Salesforce

Exam Questions MuleSoft-Integration-Architect-I

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam

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

A. Anypoint MQ Audit Logs include logs for sending, receiving, or browsing messages
B. Anypoint MQ Audit Logs include 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
What is an advantage that Anypoint Platform offers by providing universal API management and Integration-Platform-as-a-Service (iPaaS) capabilities in a unified platform?

A. Ability to use a single iPaaS to manage and integrate all API gateways
B. Ability to use a single connector to manage and integrate all APis
C. Ability to use a single control plane for both full-lifecycle AP] management and integration
D. Ability to use a single iPaaS to manage all API developer portals

Answer: C

NEW QUESTION 3
Insurance organization is planning to deploy Mule application in MuleSoft Hosted runtime plane. As a part of requirement , application should be scalable . highly available. It also has regulatory requirement which demands logs to be retained for at least 2 years. As an Integration Architect what step you will recommend in order to achieve this?

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

Answer: A

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

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

Answer: D

NEW QUESTION 6
When designing an upstream API and its implementation, the development team has been advised to not set timeouts when invoking downstream API. Because the downstream API has no SLA that can be relied upon. This is the only donwstream API dependency of that upstream API. Assume the downstream API runs uninterrupted without crashing. What is the impact of this advice?

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

Answer: B

NEW QUESTION 7
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 8
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 9
According to MuleSoft's IT delivery and operating model, which approach can an organization adopt in order to reduce the frequency of IT project delivery failures?

A. Decouple central IT projects from the innovation that happens within each line of business
B. Adopt an enterprise data model
C. Prevent technology sprawl by reducing production of API assets
D. Stop scope creep by centralizing requirements-gathering

Answer: A

NEW QUESTION 10
A Mule application is synchronizing customer data between two different database systems.
What is the main benefit of using XA transaction over local transactions to synchronize these two database system?

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

Answer: D

NEW QUESTION 10
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 13
What is true about the network connections when a Mule application uses a JMS connector to interact with a JMS provider (message broker)?

A. To complete sending a JMS message, the JMS connector must establish a network connection with the JMS message recipient
B. To receive messages into the Mule application, the JMS provider initiates a network connection to the JMS connector and pushes messages along this connection
C. The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider
D. The AMQP protocol can be used by the JMS connector to portably establish connections to various types of JMS providers

Answer: C

NEW QUESTION 18
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 23
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 26
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.

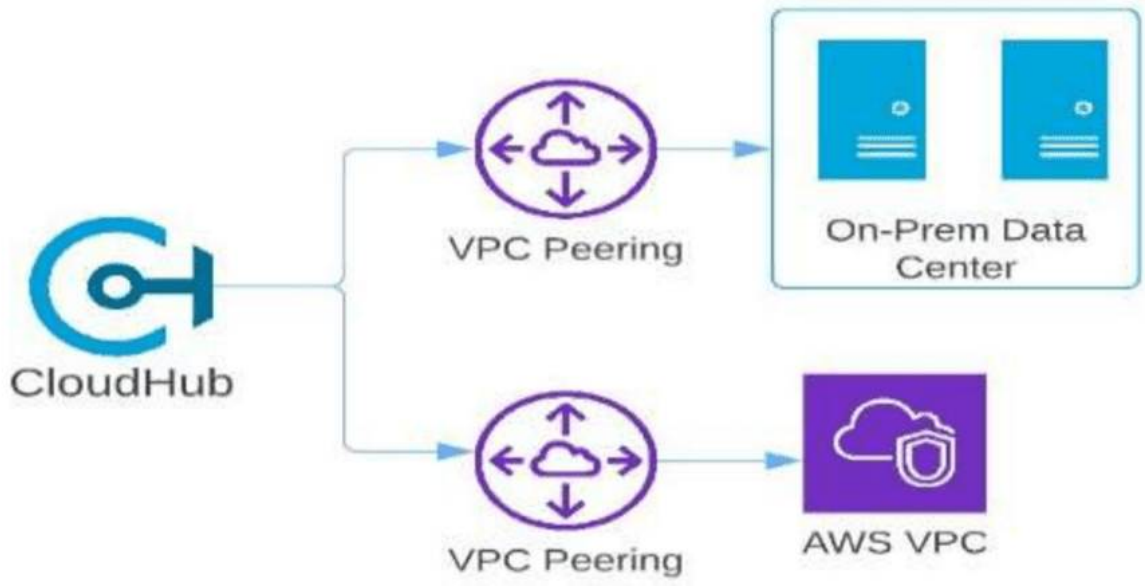

B) In the AWS account, attach the CloudHub 2.0 private space to an AWS transit gateway that routes from the CloudHub 2.0 private space to the on-premises data cente
B. InRuntime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the AWS VPC.

