Exam Questions MuleSoft-Integration-Architect-I

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam

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

NEW QUESTION 1
An organization plans to use the Anypoint Platform audit logging service to log Anypoint MQ actions.
What consideration must be kept in mind when leveraging Anypoint MQ Audit Logs?

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

Answer: C

NEW QUESTION 2
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 rout
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 rout
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 rout
G. Use a Try-Catch scope in each route to retry requests that raise timeout errors.

Answer: C

NEW QUESTION 3
Mule applications need to be deployed to CloudHub so they can access on-premises database systems. These systems store sensitive and hence tightly protected data, so are not accessible over the internet.
What network architecture supports this requirement?

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

Answer: A

NEW QUESTION 4
An API implementation is being developed to expose data from a production database via HTTP requests. The API implementation executes a database SELECT statement that is dynamically created based upon data received from each incoming HTTP request. The developers are planning to use various types of testing to make sure the Mule application works as expected, can handle specific workloads, and behaves correctly from an API consumer perspective. What type of testing would typically mock the results from each SELECT statement rather than actually execute it in the production database?

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

Answer: A

NEW QUESTION 5
Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improved throughput of the system. But it was observed over the period of time that few of the important exception log messages which were used to rollback transactions are not working as expected causing huge loss to the Organization. Organization wants to avoid these losses. Application also has constraints due to which they cant compromise on throughput much. What is the possible option in this case?

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

Answer: D

NEW QUESTION 6
An insurance provider is implementing Anypoint platform to manage its application infrastructure and is using the customer hosted runtime for its business due to certain financial requirements it must meet. It has built a number of synchronous API's and is currently hosting these on a mule runtime on one server
These applications make use of a number of components including heavy use of object stores and VM queues.
Business has grown rapidly in the last year and the insurance provider is starting to receive reports of reliability issues from its applications.
The DevOps team indicates that the API's are currently handling too many requests and this is over loading the server. The team has also mentioned that there is a significant downtime when the server is down for maintenance.
As an integration architect, which option would you suggest to mitigate these issues?

A. Add a load balancer and add additional servers in a server group configuration
B. Add a load balancer and add additional servers in a cluster configuration
C. Increase physical specifications of server CPU memory and network
D. Change applications by use an event-driven model

Answer: B

NEW QUESTION 7
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 8
A global, high-volume shopping Mule application is being built and will be deployed to CloudHub. To improve performance, the Mule application uses a Cache scope that maintains cache state in a CloudHub object store. Web clients will access the Mule application over HTTP from all around the world, with peak volume coinciding with business hours in the web client's geographic location. To achieve optimal performance, what Anypoint Platform region should be chosen for the CloudHub object store?

A. Choose the same region as to where the Mule application is deployed
B. Choose the US-West region, the only supported region for CloudHub object stores
C. Choose the geographically closest available region for each web client
D. Choose a region that is the traffic-weighted geographic center of all web clients

Answer: A

NEW QUESTION 9
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: B

NEW QUESTION 10
A Mule application name Pub uses a persistence object store. The Pub Mule application is deployed to Cloudhub and it configured to use Object Store v2.
Another Mule application name sub is being developed to retrieve values from the Pub Mule application persistence object Store and will also be deployed to cloudhub.
What is the most direct way for the Sub Mule application to retrieve values from the Pub Mule application persistence object store with the least latency?

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

Answer: D

NEW QUESTION 10
An organization uses Mule runtimes which are managed by Anypoint Platform - Private Cloud Edition. What MuleSoft component is responsible for feeding analytics data to non- MuleSoft analytics platforms?

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

Answer: D

NEW QUESTION 11
A high-volume eCommerce retailer receives thousands of orders per hour and requires notification of its order management, warehouse, and billing system for subsequent processing within 15 minutes of order submission through its website.
Which integration technology, when used for its typical and intended purpose, meets the retailer??s requirements for this use case?

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

Answer: B

NEW QUESTION 16
A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should be published in a way that is guaranteed in normal situations and also minimizes duplicate delivery in less frequent failure scenarios.
The organizational firewall is configured to only allow outbound traffic on ports 80 and 443. Some external event consumers are within the organizational network, while others are located outside the firewall.
What Anypoint Platform service is most idiomatic (used for its intended purpose) for publishing these Mule events to all external consumers while addressing the desired reliability goals?

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

Answer: B

NEW QUESTION 21
A rale limiting policy has been applied to a soap VI.2 API published in Clondhub. The API implementation catches errors in a global error handler on error propagate in the main flow for HTTP: RETRY_EXHAUSTED with HTTP status set to 429 and any with the HTTP status set to 500.
What is the expected H1TP status when the client exceeds the quota of the API calls?

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

Answer: A

NEW QUESTION 23
A gaming company has implemented an API as a Mule application and deployed the API implementation to a CloudHub 2.0 private space. The API implementation must connect to a mainframe application running in the customer??s on-premises corporate data center and also to a Kafka cluster running in an Amazon AWS VPC.
What is the most efficient way to enable the API to securely connect from its private space to the mainframe application and Kafka cluster?

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

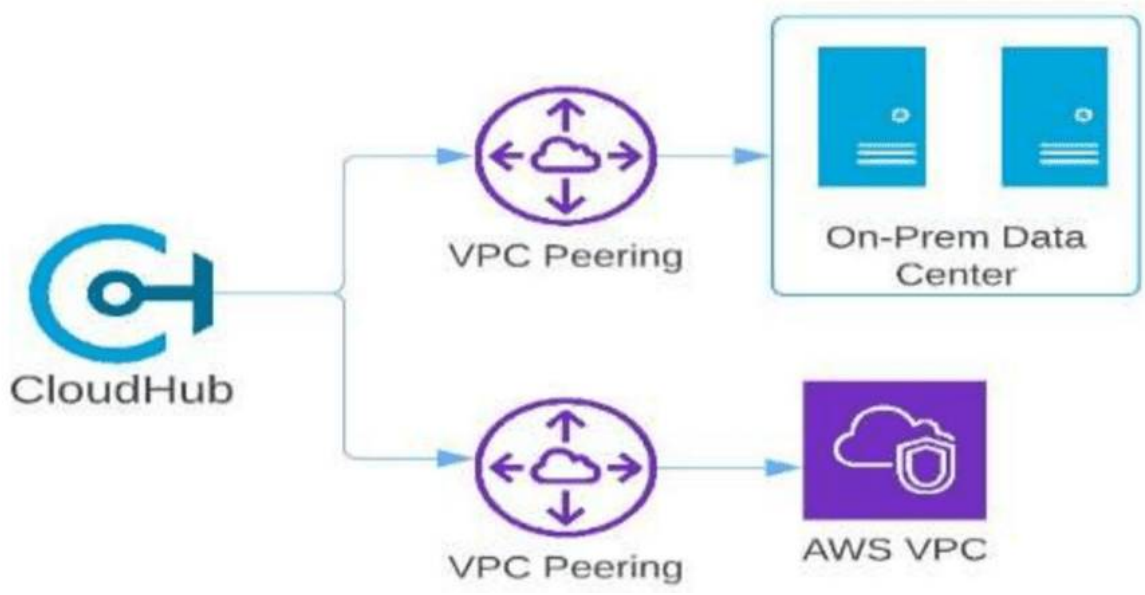

B) In the AWS account, attach the CloudHub 2.0 private space to an AWS transit gateway that routes from the CloudHub 2.0 private space to the on-premises data cente
B. InRuntime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the AWS VPC.

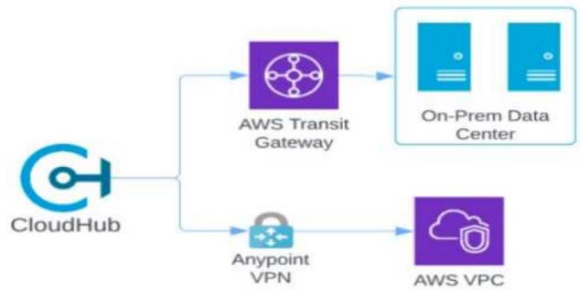

C) In Runtime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the on-premises data cente
C. In the AWS account, attach the private..

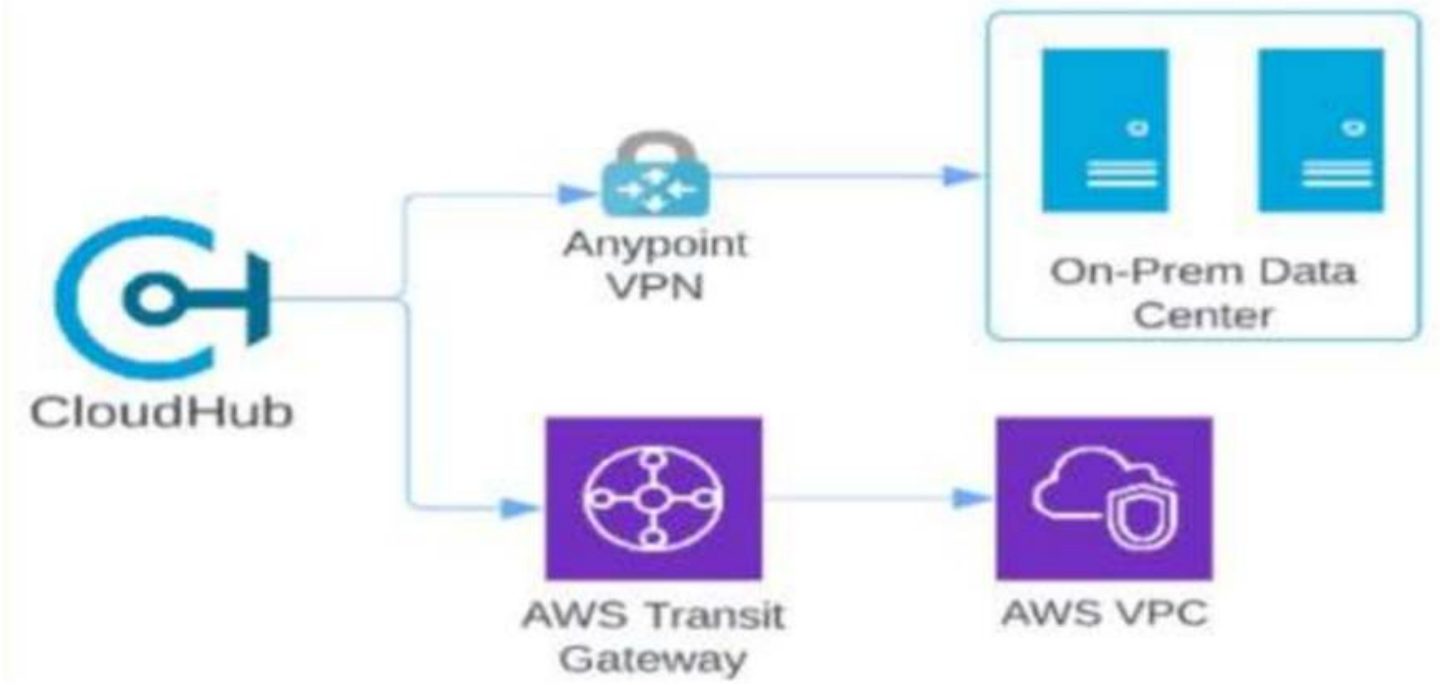

D) In the AWS account, attach the private space directly to the AWS VPC, In the AWS account, use an AWS transit gateway to route from the AWS VPC to the on-premises data center.

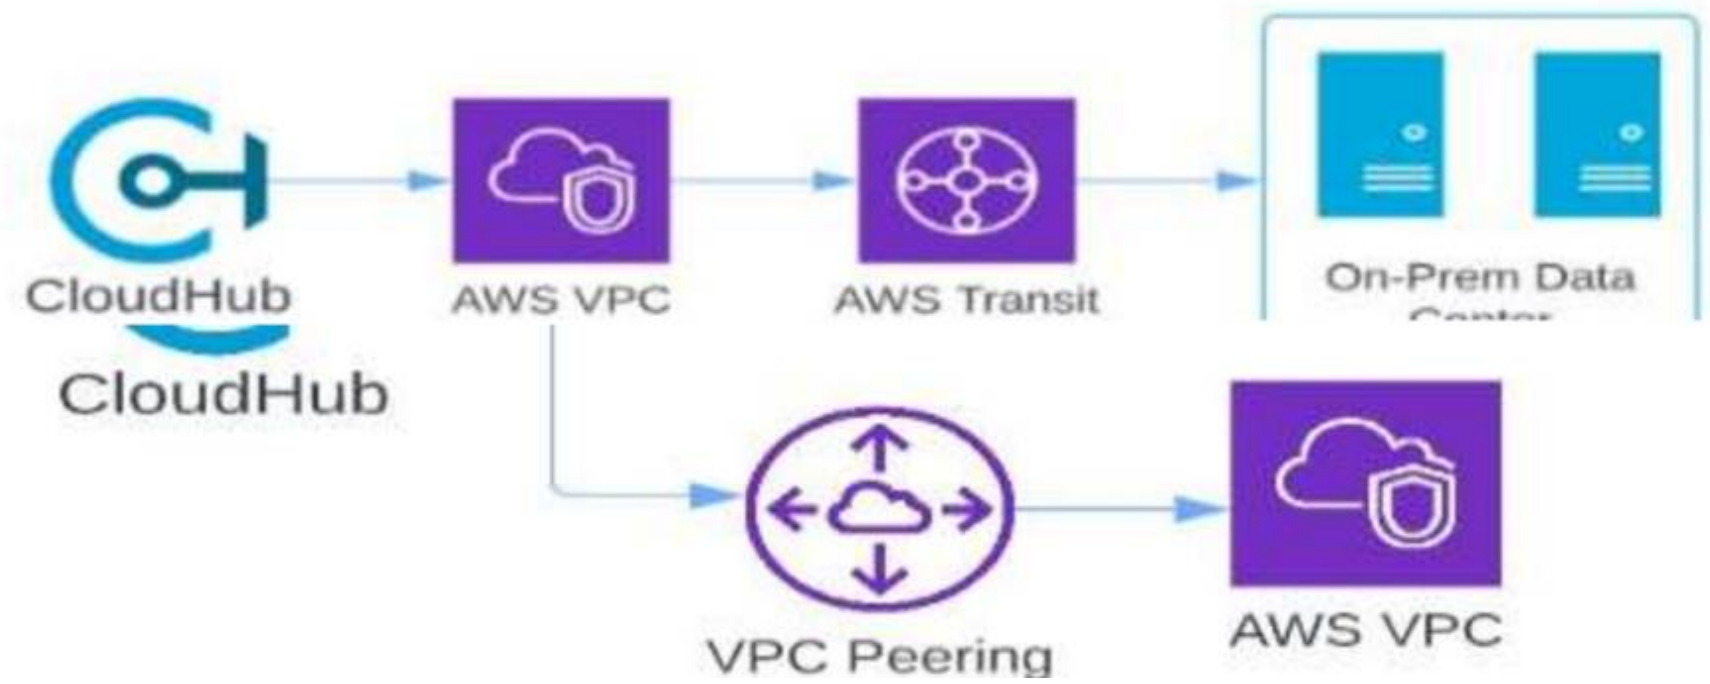

D. Option A
E. Option B
F. Option C
G. Option D

Answer: B

NEW QUESTION 24
Which Mulesoft feature helps users to delegate their access without sharing sensitive credentials or giving full control of accounts to 3rd parties?

