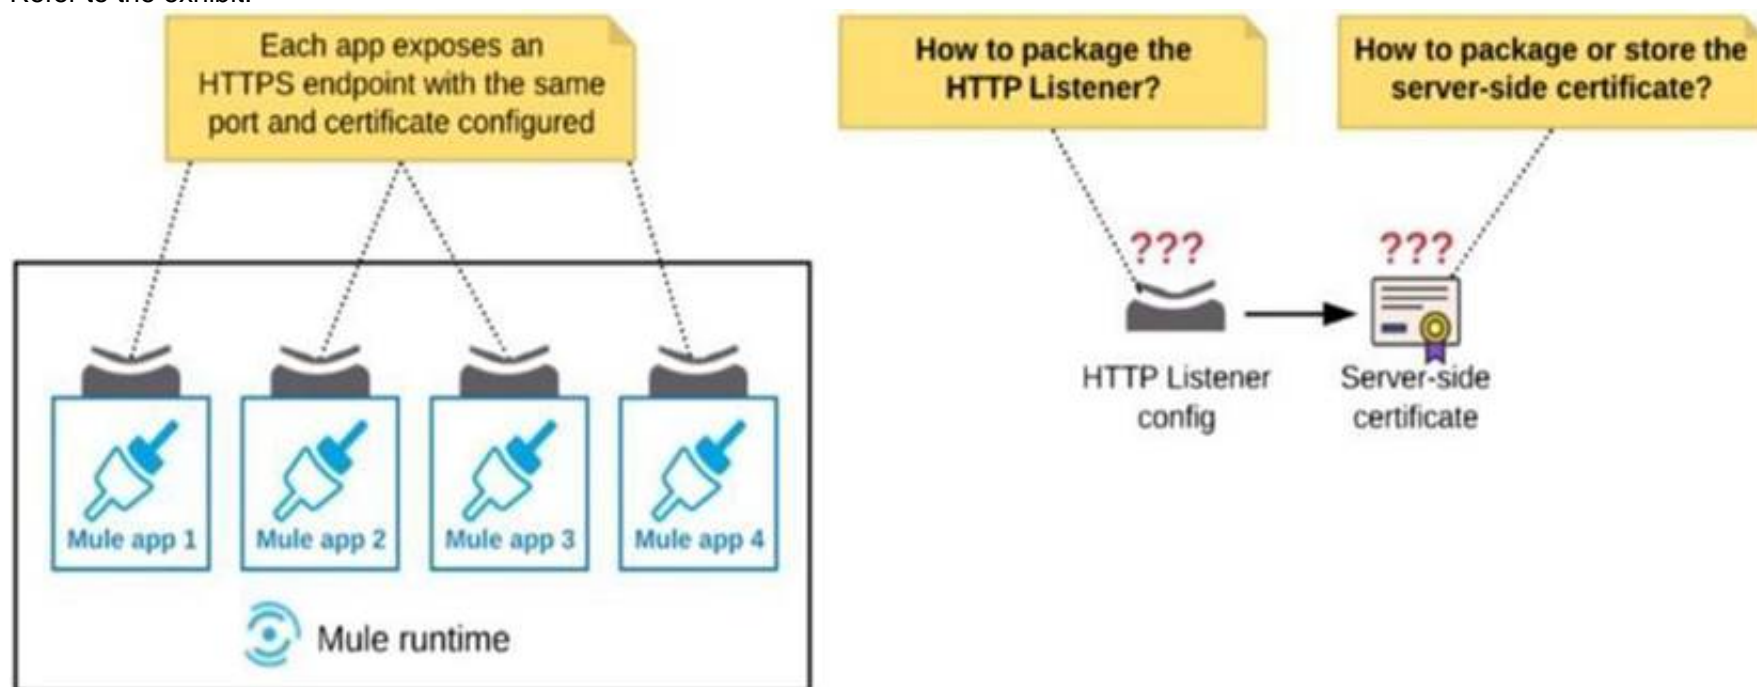
MuleSoft Certified Integration Architect - Level 1

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



NEW QUESTION 1

Refer to the exhibit.



An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a serverside certificate that rotates often.