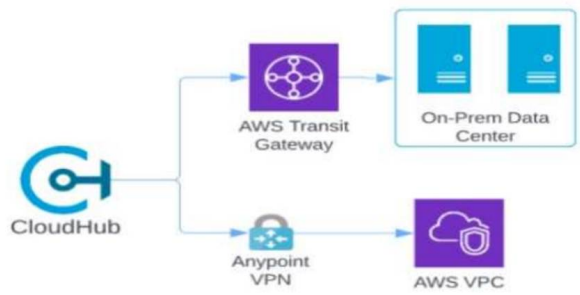

C) In Runtime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the on-premises data cente
C. In the AWS account, attach the private..

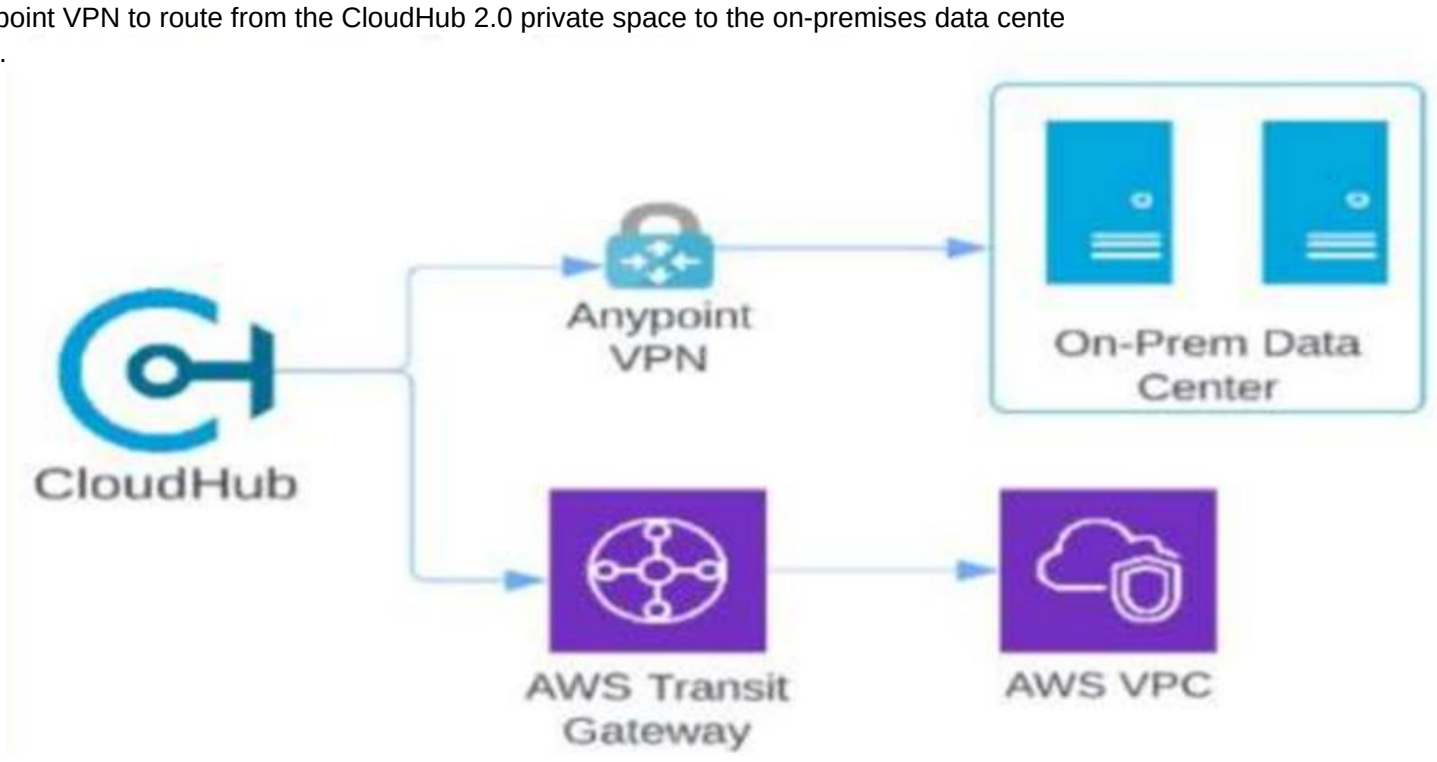

D) In the AWS account, attach the private space directly to the AWS VPC, In the AWS account, use an AWS transit gateway to route from the AWS VPC to the on-premises data center.