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

Answer: C

NEW QUESTION 27
A developer needs to discover which API specifications have been created within the organization before starting a new project.
Which Anypoint Platform component can the developer use to find and try out the currently released API specifications?

A. Anypoint Exchange
B. Runtime Manager
C. API Manager
D. Object Store

Answer: A

NEW QUESTION 29
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: CE

NEW QUESTION 31
Refer to the exhibit.

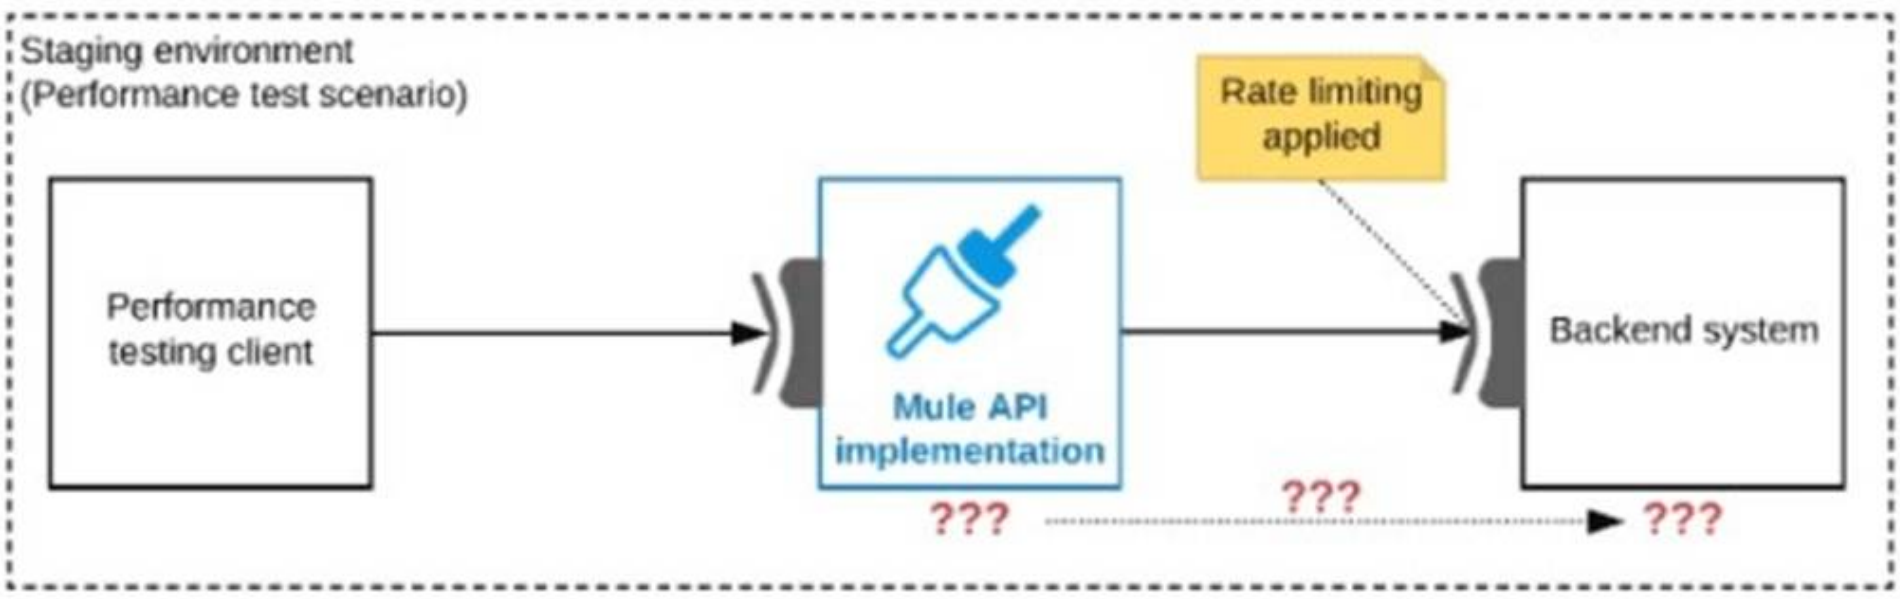

One of the backend systems invoked by an API implementation enforces rate limits on the
number of requests a particular client can make. Both the backend system and the API implementation are deployed to several non-production environments in addition to production.
Rate limiting of the backend system applies to all non-production environments. The production environment, however, does NOT have any rate limiting.
What is the most effective approach to conduct performance tests of the API implementation in a staging (non-production) environment?

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

Answer: A

NEW QUESTION 35
According to MuleSoft's API development best practices, which type of API development approach starts with writing and approving an API contract?

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

Answer: D

NEW QUESTION 39
An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/'CD pipeline. Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?

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

Answer: C

NEW QUESTION 43
An organization has decided on a cloud migration strategy to minimize the organization's own IT resources. Currently the organization has all of its new applications running on its own premises and uses an on-premises load balancer that exposes all APIs under the base URL (https://api.rutujar.com).
As part of migration strategy, the organization is planning to migrate all of its new applications and load balancer CloudHub.
What is the most straightforward and cost-effective approach to Mule application deployment and load balancing that preserves the public URL's?

A. Deploy the Mule application to CloudhubCreate a CNAME record for base URL(httpsr://api.rutujar.com) in the Cloudhub shared load balancer that points to the A record of theon-premises load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications

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

Answer: C

NEW QUESTION 48
Mule application muleA deployed in cloudhub uses Object Store v2 to share data across instances. As a part of new requirement , application muleB which is deployed in same region wants to access this Object Store.
Which of the following option you would suggest which will have minimum latency in this scenario?

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

Answer: A

NEW QUESTION 53
What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

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

Answer: A

NEW QUESTION 56
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 60
Which component of Anypoint platform belongs to the platform control plane?

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

Answer: D

NEW QUESTION 63
When the mule application using VM is deployed to a customer-hosted cluster or multiple cloudhub workers, how are messages consumed by the Mule engine?

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

Answer: A

NEW QUESTION 66
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 68
A platform architect includes both an API gateway and a service mesh in the architect of a distributed application for communication management.
Which type of communication management does a service mesh typically perform in this architecture?

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

Answer: C

NEW QUESTION 72
According to MuteSoft, which principle is common to both Service Oriented Architecture (SOA) and API-led connectivity approaches?

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

Answer: C

NEW QUESTION 73
An external web UI application currently accepts occasional HTTP requests from client web browsers to change (insert, update, or delete) inventory pricing information in an inventory system's database. Each inventory pricing change must be transformed and then synchronized with multiple customer experience systems in near real-time (in under 10 seconds). New customer experience systems are expected to be added in the future.
The database is used heavily and limits the number of SELECT queries that can be made to the database to 10 requests per hour per user.
What is the most scalable, idiomatic (used for its intended purpose), decoupled. reusable, and maintainable integration mechanism available to synchronize each inventory pricing change with the various customer experience systems in near real-time?

A. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the watermark attribute set to an appropriate database columnIn the same now, use a Scatter-Gather to call each customer experience system's REST API with transformed inventory-pricing records
B. Add a trigger to the inventory-pricing database table so that for each change to the inventory pricing database, a stored procedure is called that makes a REST call to a Mule applicationWrite the Mule application to publish each Mule event as a message to an Anypoint MQ exchangeWrite other Mule applications to subscribe to the Anypoint MQ exchange, transform each received message, and then update the Mule application's corresponding customer experience system(s)
C. Replace the external web UI application with a Mule application to accept HTTP requests from client web browsersIn 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 columnIn the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queueWrite 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 76
An insurance company is implementing a MuleSoft API to get inventory details from the two vendors. Due to network issues, the invocations to vendor applications are getting timed- out intermittently. But the transactions are successful upon reprocessing
What is the most performant way of implementing this requirement?

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

Answer: A

NEW QUESTION 78
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 83
Refer to the exhibit.

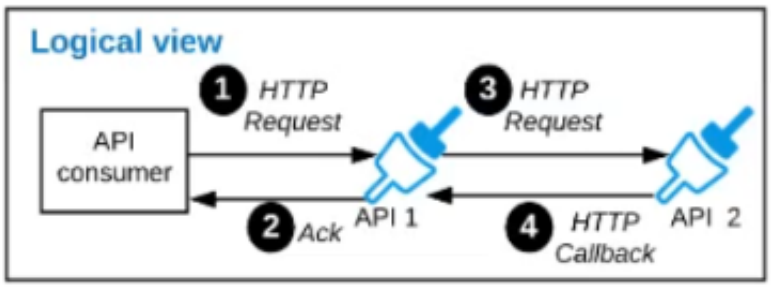

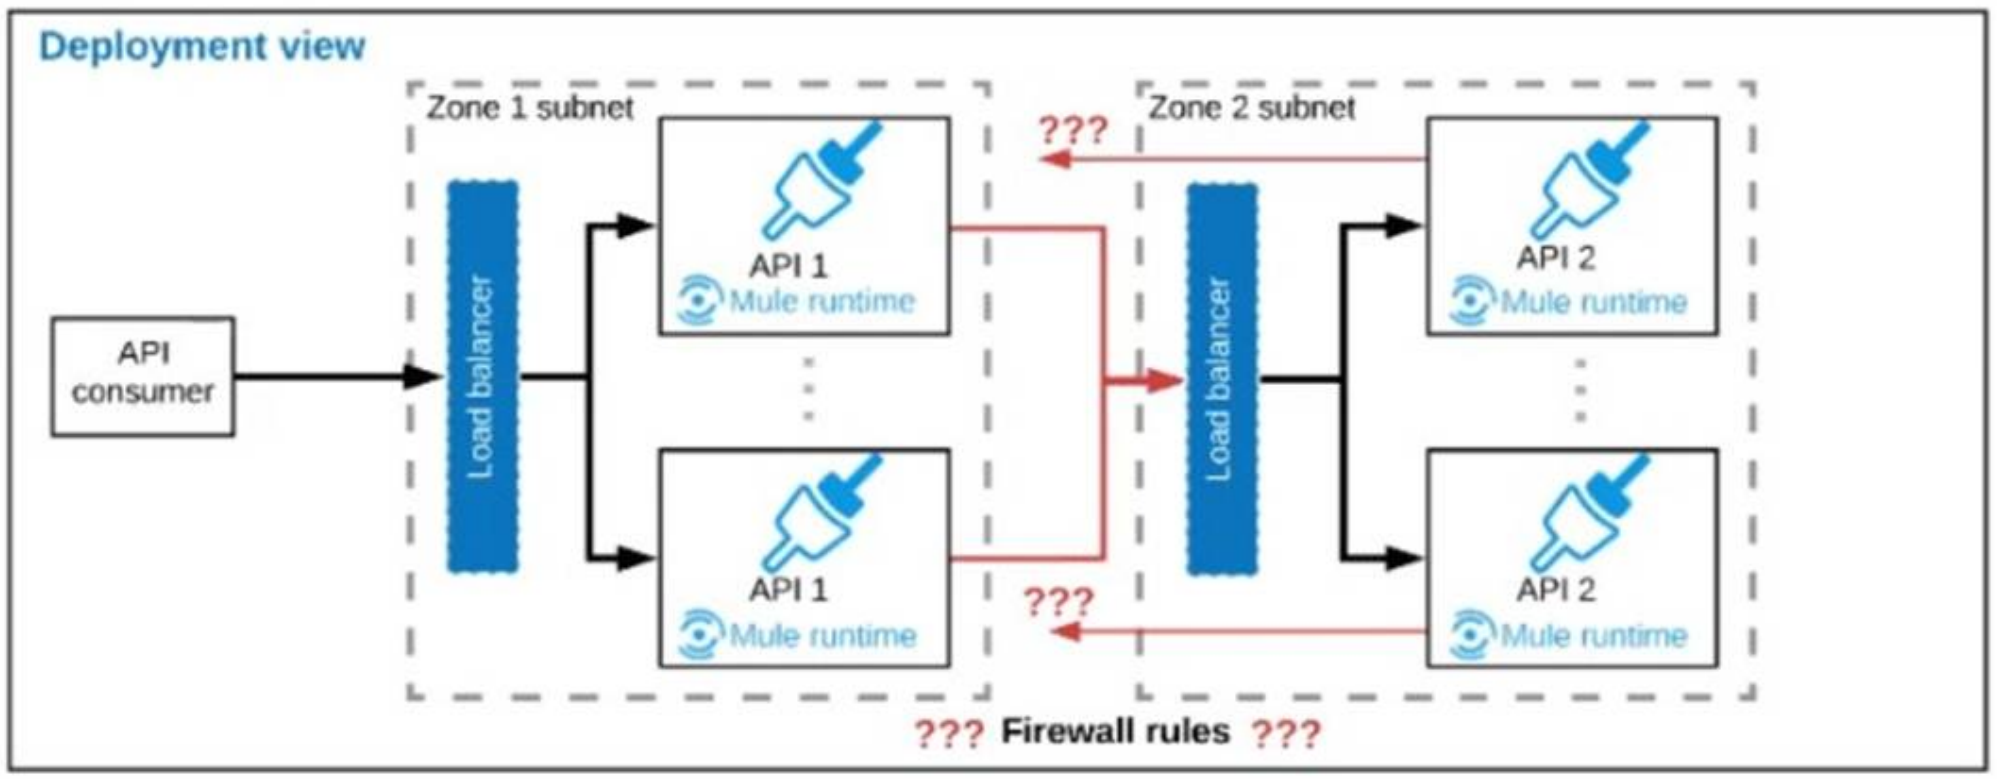

A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API l to notify about completion of the asynchronous process.
Each API Is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.
In the network architecture, how must the firewall rules be configured to enable the above Interaction between API 1 and API 2?