What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

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

Answer: B

NEW QUESTION 2

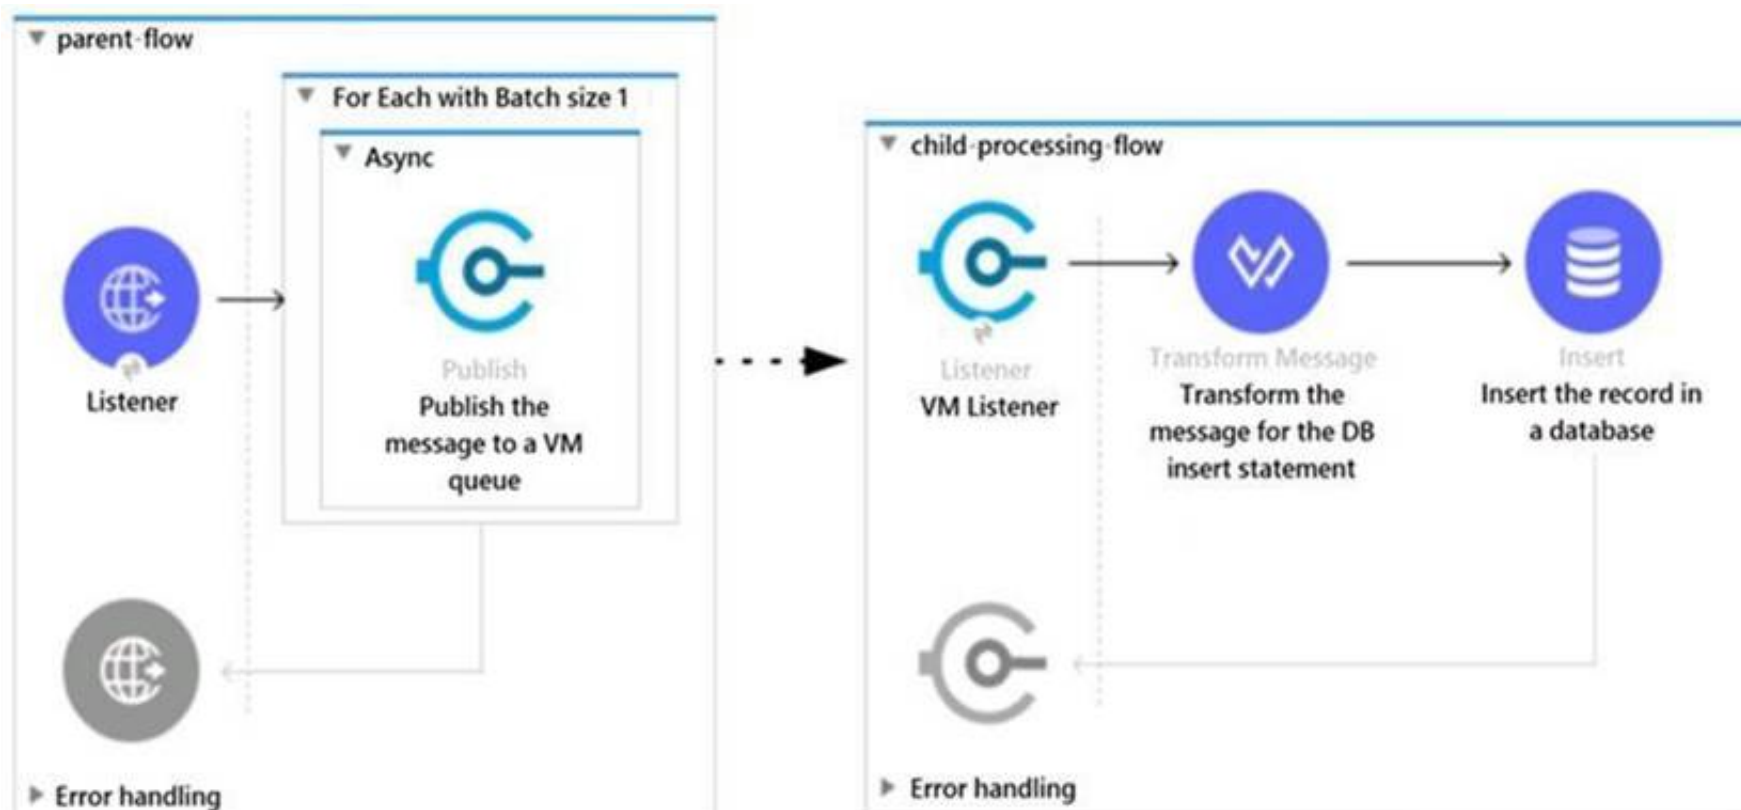
An XA transaction Is being configured that involves a JMS connector listening for Incoming JMS messages. What is the meaning of the timeout attribute of the XA transaction, and what happens after the timeout expires?

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

Answer: C

NEW QUESTION 3

Refer to the exhibit.



A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items,

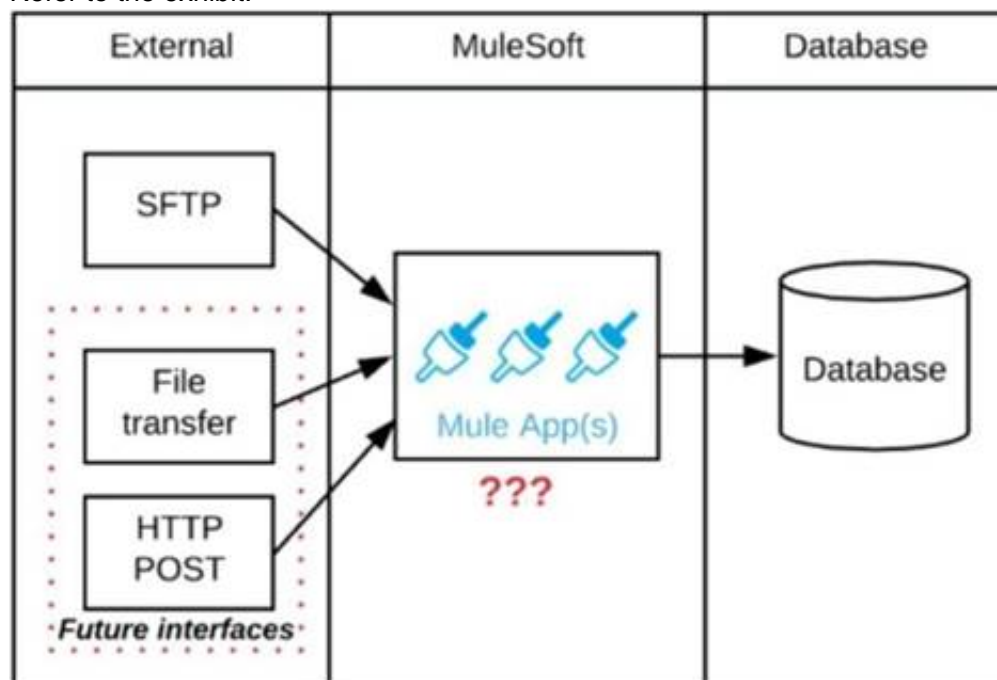
then sends each item one at a time inside an Async scope to a VM queue.
A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing the item to a database.
This Mule application is deployed to four CloudHub workers with persistent 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 4

Refer to the exhibit.