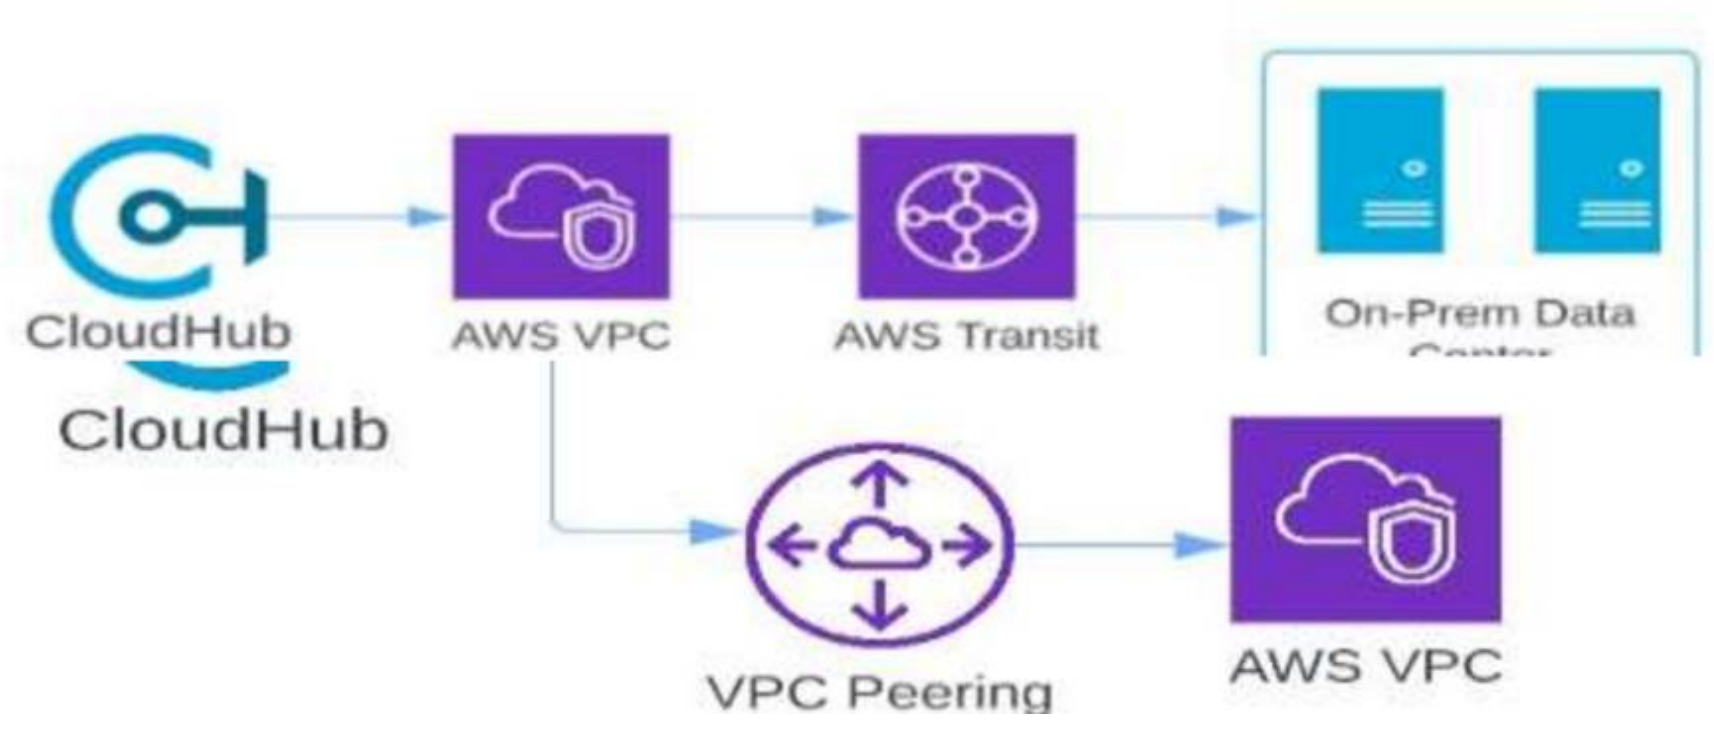

D. Option A
E. Option B
F. Option C
G. Option D

Answer: B

NEW QUESTION 30
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 35
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 40
Refer to the exhibit.