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

Answer: B

NEW QUESTION 86
What is the MuleSoft-recommended best practice to share the connector and configuration information among the APIs?

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

Answer: A

NEW QUESTION 91
An organization plans to extend its Mule APIs to the EU (Frankfurt) region.
Currently, all Mule applications are deployed to CloudHub 1.0 in the default North American region, from the North America control plane, following this naming convention: {API- name}—{environment} (for example, Orderssapi—dev, Orders-sapi-—qa, Orders-sapi-—prod, etc.).
There is no network restriction to block communications between APIs.
What strategy should be implemented in order to deploy the same Mule APIs to the CloudHub 1.0 EU region from the North America control plane, as well as to minimize latency between APIs and target users and systems in Europe?

A. In Runtime Manager, for each Mule application deployment, set the Region property to EU (Frankfurt) and reuse the same Mule application mame as in the North American region.Communicate the new urls {API-name}—{environment}.de-ci.cloudhub.io to the consuming API clients In Europe.
B. In API Manager, set the Region property to EU (Frankfurt) to create an API proxy named {API-name}—proxy—{environment} for each Mule application.Communicate the new url {API-name}—proxy—{environment}.de-c1.cloudhub.io to the consuming API clients In Europe.
C. In Runtime Manager, for each Mule application deployment, leave the Region property blank (default) and change the Mule application name to {API-name}—{environment).de-cl.Communicate the new urls {API-name}—{environment}.de-ci1.cloudhub.io to the consuming API clients in Europe.
D. In API Manager, leave the Region property blank (default) to deploy an API proxy named {API-name}~proxy~- (environment}.de-cl for each Mule applicatio
E. Communicate the new url {API-name}—proxy—{environment}.de-cl.cloudhub.io to the consuming API clients in Europe.

Answer: A

NEW QUESTION 93
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 97
An organization heeds to procure an enterprise software system to increase cross-selling opportunities and better rack prospect data.
Which category of enterprise software has these core capabilities, when used for its typical and intended purpose?

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

Answer: D

NEW QUESTION 98
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 102
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 104
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 106
A company is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTP5 POST and must be acknowledged immediately.
Once acknowledged the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to the rejections from the back-end system will need to be processed manually (outside the banking system).
The mule application will be deployed to a customer hosted runtime and will be able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization's firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.
Which combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

A. One or more On Error scopes to assist calling the back-end system An Untill successful scope containing VM components for long retries A persistent dead-letter VM queue configure in Cloud hub
B. An Until Successful scope to call the back-end system One or more ActiveMQ long-retry queues One or more ActiveMQ dead-letter queues for manual processing
C. One or more on-Error scopes to assist calling the back-end system one or more ActiveMQ long-retry queues A persistent dead-letter Object store configuration in the CloudHub object store service
D. A batch job scope to call the back in system An Untill successful scope containing Object Store components for long retrie

E. A dead-letter object store configured in the Mule application

Answer: B

NEW QUESTION 109
A Mule application is deployed to a cluster of two(2) cusomter-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 form Bob node Alice is now the primary node
C. Files are polled from Alice node Bob is the now the primary node
D. Files are polled form Bob node Bob is now the primary node

Answer: D

NEW QUESTION 110
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: C

NEW QUESTION 111
A manufacturing company plans 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 choice of runtime plane meets these organizational requirements?

A. CloudHub 2.0
B. Customer-hosted self-provisioned runtime plane
C. Anypoint Service Mesh
D. Anypoint Runtime Fabric

Answer: D

NEW QUESTION 115
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 120
An organization is designing an integration solution to replicate financial transaction data from a legacy system into a data warehouse (DWH).
The DWH must contain a daily snapshot of financial transactions, to be delivered as a CSV file. Daily transaction volume exceeds tens of millions of records, with significant spikes in volume during popular shopping periods.
What is the most appropriate integration style for an integration solution that meets the organization's current requirements?

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

Answer: D

NEW QUESTION 121
An organization is creating a Mule application that will be deployed to CloudHub. The Mule application has a property named dbPassword that stores a database user??s password.
The organization's security standards indicate that the dbPassword property must be hidden from every Anypoint Platform user after the value is set in the Runtime Manager Properties tab.
What configuration in the Mule application helps hide the dbPassword property value in Runtime Manager?

A. Use secure::dbPassword as the property placeholder name and store the cleartext (unencrypted) value in a secure properties placeholder file
B. Use secure::dbPassword as the property placeholder name and store the property encrypted value in a secure properties placeholder file
C. Add the dbPassword property to the secureProperties section of the pom.xml file
D. Add the dbPassword property to the secureProperties section of the mule-artifact.json file

Answer: B

NEW QUESTION 123
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

NEW QUESTION 128
An organization is evaluating using the CloudHub shared Load Balancer (SLB) vs creating a CloudHub dedicated load balancer (DLB). They are evaluating how this choice affects the various types of certificates used by CloudHub deployed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates. What type of restrictions exist on the types of certificates for the service that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

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

Answer: B

NEW QUESTION 133
Which role is primarily responsible for building API implementation as part of a typical MuleSoft integration project?

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

Answer: A

NEW QUESTION 134
As a part of business requirement , old CRM system needs to be integrated using Mule application. CRM system is capable of exchanging data only via SOAP/HTTP protocol. As an integration architect who follows API led approach , what is the the below step you will perform so that you can share document with CRM team?

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

Answer: C

NEW QUESTION 136
An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger of flow for processing a high-volume set of records from the source system and synchronize with the SaaS system using the Batch job scope
After processing 1000 records in a periodic synchronization of 1 lakh records, the replicas in which batch job instance was started went down due to unexpected failure in the runtime fabric environment
What is the consequence of losing the replicas that run the Batch job instance?

A. The remaining 99000 records will be lost and left and processed
B. The second replicas will take over processing the remaining 99000 records
C. A new replacement replica will be available and will be process all 1,00,000 records from scratch leading to duplicate record processing
D. A new placement replica will be available and will take or processing the remaining 99,000 records

Answer: D

NEW QUESTION 140
An architect is designing a Mule application to meet the following two requirements:
* 1. The application must process files asynchronously and reliably from an FTPS server to a back-end database using VM intermediary queues for load-balancing Mule events.
* 2. The application must process a medium rate of records from a source to a target system using a Batch Job scope.
To make the Mule application more reliable, the Mule application will be deployed to two CloudHub 1.0 workers.
Following MuleSoft-recommended best practices, how should the Mule application deployment typically be configured in Runtime Manger to best support the performance and reliability goals of both the Batch Job scope and the file processing VM queues?

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

Answer: A

NEW QUESTION 142
The company's FTPS server login username and password

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

Answer: C

NEW QUESTION 146
An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders.
The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS).
At a minimum, what must be stored in each truststore and keystore of these two Mule applications to properly support two-way TLS between the two Mule applications while properly protecting each Mule application's keys?

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

Answer: C

NEW QUESTION 148
What is not true about Mule Domain Project?

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

Answer: C

NEW QUESTION 151
Refer to the exhibit.

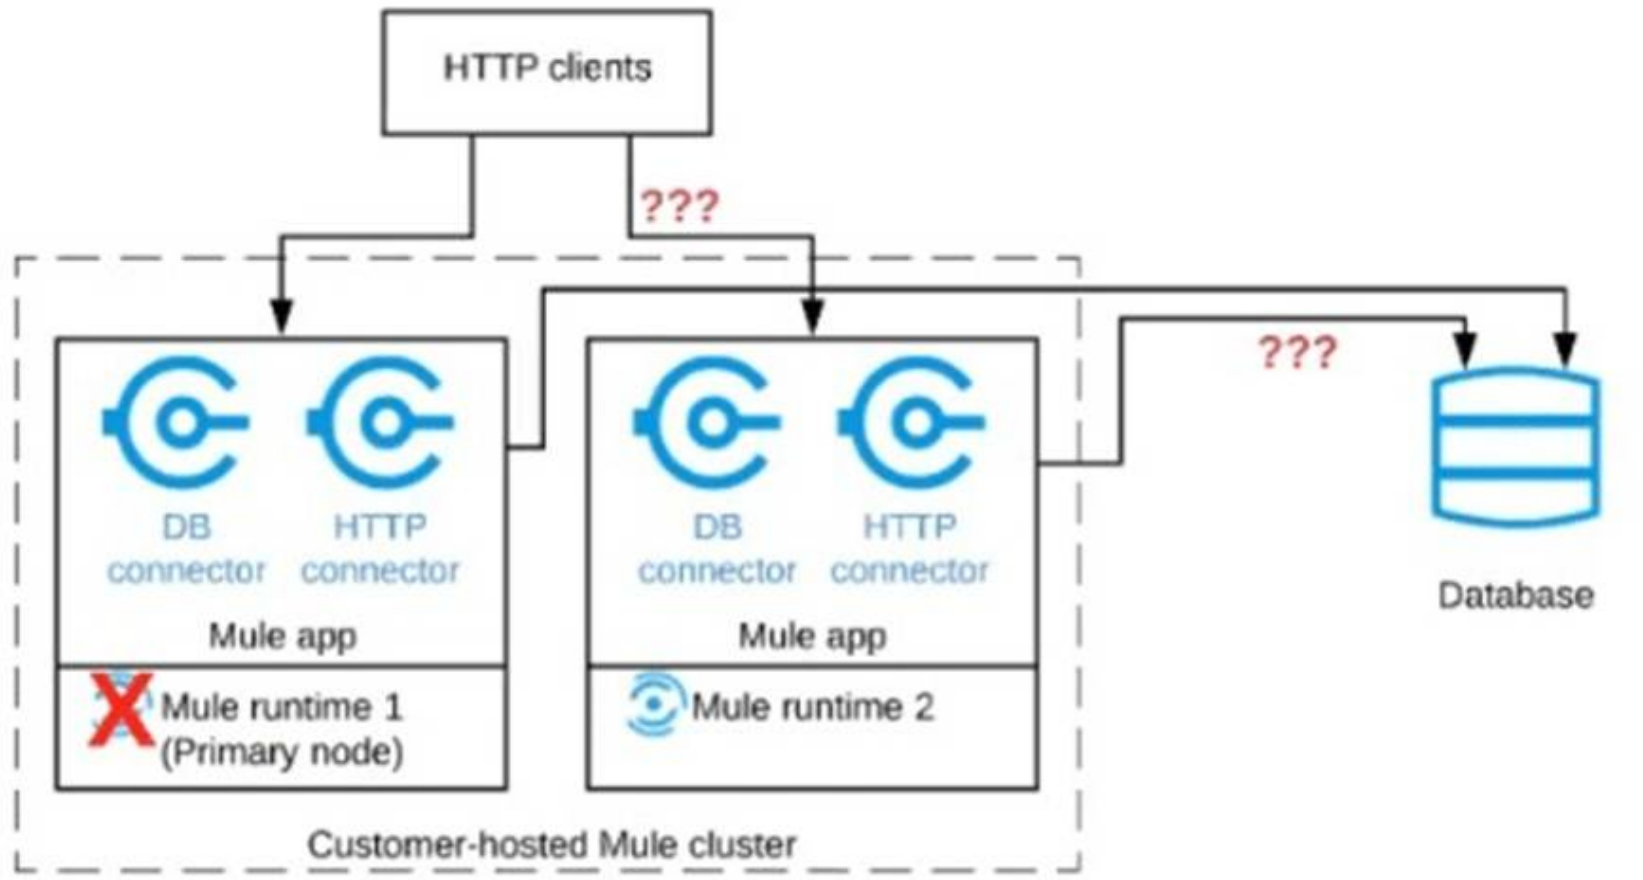

A Mule application is deployed to a cluster of two customer-hosted Mute runtimes. The Mute application has a flow that polls a database and another flow with an HTTP Listener.
HTTP clients send HTTP requests directly to individual cluster nodes.
What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has railed, 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 152
An integration team follows MuleSoft??s recommended approach to full lifecycle API development.
Which activity should this team perform during the API implementation phase?