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

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

- A. Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- B. Create a Process API to receive the file and process it using a MuleSoft Batch Job while delegating the data save process to a System API
- C. Create an API that receives the file and invokes a Process API with the data contained In the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed
- D. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing

Answer: C

NEW QUESTION 5

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

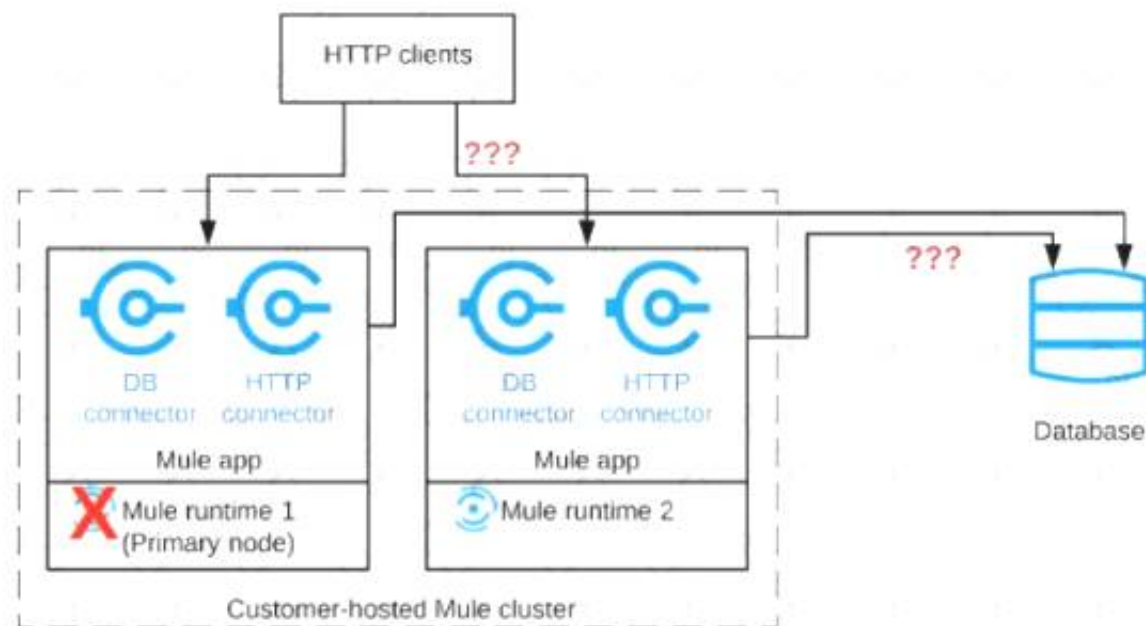
What key elements should be included in the integration design specification when requesting approval for this Mule application?

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

Answer: A

NEW QUESTION 6

Refer to the exhibit.



A Mule application is deployed to a cluster of two customer-hosted Mule runtimes. The Mule application has a flow that polls a database and another flow with an HTTP Listener.

HTTP clients send HTTP requests directly to individual cluster nodes.

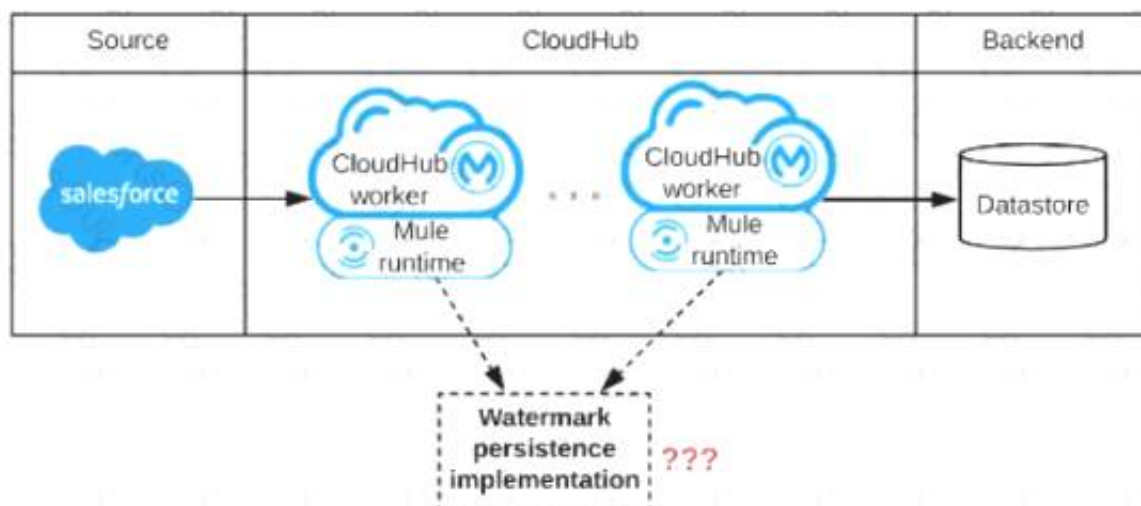
What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has failed, but before that node is restarted?