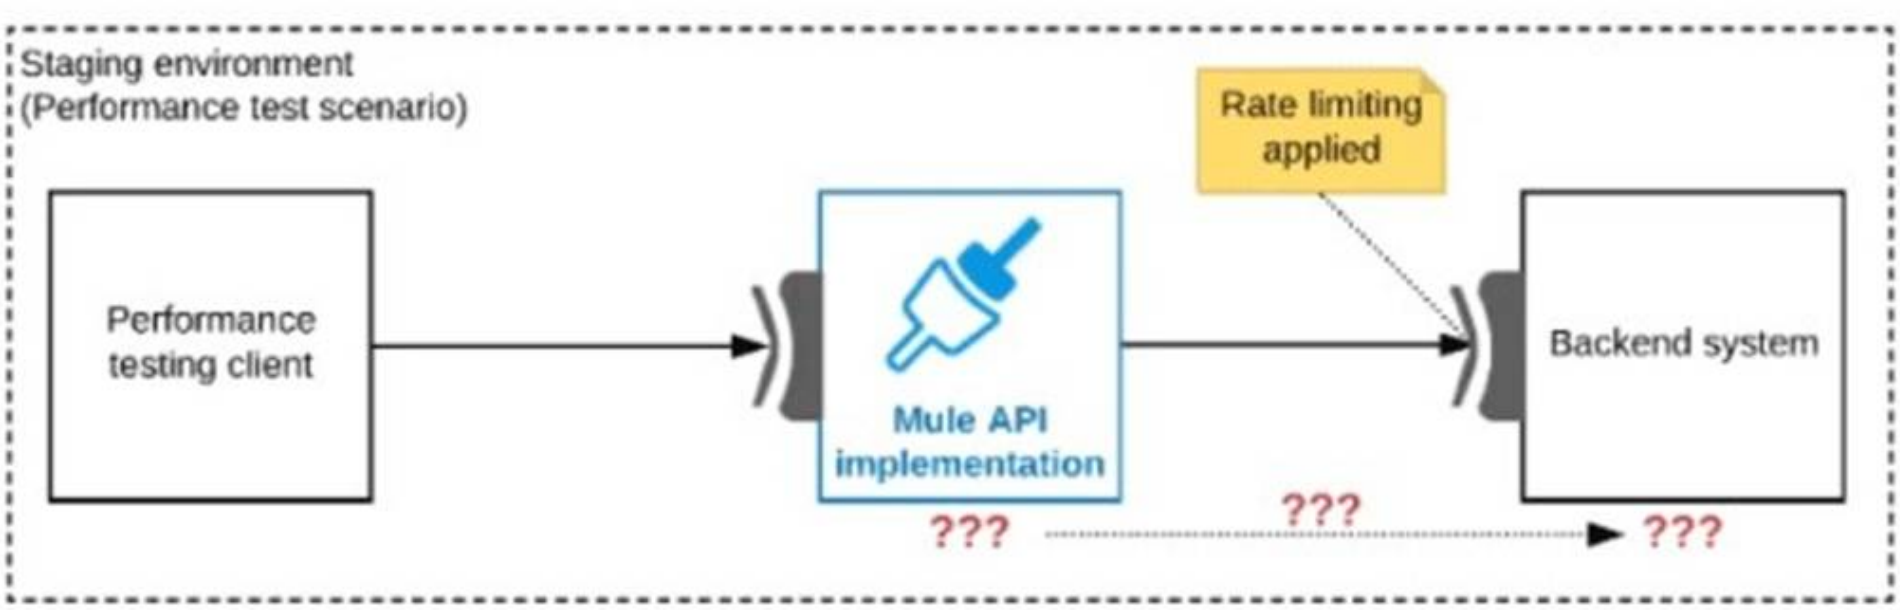

One of the backend systems invoked by an API implementation enforces rate limits on the
number of requests a particular client can make. Both the backend system and the API implementation are deployed to several non-production environments in addition to production.
Rate limiting of the backend system 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 44
An organization has strict unit test requirement that mandate every mule application must have an MUnit test suit with a test case defined for each flow and a minimum test coverage of 80%.
A developer is building Munit test suit for a newly developed mule application that sends API request to an external rest API.
What is the effective approach for successfully executing the Munit tests of this new application while still achieving the required test coverage for the Munit tests?

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

Answer: C

NEW QUESTION 49
When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

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

Answer: B

NEW QUESTION 51
Refer to the exhibit.