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

Answer: B

NEW QUESTION 154
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 158
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 163
An API client makes an HTTP request to an API gateway with an Accept header containing the value???? application????.
What is a valid HTTP response payload for this request in the client requested data format?

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

Answer: B

NEW QUESTION 165
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 167
An organization is designing the following two Mule applications that must share data via a common persistent object store instance:
- Mule application P will be deployed within their on-premises datacenter.
- Mule application C will run on CloudHub in an Anypoint VPC.
The object store implementation used by CloudHub is the Anypoint Object Store v2 (OSv2).
what type of object store(s) should be used, and what design gives both Mule applications access to the same object store instance?

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

Answer: C

NEW QUESTION 171
A Mule application is being designed to do the following:
Step 1: Read a SalesOrder message from a JMS queue, where each SalesOrder consists of a header and a list of SalesOrderLineItems.
Step 2: Insert the SalesOrder header and each SalesOrderLineItem into different tables in an RDBMS.
Step 3: Insert the SalesOrder header and the sum of the prices of all its SalesOrderLineItems into a table In a different RDBMS.
No SalesOrder message can be lost and the consistency of all SalesOrder-related information in both RDBMSs must be ensured at all times.
What design choice (including choice of transactions) and order of steps addresses these requirements?

A. 1) Read the JMS message (NOT in an XA transaction)2) Perform BOTH DB inserts in ONE DB transaction3) Acknowledge the JMS message
B. 1) Read the JMS message (NOT in an XA transaction)2) Perform EACH DB insert in a SEPARATE DB transaction3) Acknowledge the JMS message
C. 1) Read the JMS message in an XA transaction2) In the SAME XA transaction, perform BOTH DB inserts but do NOT acknowledge the JMS message
D. 1) Read and acknowledge the JMS message (NOT in an XA transaction)2) In a NEW XA transaction, perform BOTH DB inserts

Answer: A

NEW QUESTION 172
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 177
According to MuleSoft, a synchronous invocation of a RESTful API using HTTP to get an individual customer record from a single system is an example of which system integration interaction pattern?

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

Answer: A

NEW QUESTION 179
A corporation has deployed Mule applications to different customer-hosted Mule runtimes. Mule applications deployed to these Mule runtimes are managed by Anypoint Platform.
What needs to be installed or configured (if anything) to monitor these Mule applications from Anypoint Monitoring, and how is monitoring data from each Mule application sent to Anypoint Monitoring?

A. Enable monitoring of individual Mule applications from the Runtime Manager application settings.Runtime Manager sends monitoring data to Anypoint Monitoring for each deployed Mule application.
B. Install a Runtime Manager agent on each Mule runtime.Each Runtime Manager agent sends monitoring data from the Mule applications running in its Mule runtime to Runtime Manager, then Runtime Manager sends monitoring data to Anypoint Monitoring.
C. Leave the out-of-the-box Anypoint Monitoring agent unchanged in its default Mule runtime installation.Each Anypoint Monitoring agent sends monitoring data from the Mule applications running in its Mule runtime to Runtime Manager, then Runtime Manager sends monitoring data to Anypoint Monitoring.
D. Install an Anypoint Monitoring agent on each Mule runtime.Each Anypoint Monitoring agent sends monitoring data from the Mule applications running in its Mule runtime to Anypoint Monitoring.

Answer: D

NEW QUESTION 182
An organization is not meeting its growth and innovation objectives because IT cannot deliver projects last enough to keep up with the pace of change required by the business.
According to MuleSoft??s IT delivery and operating model, which step should the organization lake to solve this problem?

A. Modify IT governance and security controls so that line of business developers can have direct access to the organization's systems of record
B. Switch from a design-first to a code-first approach for IT development
C. Adopt a new approach that decouples core IT projects from the innovation that happens within each line of business
D. Hire more |T developers, architects, and project managers to increase IT delivery

Answer: C

NEW QUESTION 183
An organization if struggling frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects.

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

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

Answer: D

NEW QUESTION 188
According to the Internet Engineering Task Force (IETF), which supporting protocol does File Transfer Protocol (FTP) use for reliable communication?

A. A Secure Sockets Layer (SSL)
B. B Transmission Control Protocol (TCP)
C. Lightweight Directory Access Protocol (LDAP)
D. Hypertext Transfer Protocol (HTTP)

Answer: B

NEW QUESTION 191
A Mule application contains a Batch Job with two Batch Steps (Batch_Step_l and Batch_Step_2). A payload with 1000 records is received by the Batch Job.
How many threads are used by the Batch Job to process records, and how does each Batch Step process records within the Batch Job?

A. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and RECORDS are processed IN PARALLEL within and between the two Batch Steps
B. Each Batch Job uses a SINGLE THREAD for all Batch steps Each Batch step instance receives ONE record at a time as the payload, and RECORDS are processed IN ORDER, first through Batch_Step_l and then through Batch_Step_2
C. Each Batch Job uses a SINGLE THREAD to process a configured block size of record Each Batch Step instance receives A BLOCK OF records as the payload, and BLOCKS of records are processed IN ORDER
D. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and BATCH STEP INSTANCES execute IN PARALLEL to process records and Batch Steps in ANY order as fast as possible

Answer: A

NEW QUESTION 196
A leading bank implementing new mule API.
The purpose of API to fetch the customer account balances from the backend application and display them on the online platform the online banking platform. The online banking platform will send an array of accounts to Mule API get the account balances.
As a part of the processing the Mule API needs to insert the data into the database for auditing purposes and this process should not have any performance related implications on the account balance retrieval flow
How should this requirement be implemented to achieve better throughput?

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

Answer: C

NEW QUESTION 198
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 200
Refer to the exhibit.