- A. Database polling continues Only HTTP requests sent to the remaining node continue to be accepted
- B. Database polling stops All HTTP requests continue to be accepted
- C. Database polling continues All HTTP requests continue to be accepted, but requests to the failed node incur increased latency
- D. Database polling stops All HTTP requests are rejected

Answer: A

NEW QUESTION 7

Refer to the exhibit.



A Mule application is being designed to be deployed to several CloudHub workers. The Mule application's integration logic is to replicate changed Accounts from Salesforce to a backend system every 5 minutes.

A watermark will be used to only retrieve those Salesforce Accounts that have been modified since the last time the integration logic ran.

What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?

- A. Persistent Anypoint MQ Queue
- B. Persistent Object Store
- C. Persistent Cache Scope
- D. Persistent VM Queue

Answer: B

NEW QUESTION 8

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

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

Answer: B

NEW QUESTION 9

An organization currently uses a multi-node Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications. The organization is planning to transition to a deployment model based on Docker containers in a Kubernetes cluster. The organization has already created a standard Docker image containing a Mule runtime and all required dependencies (including a JVM), but excluding the Mule application itself.

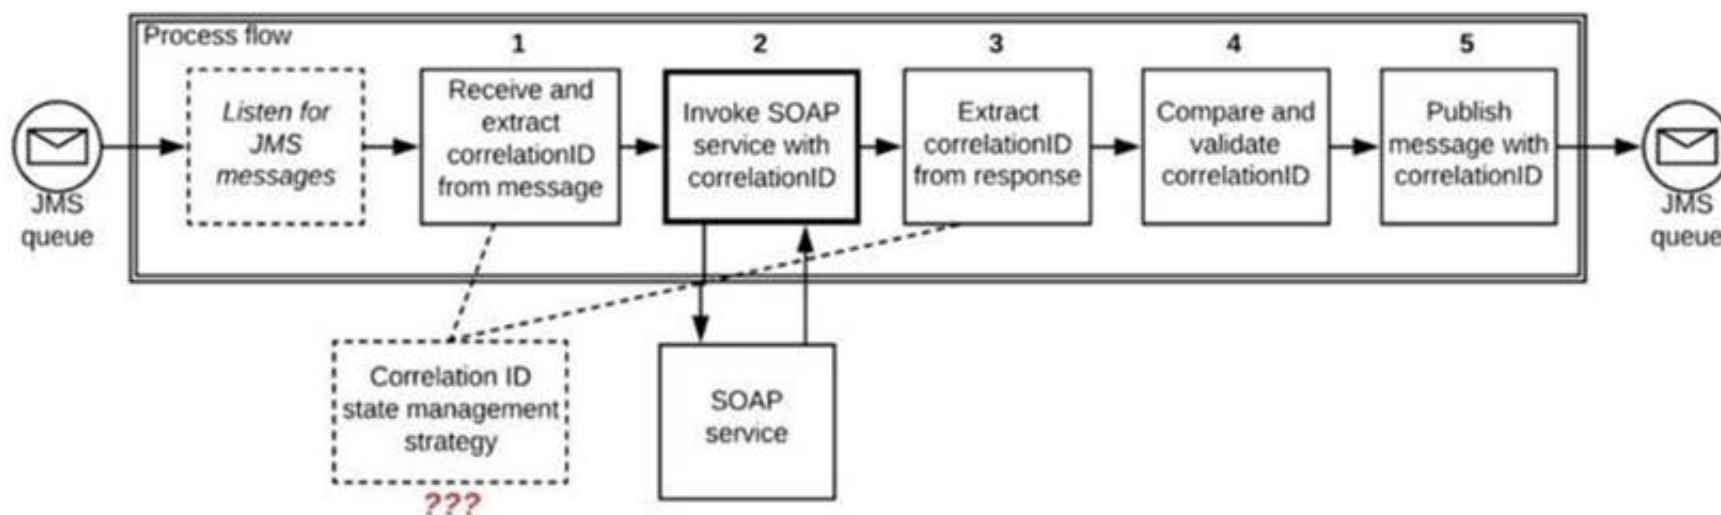
What is an expected outcome of this transition to container-based Mule application deployments?

- A. Required redesign of Mule applications to follow microservice architecture principles
- B. Required migration to the Docker and Kubernetes-based Anypoint Platform - Private Cloud Edition
- C. Required change to the URL endpoints used by clients to send requests to the Mule applications
- D. Guaranteed consistency of execution environments across all deployments of a Mule application

Answer: A

NEW QUESTION 10

Refer to the exhibit.



A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to receive JMS messages from a JMS queue. To process each received JMS message, the following steps are performed in a flow:

Step 1: The JMS Correlation ID header is read from the received JMS message.

Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.

Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.

Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.

Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.