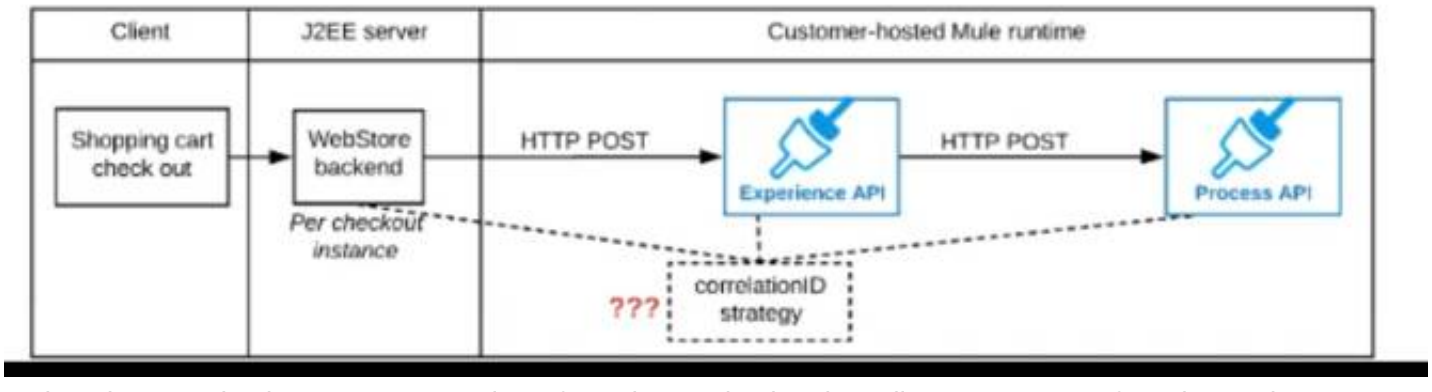

A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

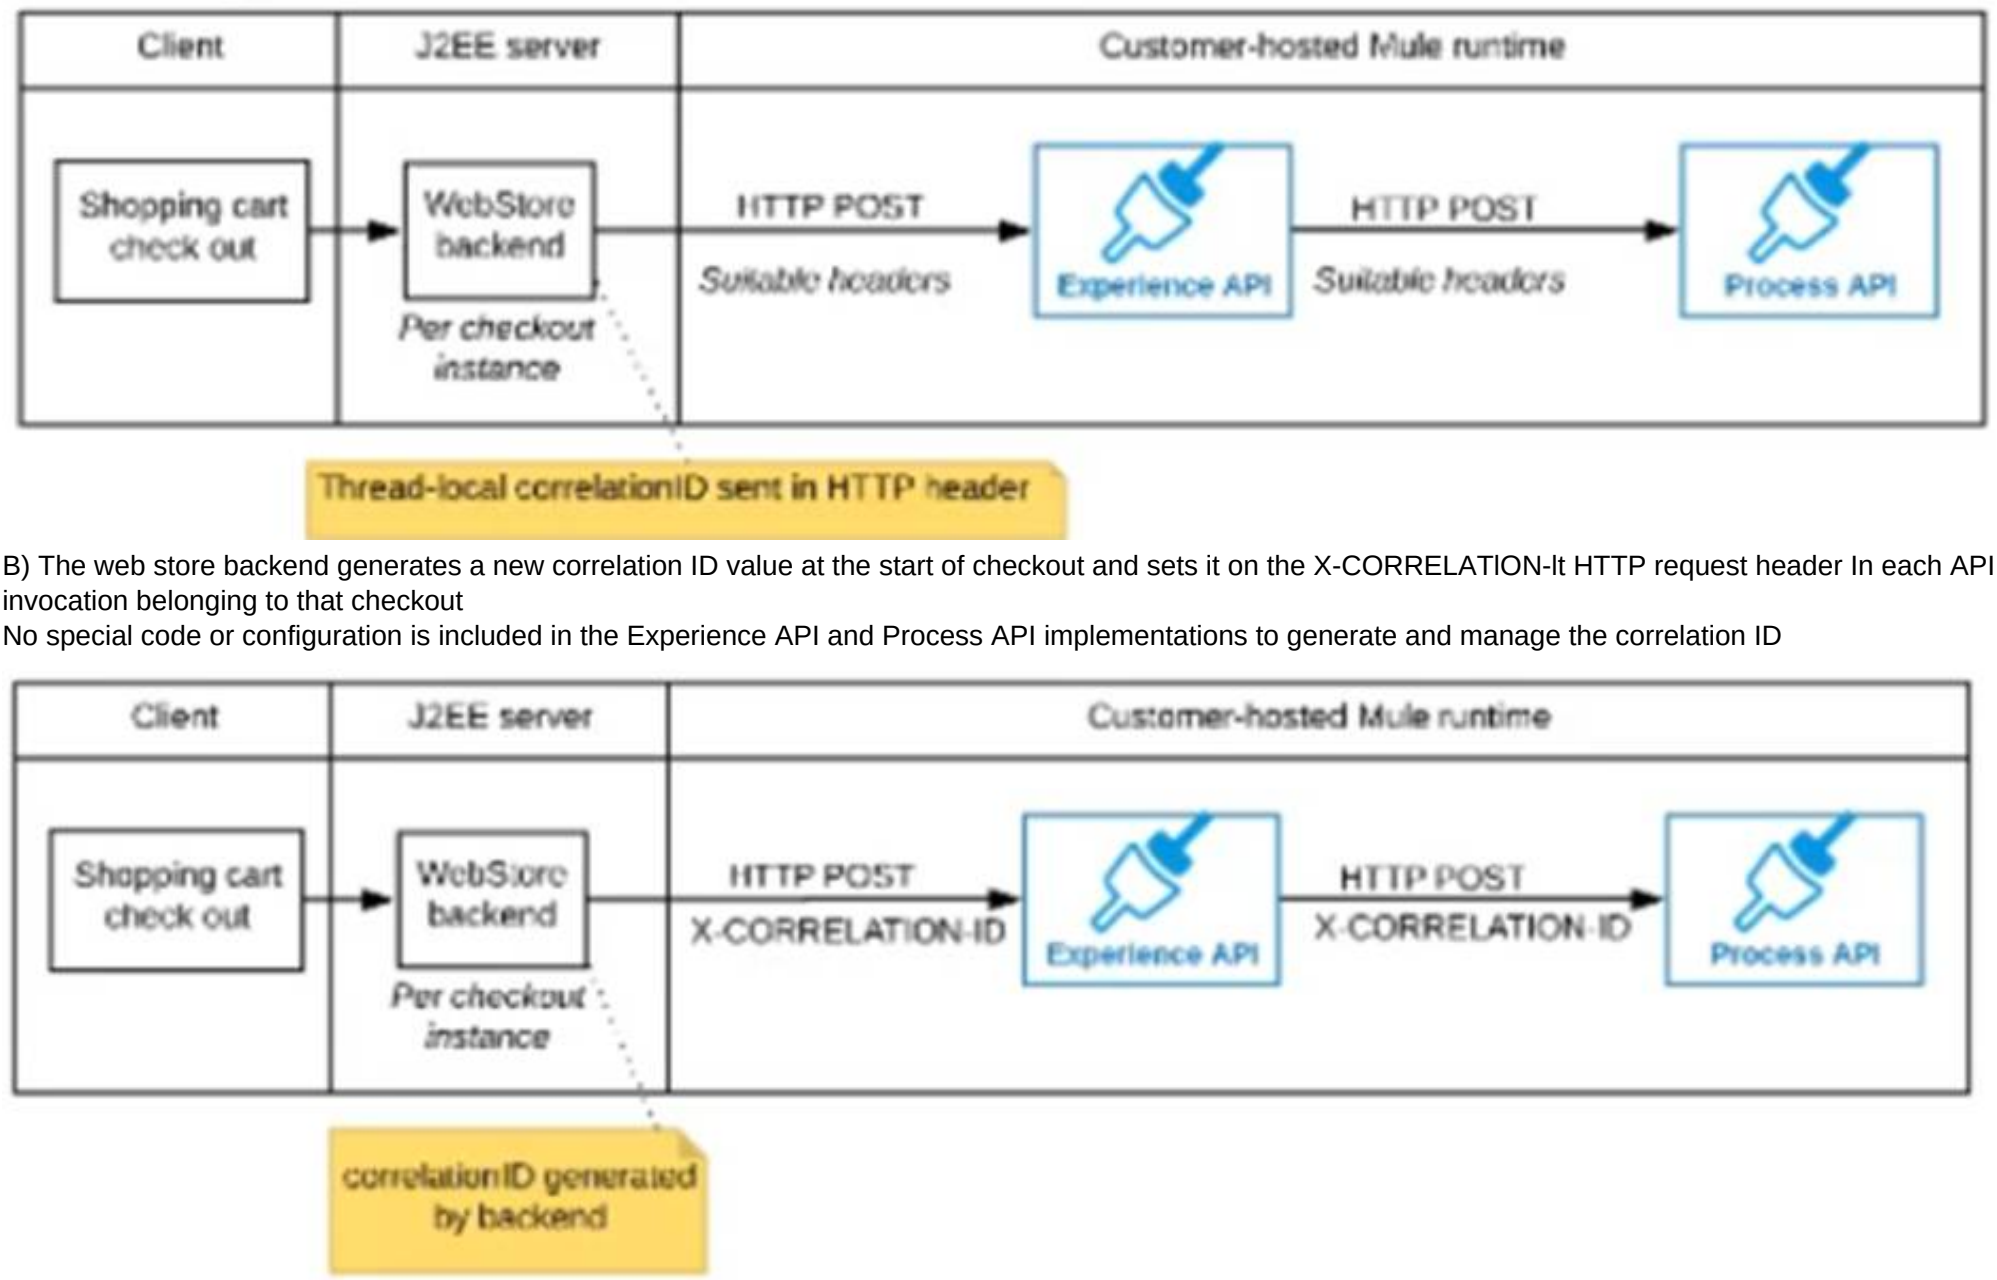

B) The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-It HTTP request header In each API invocation belonging to that checkout
No special code or configuration is included in the Experience API and Process API implementations to generate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

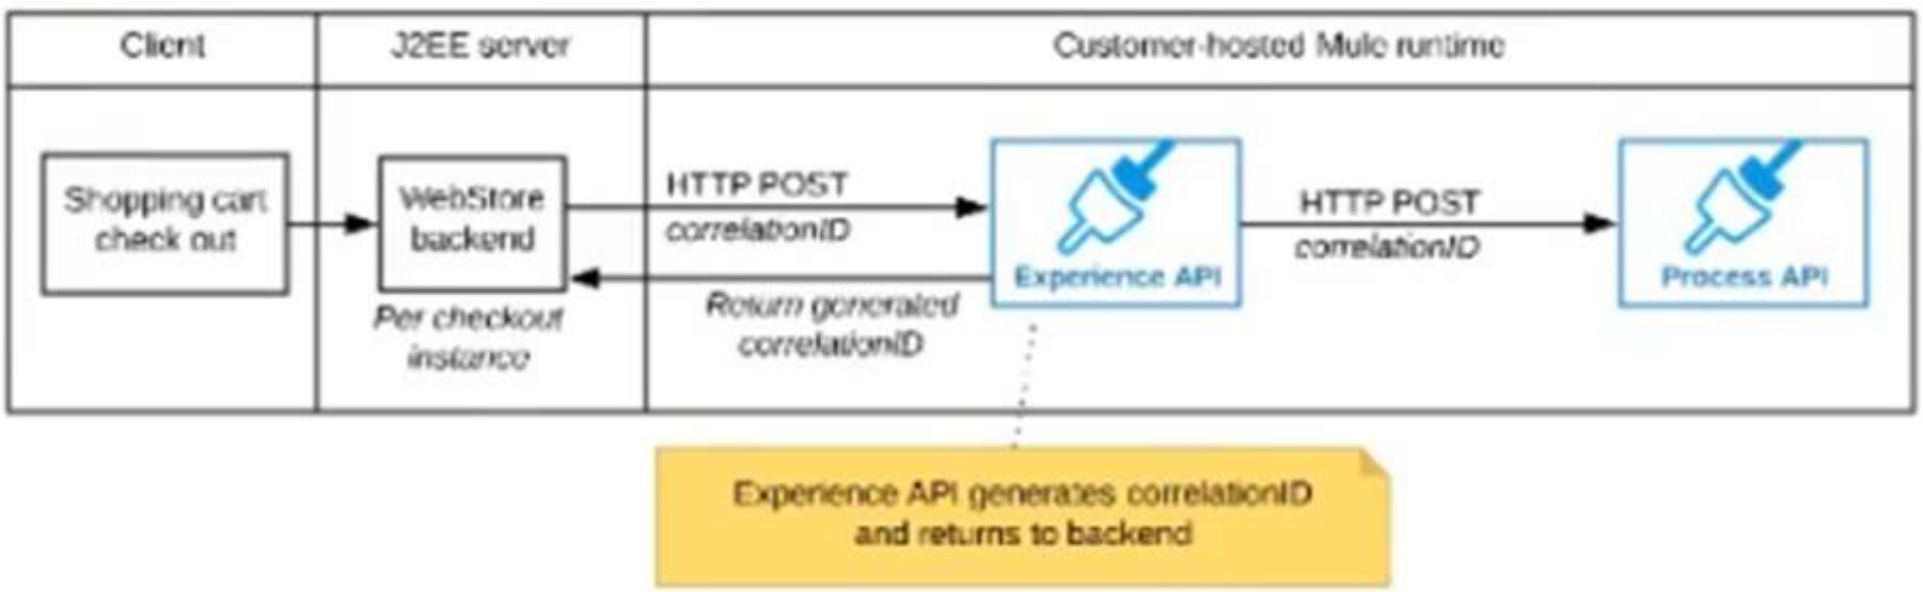

D) The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API
The Experience API and Process API implementations must be coded to receive the custom correlation ID In the HTTP requests and propagate It in suitable HTTP request headers

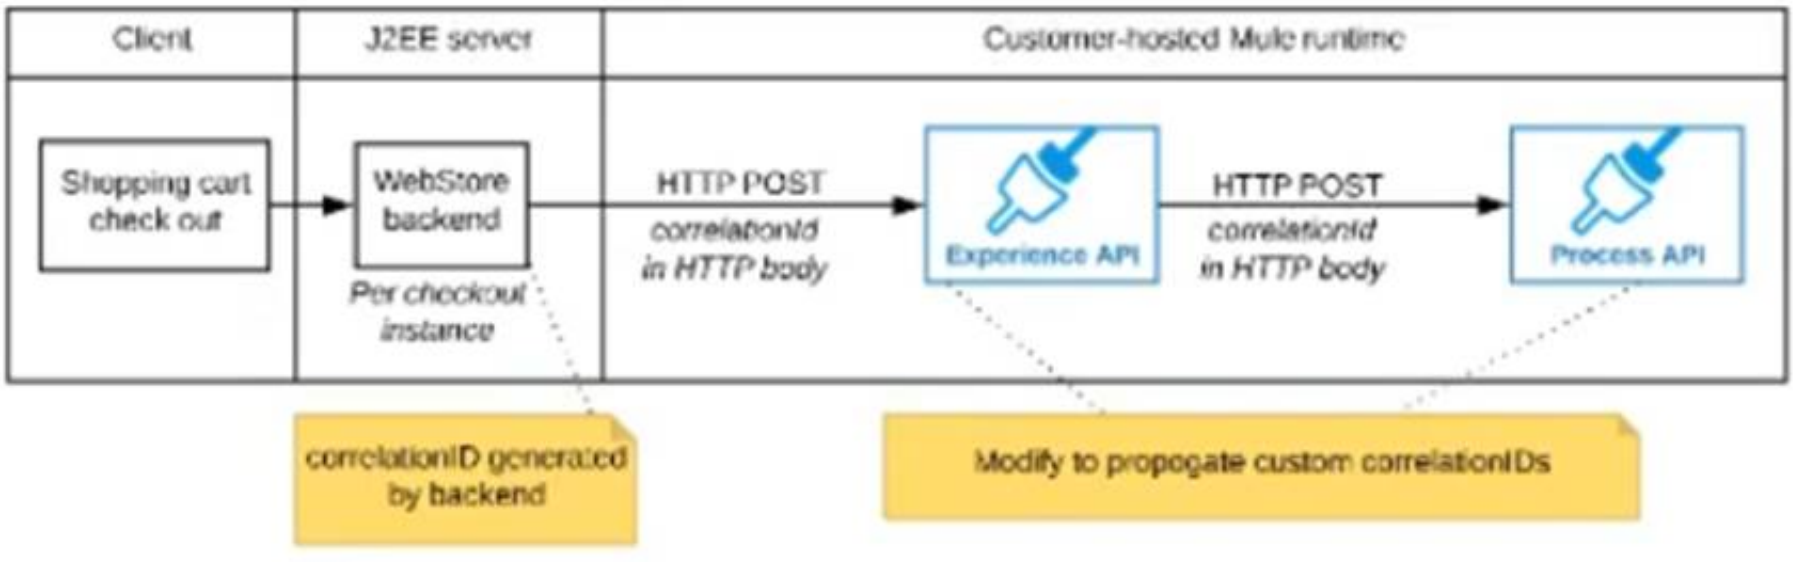

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

NEW QUESTION 204
Refer to the exhibit.

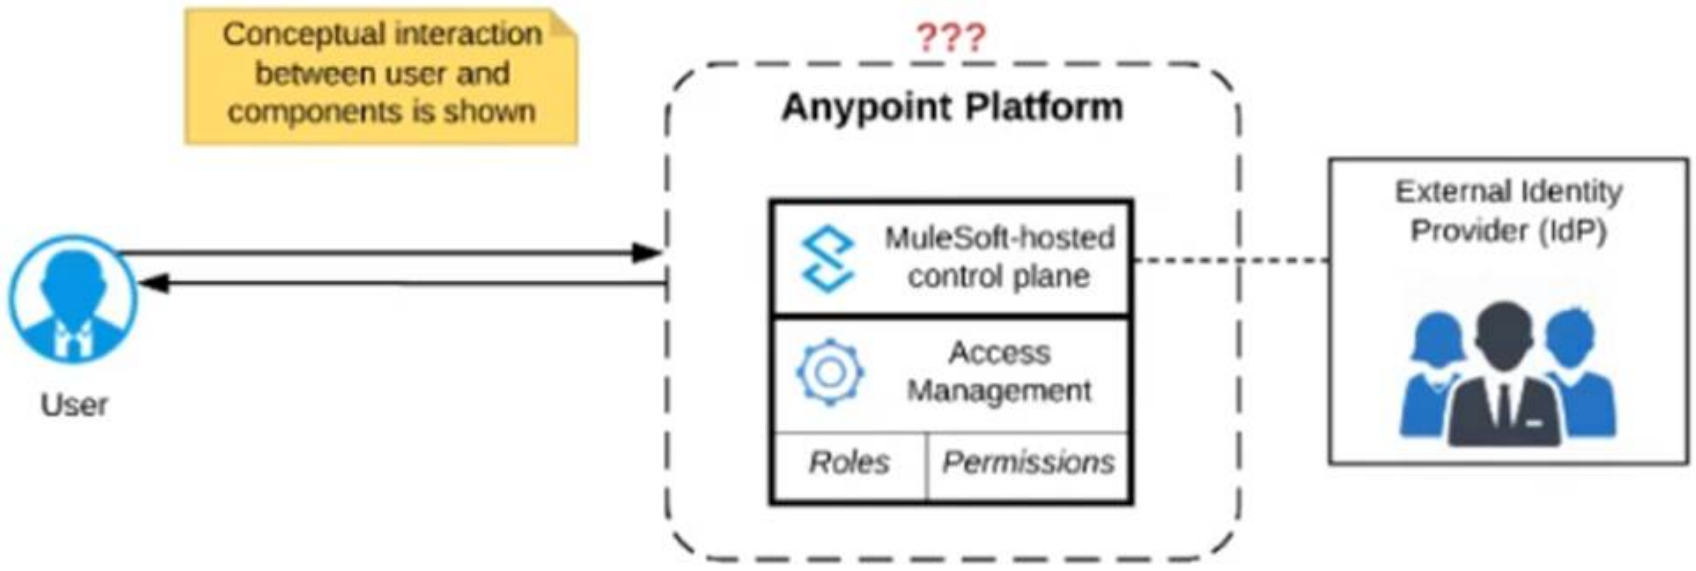

Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.
What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled 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 206
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 208
An integration architect is designing an API that must accept requests from API clients for both XML and JSON content over HTTP/1.1 by default.
Which API architectural style, when used for its intended and typical purposes, should the architect choose to meet these requirements?

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

Answer: C

NEW QUESTION 211
An Integration Mule application is being designed to synchronize customer data between two systems. One system is an IBM Mainframe and the other system is a Salesforce Marketing Cloud (CRM) instance. Both systems have been deployed in their typical configurations, and are to be invoked using the native protocols provided by Salesforce and IBM.
What interface technologies are the most straightforward and appropriate to use in this Mute application to interact with these systems, assuming that Anypoint Connectors exist that implement these interface technologies?