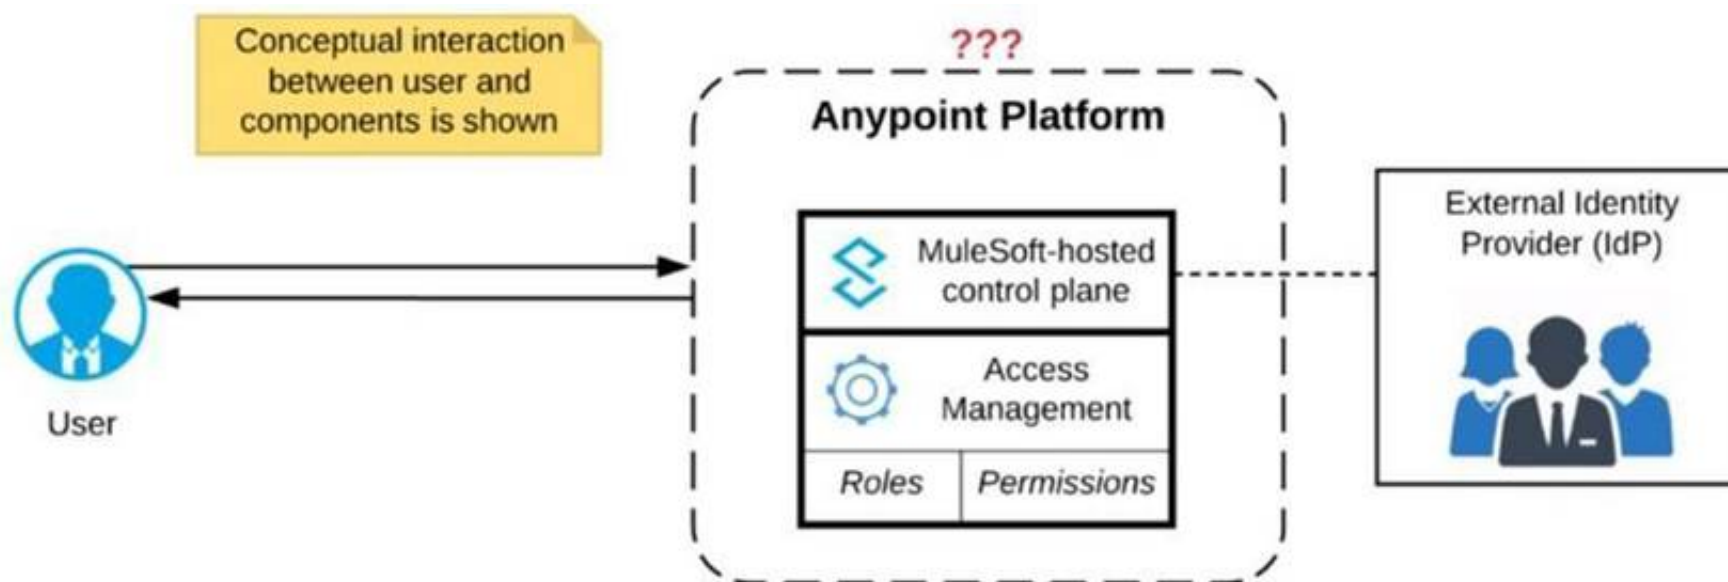
Where should the Mule application store the JMS Correlation ID values received in Step 1 and Step 3 so that the validation in Step 4 can be performed, while also making the overall Mule application highly available, fault-tolerant, performant, and maintainable?

- A. Both Correlation ID values should be stored in a persistent object store
- B. Both Correlation ID values should be stored in a non-persistent object store
- C. The Correlation ID value in Step 1 should be stored in a persistent object storeThe Correlation ID value in step 3 should be stored as a Mule event variable/attribute
- D. Both Correlation ID values should be stored as Mule event variables/attributes

Answer: C

NEW QUESTION 10

Refer to the exhibit.



Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.

What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled via the external Identity Provider?

- A. Controlling the business group within Anypoint Platform to which the user belongs
- B. Assigning Anypoint Platform permissions to a role
- C. Assigning Anypoint Platform role(s) to a user
- D. Removing a user's access to Anypoint Platform when they no longer work for the organization

Answer: B

NEW QUESTION 11

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

- A. The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider
- B. The AMQP protocol can be used by the JMS connector to portably establish connections to various types of JMS providers

- C. To receive messages into the Mule application, the JMS provider initiates a network connection to the JMS connector and pushes messages along this connection
- D. To complete sending a JMS message, the JMS connector must establish a network connection with the JMS message recipient

Answer: D

NEW QUESTION 15

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 RAML archetype and reusable trait definitions to be reused across API implementations
- B. A custom asset with the default API implementation
- 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

NEW QUESTION 16

What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

- A. Both a DLB and the SLB implement load balancing by sending HTTP requests to workers with the lowest workloads
- B. Both a DLB and the SLB allow the configuration of access control via IP whitelists
- C. Only a DLB allows the configuration of a custom TLS server certificate
- D. Only the SLB can forward HTTP traffic to the VPC-internal ports of the CloudHub workers

Answer: B

NEW QUESTION 17

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

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

Answer: D

NEW QUESTION 20

What Anypoint Connectors support transactions?

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

Answer: A

NEW QUESTION 23

An integration Mule application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mule application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system.

Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system).

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

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

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

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

Answer: A

NEW QUESTION 25

A retailer is designing a data exchange interface to be used by its suppliers. The interface must support secure communication over the public internet. The interface must also work with a wide variety of programming languages and IT systems used by suppliers.

What are suitable interface technologies for this data exchange that are secure, cross-platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

- A. EDJFACT XML over SFTP JSON/REST over HTTPS
- B. SOAP over HTTPS HOP over TLS gRPC over HTTPS
- C. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
- D. CSV over FTP YAML over TLS JSON over HTTPS

Answer: B

NEW QUESTION 29

An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to customhosted Mule runtimes (on-premises) while others execute in theMuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems.