A. IBM: DB access CRM: gRPC
B. IBM: REST CRM:REST
C. IBM: Active MQ CRM: REST
D. IBM: CICS CRM: SOAP

Answer: D

NEW QUESTION 213
A company is implementing a new Mule application that supports a set of critical functions driven by a rest API enabled, claims payment rules engine hosted on oracle ERP. As designed the mule application requires many data transformation operations as it performs its batch processing logic.
The company wants to leverage and reuse as many of its existing java-based capabilities (classes, objects, data model etc.) as possible
What approach should be considered when implementing required data mappings and transformations between Mule application and Oracle ERP in the new Mule application?

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

Answer: C

NEW QUESTION 214
A company is building an application network and has deployed four Mule APIs: one experience API, one process API, and two system APIs. The logs from all the APIs are aggregated in an external log aggregation tool. The company wants to trace messages that are exchanged between multiple API implementations. What is the most idiomatic (based on its intended use) identifier that should be used to implement Mule event tracing across the multiple API implementations?

A. Mule event ID
B. Mule correlation ID
C. Client's IP address
D. DataWeave UUID

Answer: B

NEW QUESTION 216
What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

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

Answer: D

NEW QUESTION 217
An organization is choosing between API-led connectivity and other integration approaches.
According to MuleSoft, which business benefits is associated with an API-led connectivity approach using Anypoint Platform?

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

Answer: B

NEW QUESTION 221
A mule application designed to fulfil two requirements
a) Processing files are synchronously from an FTPS server to a back-end database using VM intermediary queues for load balancing VM events
b) Processing a medium rate of records from a source to a target system using batch job scope
Considering the processing reliability requirements for FTPS files, how should VM queues be configured for processing files as well as for the batch job scope if the application is deployed to Cloudhub workers?

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

Answer: A

NEW QUESTION 226
An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?

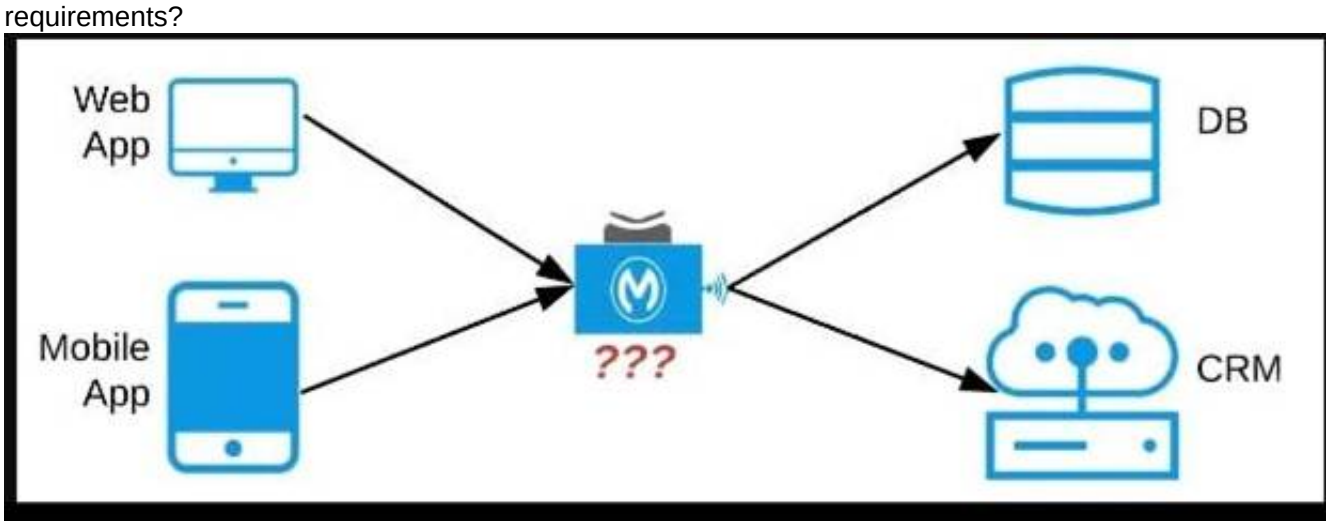

A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
B. One set of APIs (Experience API, Process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

Answer: C

NEW QUESTION 228
An auto mobile company want to share inventory updates with dealers Dl and D2 asynchronously and concurrently via queues Q1 and Q2. Dealer Dl must consume the message from the queue Q1 and dealer D2 to must consume a message from the queue Q2.
Dealer D1 has implemented a retry mechanism to reprocess the transaction in case of any
errors while processing the inventers updates. Dealer D2 has not implemented any retry mechanism.
How should the dealers acknowledge the message to avoid message loss and minimize impact on the current implementation?

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

Answer: D

NEW QUESTION 233
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 234
A new upstream API Is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity. The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms. If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API's desired SLA?

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

Answer: D

NEW QUESTION 238
Which Exchange asset type represents a complete API specification in RAML or OAS format?

A. Connectors
B. REST APIs
C. API Spec Fragments
D. SOAP APIs

Answer: B

NEW QUESTION 243
An insurance company is using a CIoudHub runtime plane. As a part of requirement, email alert should be sent to internal operations team every time of policy applied to an API instance is deleted As an integration architect suggest on how this requirement be met?

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

Answer: D

NEW QUESTION 247
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 251
An ABC Farms project team is planning to build a new API that is required to work with data from different domains across the organization.
The organization has a policy that all project teams should leverage existing investments by reusing existing APIs and related resources and documentation that other project teams have already developed and deployed.
To support reuse, where on Anypoint Platform should the project team go to discover and read existing APIs, discover related resources and documentation, and interact with mocked versions of those APIs?

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

Answer: D

NEW QUESTION 252
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 253
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 255
An organization is building out a test suite for their application using MUnit.
The Integration Architect has recommended using Test Recorder in Anypoint Studio to record the processing flows and then configure unit tests based on the captured events.
What Is a core consideration that must be kept In mind while using Test Recorder?

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

Answer: D

NEW QUESTION 259
An IT integration tram followed an API-led connectivity approach to implement an order- fulfillment business process. It created an order processing AP that coordinates stateful interactions with a variety of microservices that validate, create, and fulfill new product orders
Which interaction composition pattern did the integration architect who designed this order processing AP| use?

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

Answer: A

NEW QUESTION 262
According to MuleSoft??s recommended REST conventions, which HTTP method should an API use to specify how AP\ clients can request data from a specified resource?

A. POST
B. PUT
C. PATCH
D. GET

Answer: D

NEW QUESTION 265
To implement predictive maintenance on its machinery equipment, ACME Tractors has installed thousands of IoT sensors that will send data for each machinery asset as sequences of JMS messages, in near real-time, to a JMS queue named SENSOR_DATA on a JMS server. The Mule application contains a JMS Listener operation configured to receive incoming messages from the JMS servers SENSOR_DATA JMS queue. The Mule application persists each received JMS message, then sends a transformed version of the corresponding Mule event to the machinery equipment back-end systems.
The Mule application will be deployed to a multi-node, customer-hosted Mule runtime cluster. Under normal conditions, each JMS message should be processed exactly once.
How should the JMS Listener be configured to maximize performance and concurrent message processing of the JMS queue?

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

Answer: D

NEW QUESTION 266
49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58. https://api.acme.com/partnefs/vl. The API implementation is estimated to require deployment to 5 CloudHub workers.
The retailer has obtained a public X.509 certificate for the name apl.acme.com, signed by a reputable CA, to be used as the server certificate.
Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

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

Answer: C

NEW QUESTION 267
A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction.
What is the behavior of the Scatter-Gather when running within this local transaction?

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

Answer: A

NEW QUESTION 268
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 271
An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform. What feature of Anypoint Platform should be used to meet this requirement?

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

Answer: C

NEW QUESTION 275
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 279
An Organization has previously provisioned its own AWS VPC hosting various servers. The organization now needs to use Cloudhub to host a Mule application that will implement a REST API once deployed to Cloudhub, this Mule application must be able to communicate securely with the customer-provisioned AWS VPC resources within the same region, without being interceptable on the public internet.
What Anypoint Platform features should be used to meet these network communication requirements between Cloudhub and the existing customer-provisioned AWS VPC?

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

Answer: A

NEW QUESTION 284
An organization has just developed a Mule application that implements a REST API. The mule application will be deployed to a cluster of customer hosted Mule runtimes.
What additional infrastructure component must the customer provide in order to distribute inbound API requests across the Mule runtimes of the cluster?

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

Answer: B

NEW QUESTION 285
During a planning session with the executive leadership, the development team director presents plans for a new API to expose the data in the company??s order database. An earlier effort to build an API on top of this data failed, so the director is recommending a design-first approach.
Which characteristics of a design-first approach will help make this API successful?

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

Answer: D

NEW QUESTION 287
The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure.
Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).
Given the company's current setup and requirements, what is the most idiomatic (used for its intended purpose) way to send Mule application logs to the external LMT?

A. In RTF-VM, install and configure the external LTM's log-forwarding agent
B. In RTF-VM, edit the pod configuration to automatically install and configure an Anypoint Monitoring agent
C. In each Mule application, configure custom Log4j settings

D. In RTF-V
E. configure the out-of-the-box external log forwarder

Answer: A

NEW QUESTION 288
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 292
A mule application is being designed to perform product orchestration. The Mule application needs to join together the responses from an inventory API and a Product Sales History API with the least latency.
To minimize the overall latency. What is the most idiomatic (used for its intended purpose) design to call each API request in the Mule application?

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

Answer: B

NEW QUESTION 294
Refer to the exhibit.

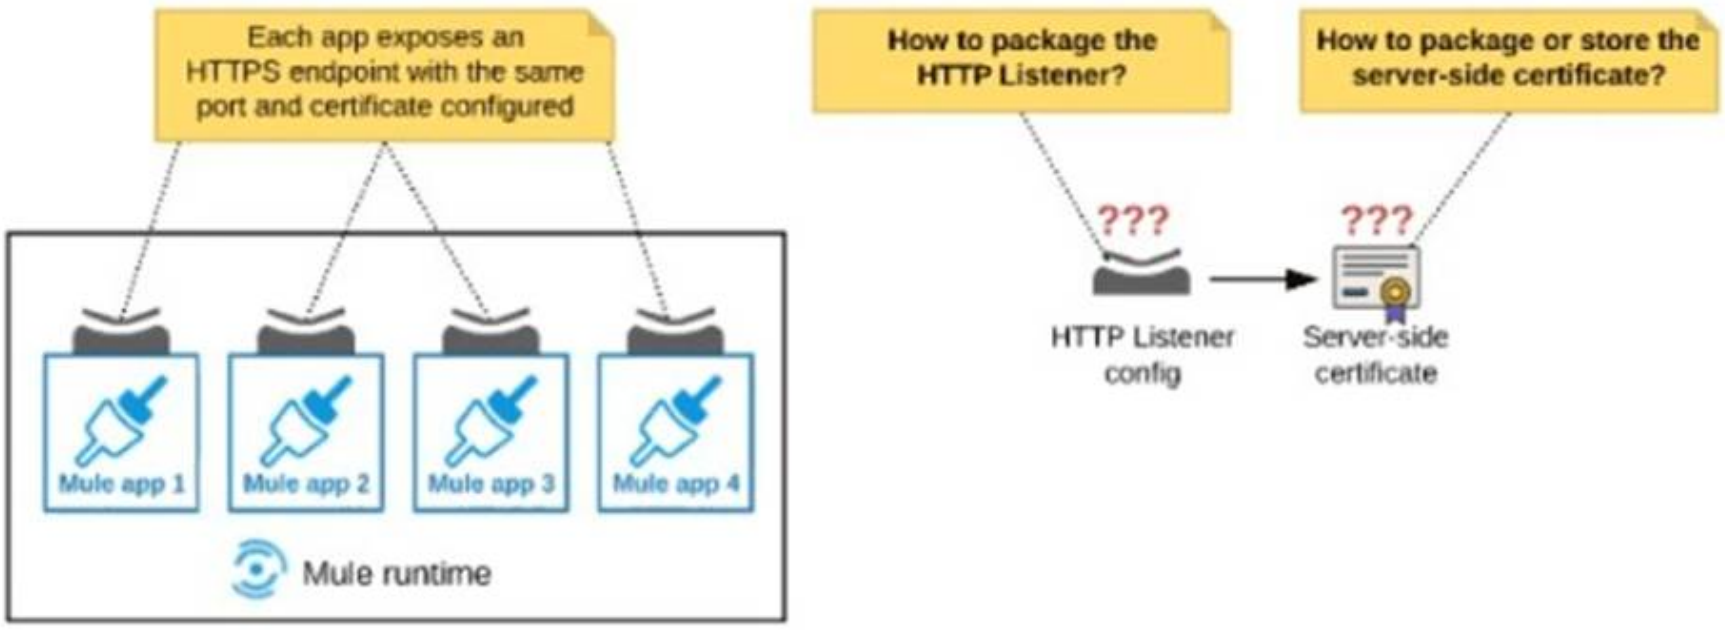

An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a server-side certificate that rotates often.
What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

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

Answer: B

NEW QUESTION 296
Which Salesforce API is invoked to deploy, retrieve, create, update, or delete customization information, such as custom object definitions using Mule Salesforce Connectors in a Mule application?

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

Answer: C

NEW QUESTION 299
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 303
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 306
An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments.
What must the testing team do before they can start integration testing the API in the Staging environment?