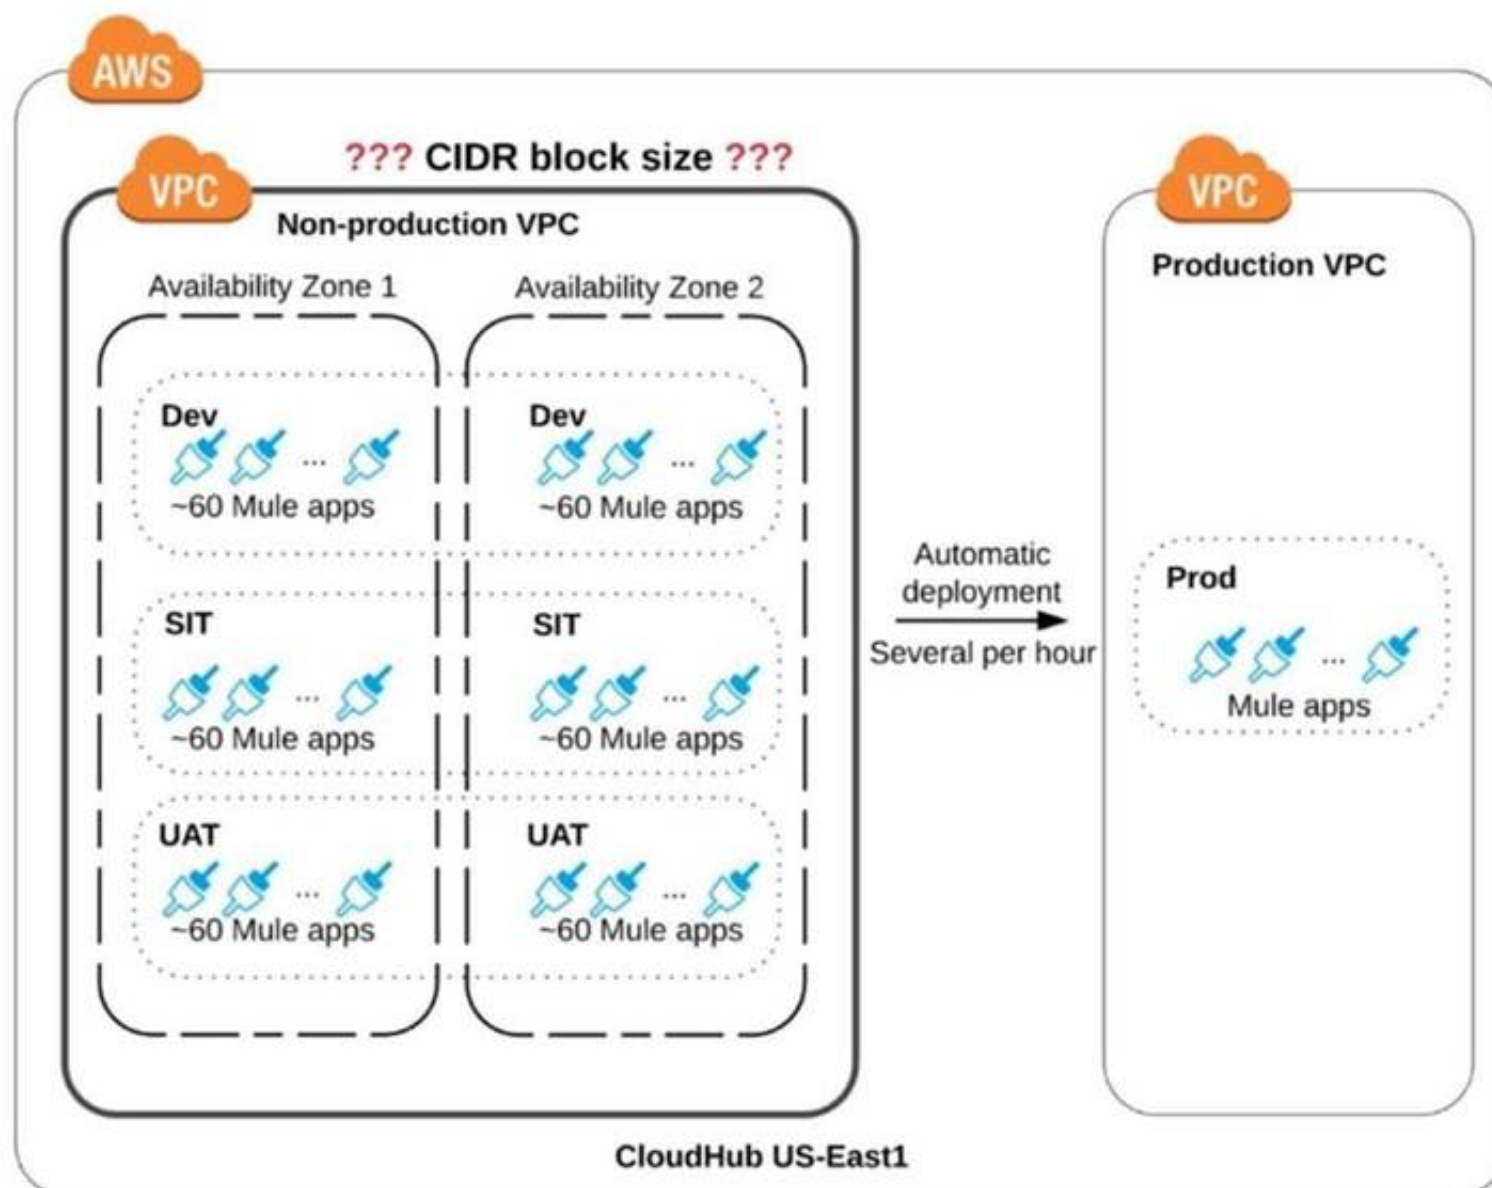
How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access thebackend systems?