A. They must access the API portal and create an API notebook using the Client ID and Client Secret supplied by the API portal in the Staging environment
B. They must request access to the API instance in the Staging environment and obtain a Client ID and Client Secret to be used for testing the API
C. They must be assigned as an API version owner of the API in the Staging environment
D. They must request access to the Staging environment and obtain the Client ID and Client Secret for that environment to be used for testing the API

Answer: B

NEW QUESTION 310
Refer to the exhibit.

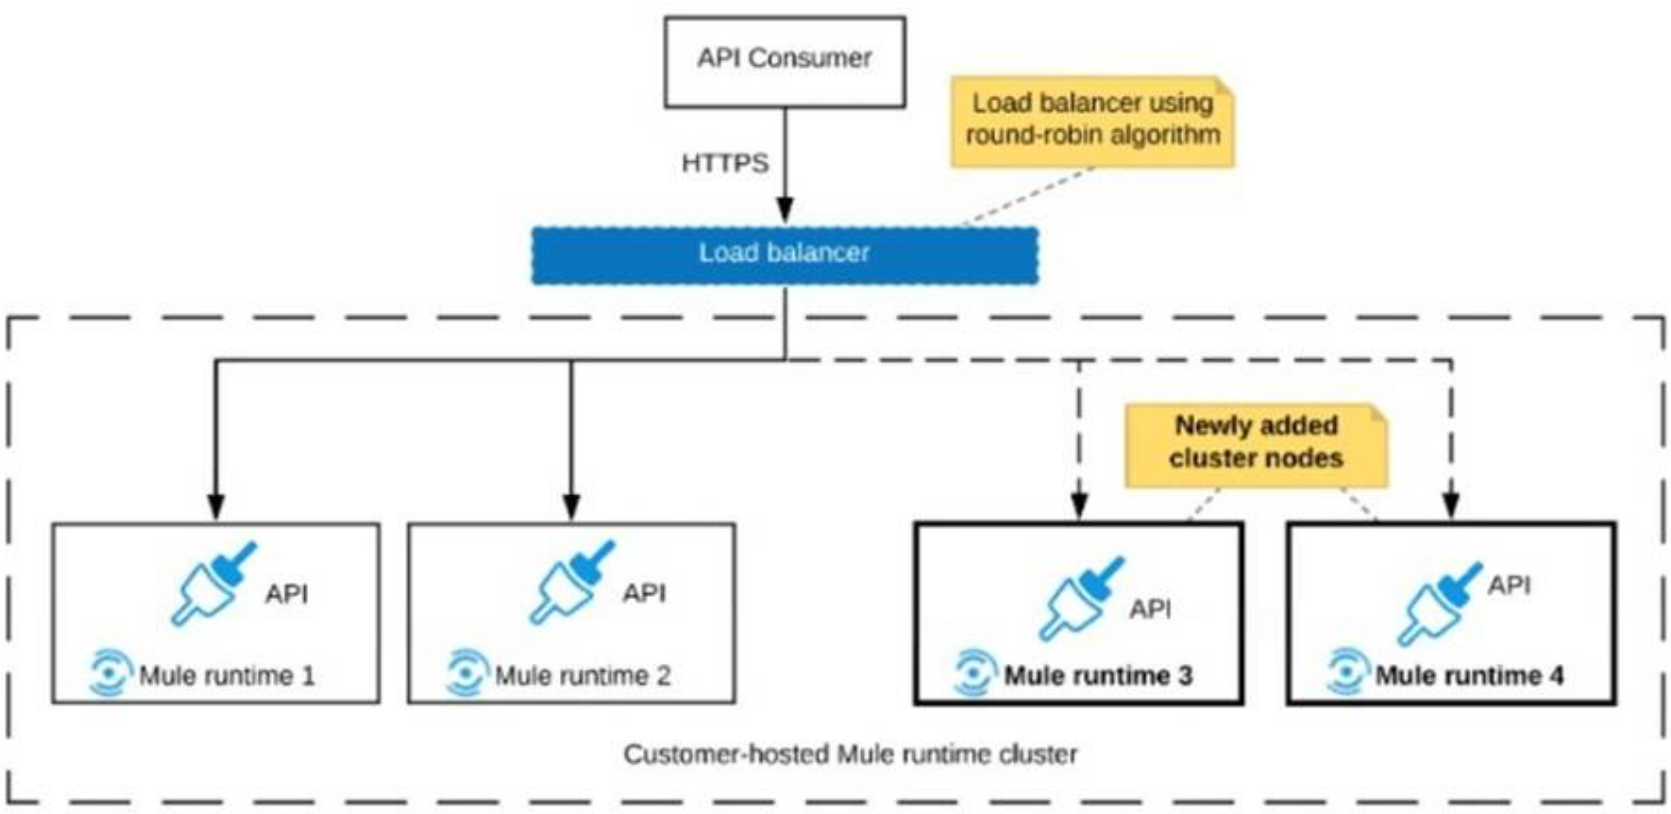

An organization uses a 2-node Mute runtime cluster to host one stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution.
Two additional nodes have been added to the cluster and the load balancer has been configured to recognize the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

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

Answer: D

NEW QUESTION 313
The AnyAirline organization's passenger reservations center is designing an integration solution that combines invocations of three different System APIs (bookFlight, bookHotel, and bookCar) in a business transaction. Each System API makes calls to a single database.
The entire business transaction must be rolled back when at least one of the APIs fails. What is the most idiomatic (used for its intended purpose) way to integrate these APIs in

near real-time that provides the best balance of consistency, performance, and reliability?

A. Implement eXtended Architecture (XA) transactions between the API implementations Coordinate between the API implementations using a Saga patternImplement caching in each API implementation to improve performance
B. Implement local transactions within each API implementationConfigure each API implementation to also participate in the same eXtended Architecture (XA) transactionImplement caching in each API implementation to improve performance
C. Implement local transactions in each API implementation Coordinate between the API implementations using a Saga patternApply various compensating actions depending on where a failure occurs
D. Implement an eXtended Architecture (XA) transaction manager in a Mule application using a Saga patternConnect each API implementation with the Mule application using XA transactionsApply various compensating actions depending on where a failure occurs

Answer: C

NEW QUESTION 318
An automation engineer needs to write scripts to automate the steps of the API lifecycle, including steps to create, publish, deploy and manage APIs and their implementations in Anypoint Platform.
What Anypoint Platform feature can be used to automate the execution of all these actions in scripts in the easiest way without needing to directly invoke the Anypoint Platform REST APIs?

A. Automated Policies in API Manager
B. Runtime Manager agent
C. The Mule Maven Plugin
D. Anypoint CLI

Answer: D

NEW QUESTION 322
What aspects of a CI/CD pipeline for Mule applications can be automated using MuleSoft- provided Maven plugins?

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

Answer: A

NEW QUESTION 323
Refer to the exhibit.

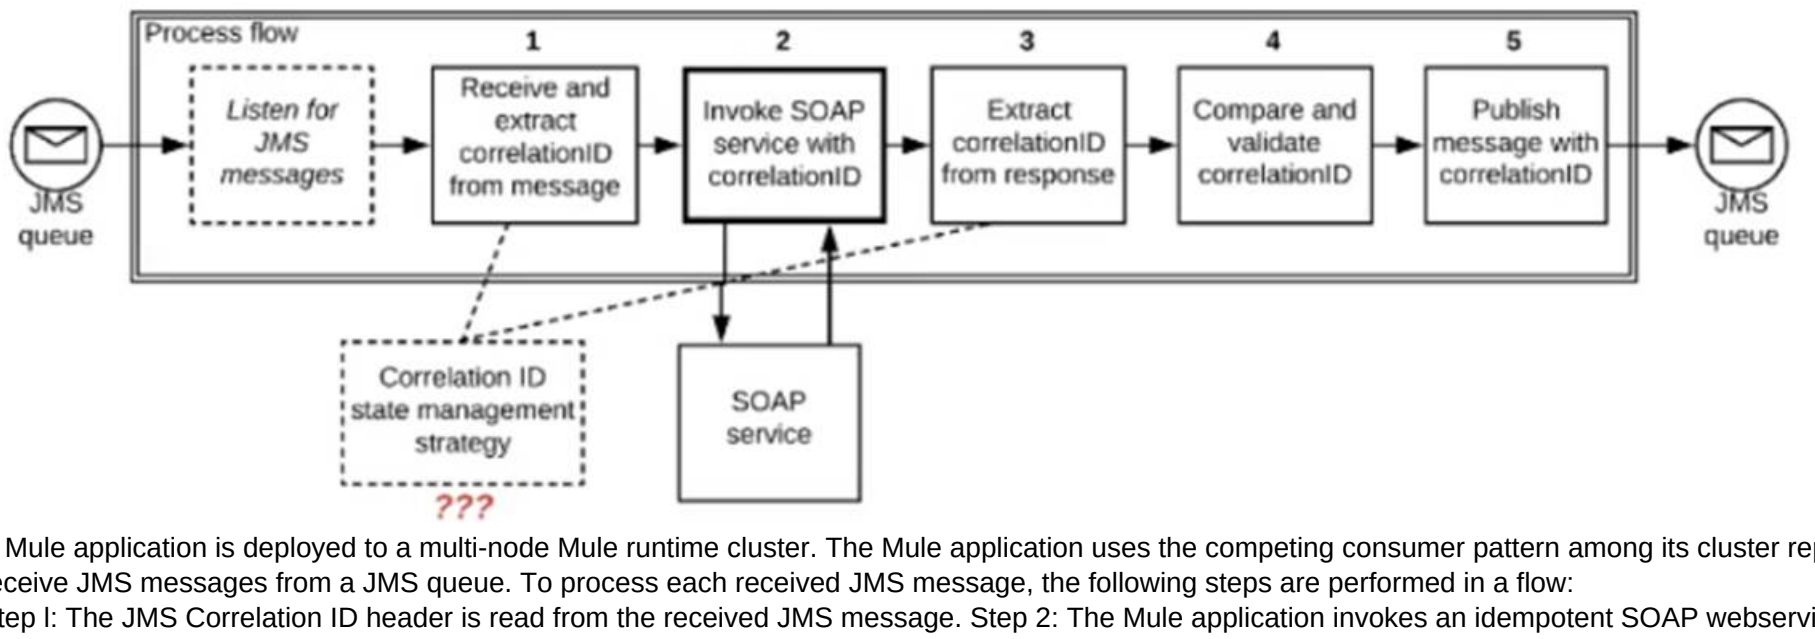

A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to receive JMS messages from a JMS queue. To process each received JMS message, the following steps are performed in a flow:
Step l: The JMS Correlation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

Answer: C

NEW QUESTION 324
Refer to the exhibit.