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

Answer: D

NEW QUESTION 34

Refer to the exhibit.



An organization is sizing an Anypoint VPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization currently has three non-production environments (DEV, SIT and UAT) that share this VPC. The AWS region of the VPC has two AZs.

The organization has a very mature DevOps approach which automatically progresses each application through all non-production environments before automatically deploying to production. This process results in several Mule application deployments per hour, using CloudHub's normal zero-downtime deployment feature.

What is a CIDR block for this VPC that results in the smallest usable private IP address range?

- A. 10.0.0.0/26 (64 IPS)
- B. 10.0.0.0/25 (128 IPs)
- C. 10.0.0.0/24 (256 IPs)
- D. 10.0.0.0/22 (1024 IPs)

Answer: D

NEW QUESTION 37

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

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

- A. Use Anypoint Design Center to implement the Mule applications and APIs and give the various stakeholders access to these Design Center projects, so they can collaborate and provide feedback
- B. Create Anypoint Exchange entries with pages elaborating the integration design, including API notebooks (where applicable) to help the stakeholders understand and interact with the Mule applications and APIs at various levels of technical depth

- C. Use Anypoint Exchange to register the various Mule applications and APIs and share the RAML definitions with the stakeholders, so they can be discovered
D. Capture documentation about the Mule applications and APIs inline within the Mule integration flows and use Anypoint Studio's Export Documentation feature to provide an HTML version of this documentation to the stakeholders

Answer: B

NEW QUESTION 39

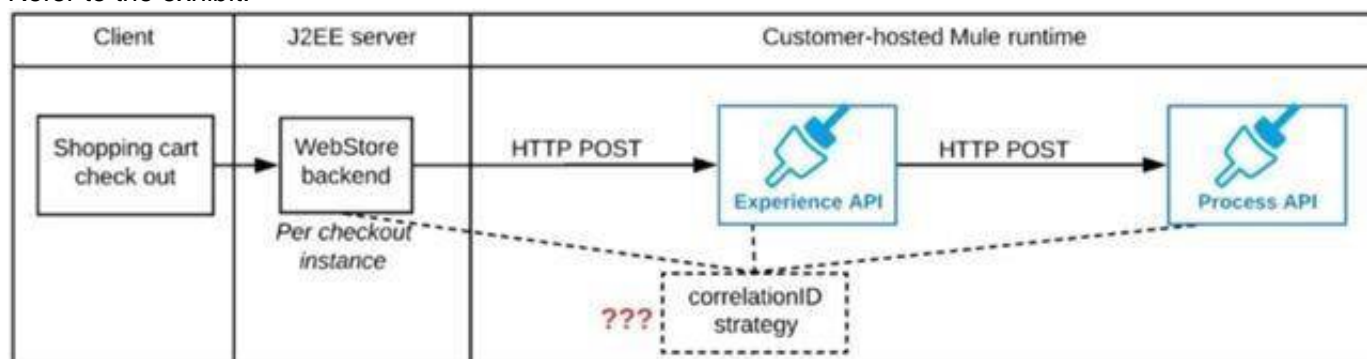
What metrics about API invocations are available for visualization in custom charts using Anypoint Analytics?

- A. Request size, request HTTP verbs, response time
B. Request size, number of requests, JDBC Select operation result set size
C. Request size, number of requests, JDBC Select operation response time
D. Request size, number of requests, response size, response time

Answer: D

NEW QUESTION 40

Refer to the exhibit.

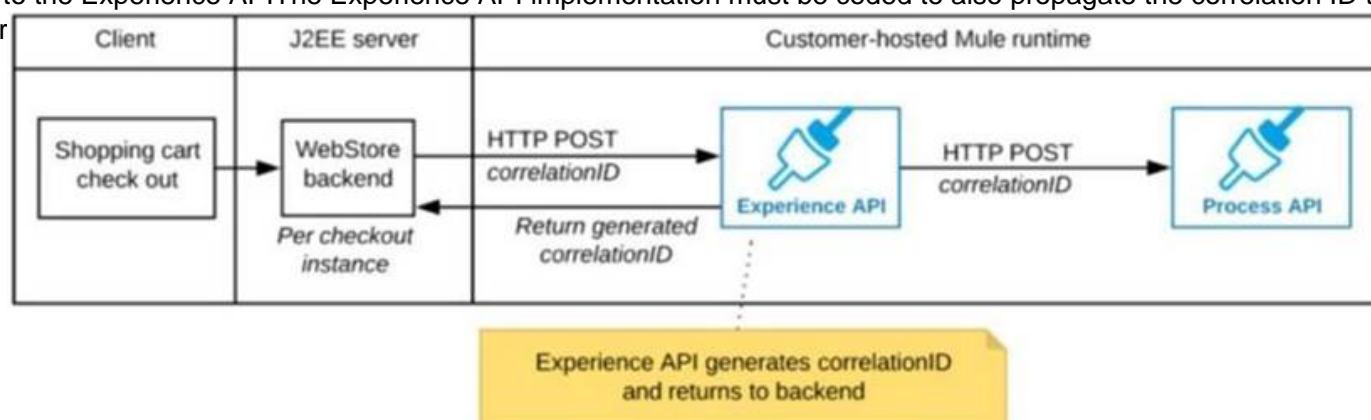


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

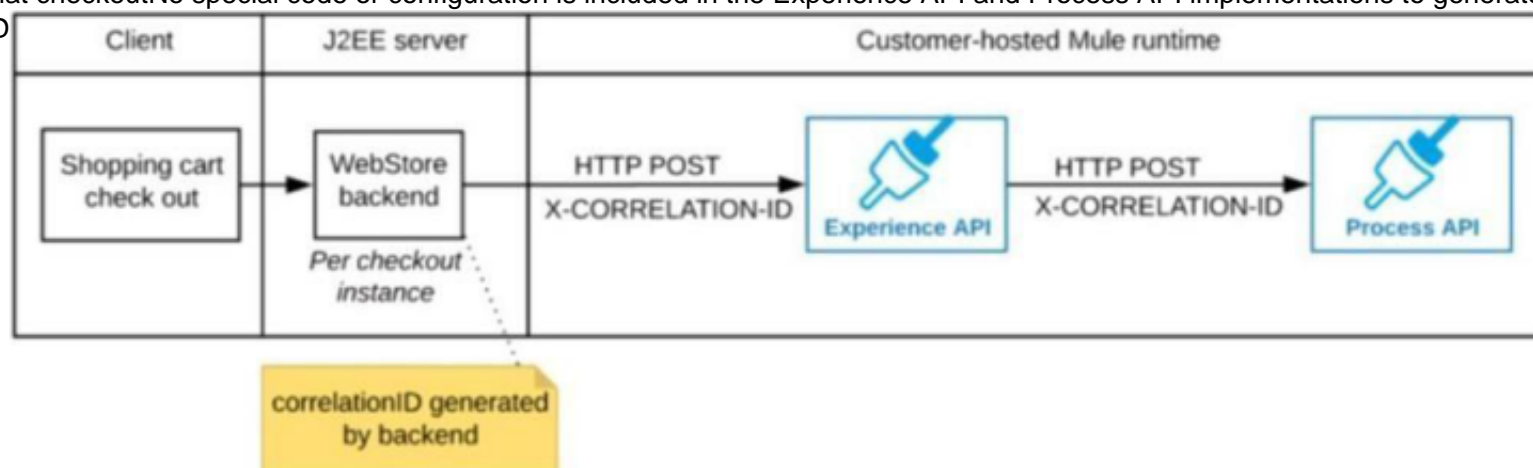
End-to-end correlation of all HTTP requests and responses belonging to each individual checkout instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the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

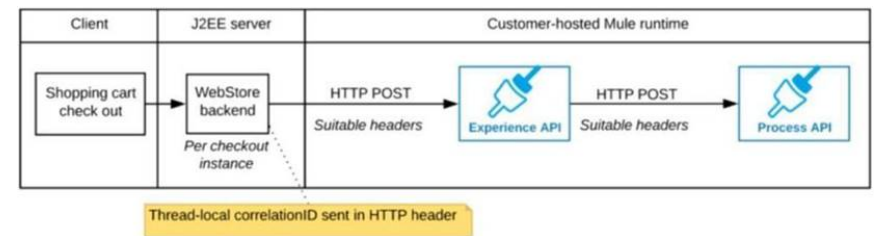
- A. The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API. The Experience API implementation must be coded to also propagate the correlation ID to the Process API in a suitable HTTP request header.



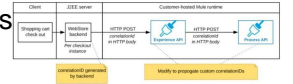
- B. The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-ID HTTP request header in each API invocation belonging to that checkout. No special code or configuration is included in the Experience API and Process API implementations to generate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



- C. The 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.



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.



Answer: B

NEW QUESTION 45

Mule application A receives a request Anyoint 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 Anyoint 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 Anyoint 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. 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
- B. Use a Scatter-Gather within the For Each scope to ensure response message order. Configure the Scatter-Gather with a persistent object store
- C. Keep track of the list length and all object indices in REQU, both in the For Each scope and in all communication involving service S
- D. Use persistent storage when creating RESP
- E. 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

Answer: B

NEW QUESTION 49

What is required before an API implemented using the components of Anyoint Platform can be managed and governed (by applying API policies) on Anyoint Platform?

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

Answer: A

NEW QUESTION 54

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 Anyoint MQ exchange. OrderFulfilled messages are sent to an Anyoint MQ queue. Both microservices interact with Anyoint 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. 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
- C. Order messages are sent directly to the Fulfillment microservice. 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
- D. Order messages are sent to a JMS queue. 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: D

NEW QUESTION 59

An organization uses a set of customer-hosted Mule runtimes that are managed using the Mulesoft-hosted control plane.

What is a condition that can be alerted on from Anyoint Runtime Manager without any custom components or custom coding?

- A. When an SSL certificate used by one of the deployed Mule applications is about to expire
- B. When a Mule runtime on a given customer-hosted server is experiencing high memory consumption during certain periods
- C. When a Mule runtime's customer-hosted server is about to run out of disk space
- D. When the Mule runtime license installed on a Mule runtime is about to expire

Answer: A

NEW QUESTION 62

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

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

Answer: A

NEW QUESTION 64

An integration Mule application consumes and processes a list of rows from a CSV file. Each row must be read from the CSV file, validated, and the row data sent to a JMS queue, in the exact order as in the CSV file.

If any processing step for a row fails, then a log entry must be written for that row, but processing of other rows must not be affected.

What combination of Mule components is most idiomatic (used according to their intended purpose) when implementing the above requirements?

- A. Scatter-Gather componentOn Error Continue scope
- B. VM connectorFirst Successful scopeOn Error Propagate scope
- C. Async scopeOn Error Propagate scope
- D. For Each scopeOn Error Continue scope

Answer: D

NEW QUESTION 65

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your MCIA-Level-1 Exam with Our Prep Materials Via below:

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