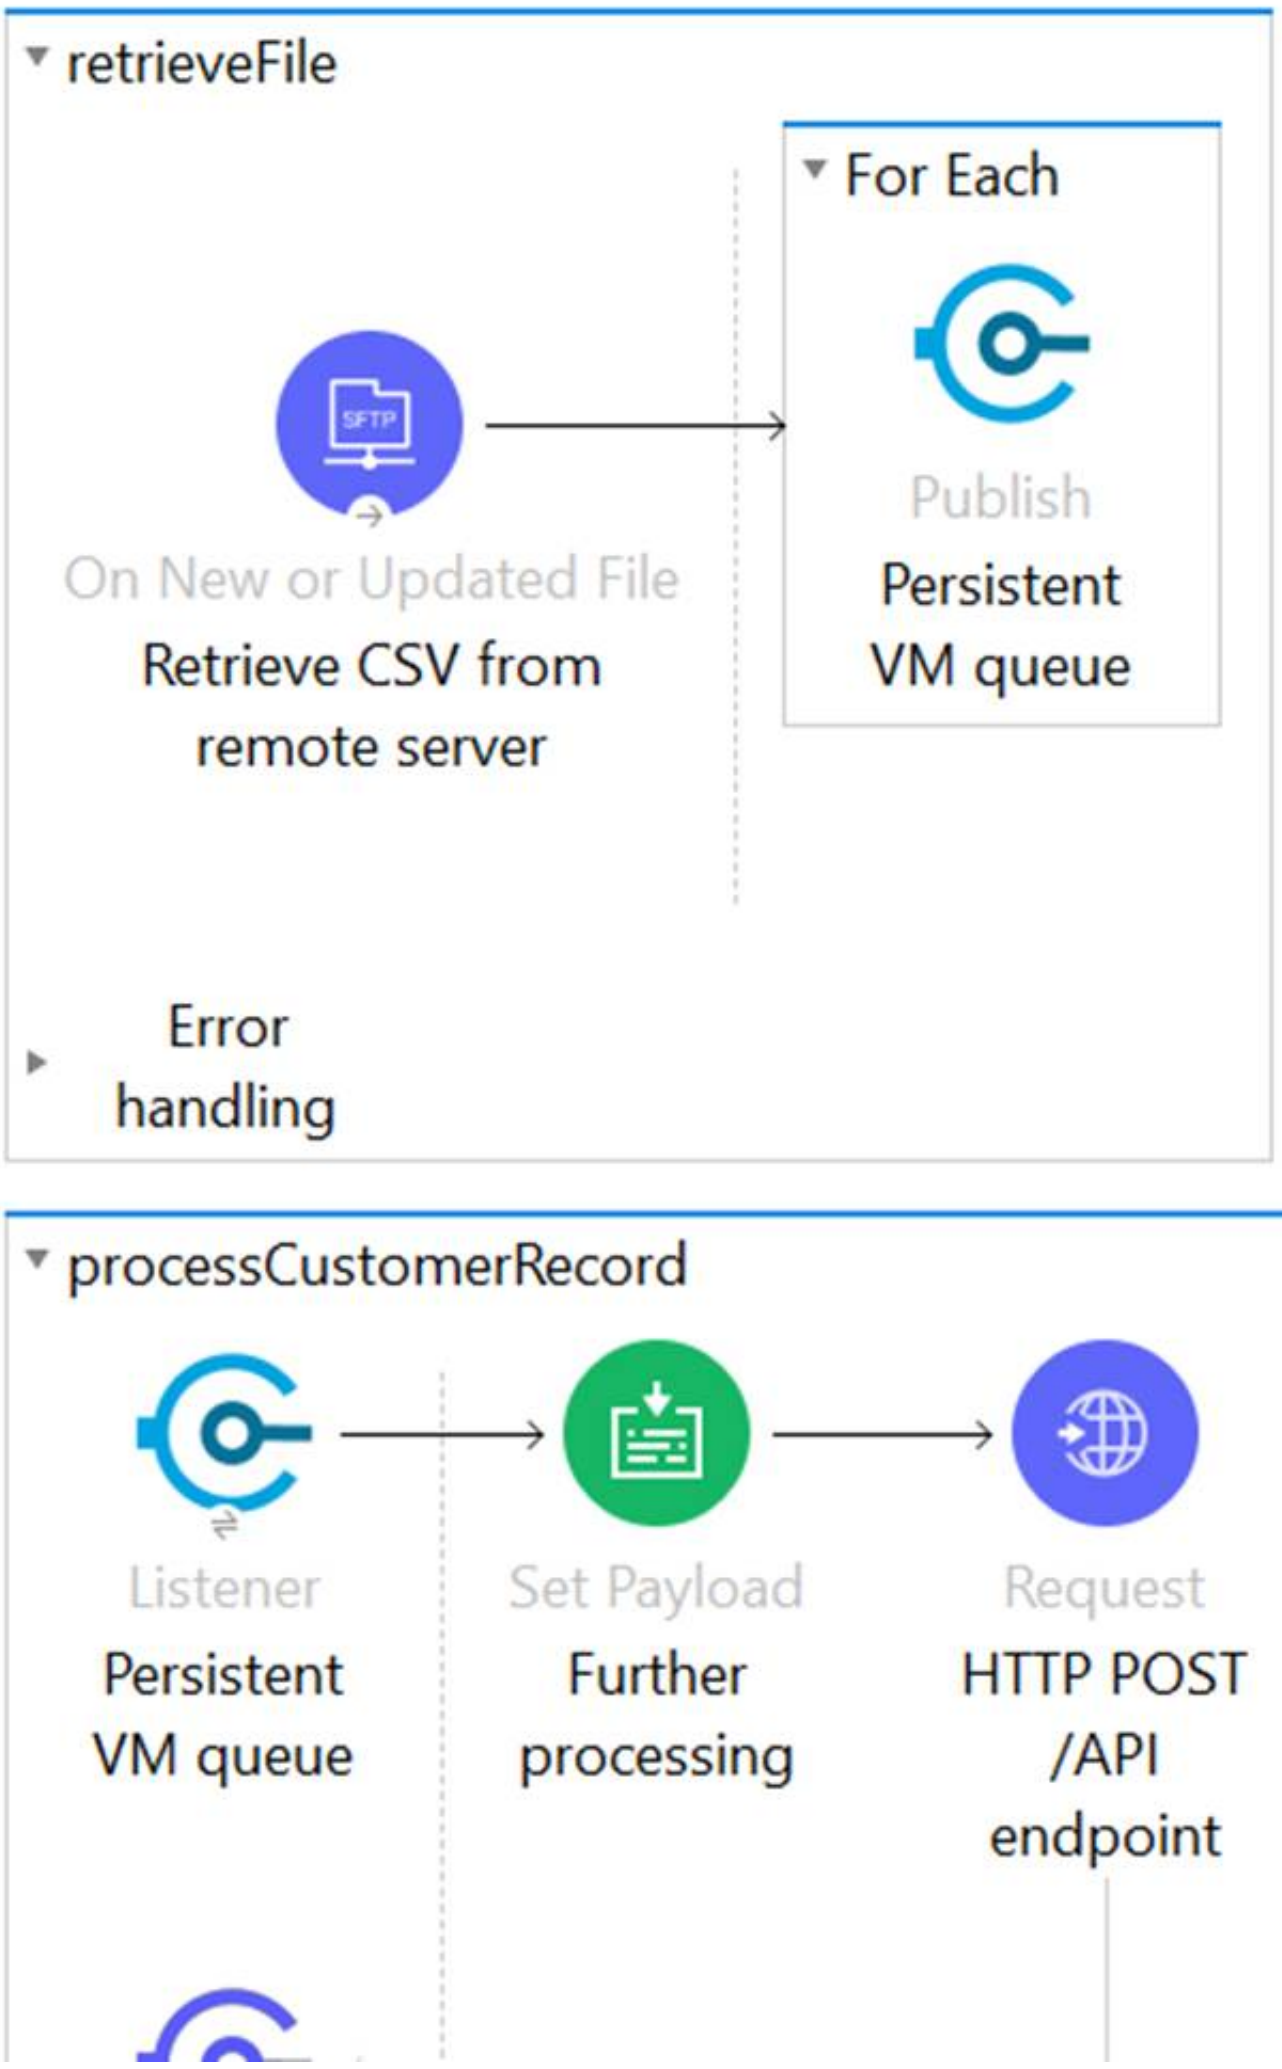

This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow??s VM Listner receives messages from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

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

Answer: B

NEW QUESTION 328
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 330
A developer is examining the responses from a RESTful web service that is compliant with the Mypertext Transfer Protocol (HTTP/1.1) a8 defined by the Internet Engineering Task Force (IETF).
In this HTTP/1.1-compliant web service, which class of HTTP response status codes should be specified to indicate when client requests are successfully received, understood, and accepted by the web service?

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

Answer: B

NEW QUESTION 331
Refer to the exhibit.

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

What is the type data format shown in the exhibit?

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

Answer: C

NEW QUESTION 332
Refer to the exhibit.

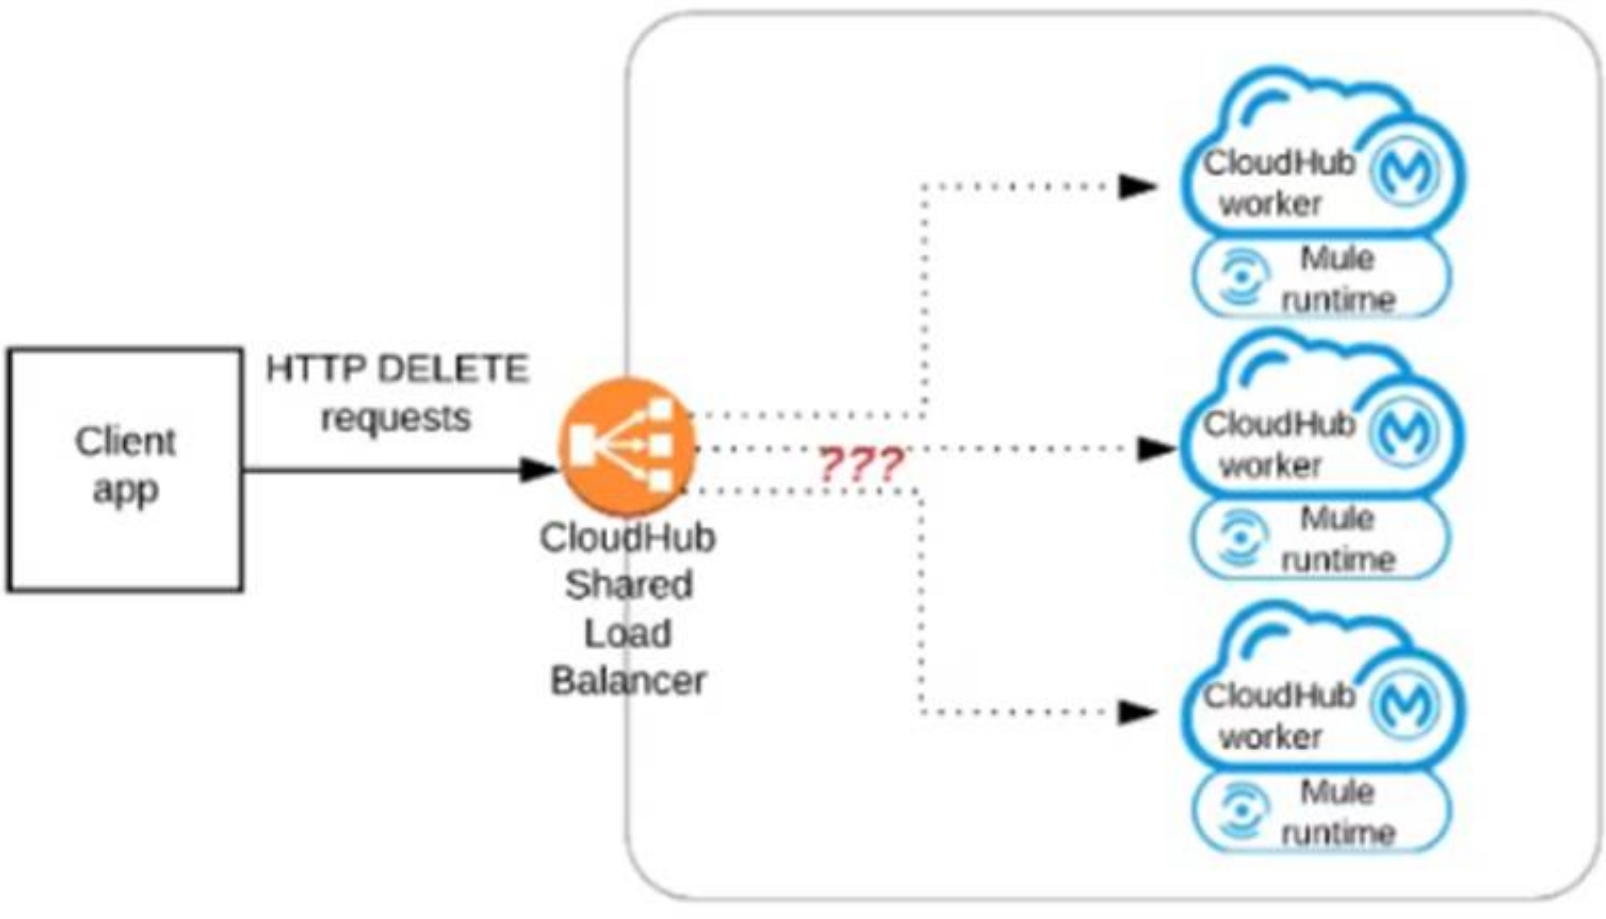

A Mule application has an HTTP Listener that accepts HTTP DELETE requests. This Mule application Is deployed to three CloudHub workers under the control of the CloudHub Shared Load Balancer.

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 337
A new Mule application under development must implement extensive data transformation logic. Some of the data transformation functionality is already available as external transformation services that are mature and widely used across the organization; the rest is highly specific to the new Mule application.
The organization follows a rigorous testing approach, where every service and application must be extensively acceptance tested before it is allowed to go into production.
What is the best way to implement the data transformation logic for this new Mule application while minimizing the overall testing effort?

A. Implement and expose all transformation logic as mlaoservices using DataWeave, so it can be reused by any application component that needs it, including the new Mule application
B. Implement transformation logic in the new Mute application using DataWeave, replicating the transformation logic of existing transformation services
C. Extend the existing transformation services with new transformation logic and Invoke them from the new Mule application
D. Implement transformation logic in the new Mute application using DataWeave, invoking existing transformation services when possible

Answer: D

NEW QUESTION 339
What metrics about API invocations are available for visualization in custom charts using Anypoint Analytics?

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

Answer: C

NEW QUESTION 340
An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events
What are the two considerations that must be kept in mind while using test recorder (Choose two answers)

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

Answer: AD

NEW QUESTION 344
An organization has implemented a continuous integration (CI) lifecycle that promotes Mule applications through code, build, and test stages. To standardize the organization's CI journey, a new dependency control approach is being designed to store artifacts that include information such as dependencies, versioning, and build promotions.
To implement these process improvements, the organization will now require developers to maintain all dependencies related to Mule application code in a shared location.
What is the most idiomatic (used for its intended purpose) type of system the organization should use in a shared location to standardize all dependencies related to Mule application code?

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

Answer: C

NEW QUESTION 347
What is a key difference between synchronous and asynchronous logging from Mule applications?

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

Answer: B

NEW QUESTION 352
An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub.
For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub??s

default zero- downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers.This is expected to result in several Mule application deployments per hour.

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

Answer: A

NEW QUESTION 353
A REST API is being designed to implement a Mule application.
What standard interface definition language can be used to define REST APIs?

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

Answer: B

NEW QUESTION 356
Anypoint Exchange is required to maintain the source code of some of the assets committed to it, such as Connectors, Templates, and API specifications.
What is the best way to use an organization's source-code management (SCM) system in this context?

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

Answer: B

NEW QUESTION 358
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

MuleSoft-Integration-Architect-I Practice Exam Features:

* MuleSoft-Integration-Architect-I Questions and Answers Updated Frequently

* MuleSoft-Integration-Architect-I Practice Questions Verified by Expert Senior Certified Staff

* MuleSoft-Integration-Architect-I Most Realistic Questions that Guarantee you a Pass on Your FirstTry

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