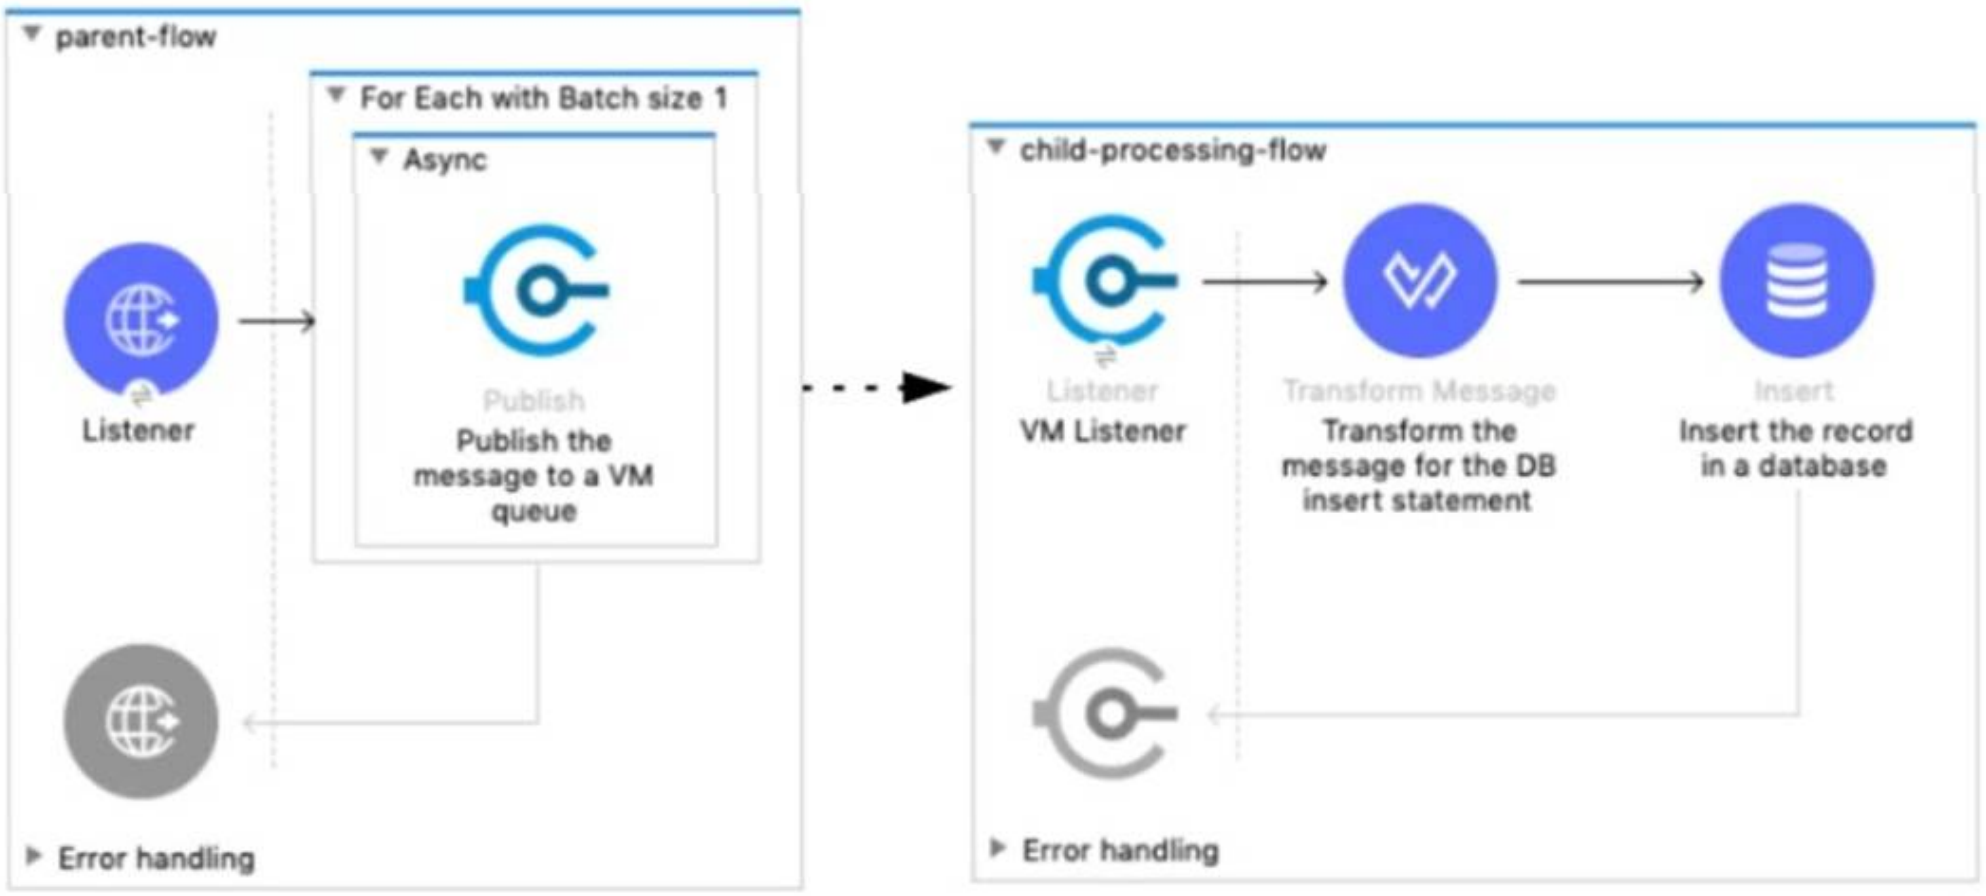

A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items, then sends each item one at a time inside an Async scope to a VM queue.
A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing the item to a database.
This Mule application is deployed to four CloudHub workers with persistent queues enabled.
What message processing guarantees are provided by the VM queue and the CloudHub workers, and how are VM messages routed among the CloudHub workers for each invocation of the parent flow under normal operating conditions where all the CloudHub workers remain online?

A. EACH item VM message is processed AT MOST ONCE by ONE CloudHub worker, with workers chosen in a deterministic round-robin fashion Each of the four CloudHub workers can be expected to process 1/4 of the Item VM messages (about 50 items)
B. EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker Each of the four CloudHub workers can be expected to process some item VM messages
C. ALL Item VM messages are processed AT LEAST ONCE by the SAME CloudHub worker where the parent flow was invokedThis one CloudHub worker processes ALL 200 item VM messages
D. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub workerThis one CloudHub worker processes ALL 200 item VM messages

Answer: B

NEW QUESTION 55
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 56
Which Anypoint Platform component should a MuleSoft developer use to create an API specification prior to building the API implementation?

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

Answer: B

NEW QUESTION 60
An external API frequently invokes an Employees System API to fetch employee data from a MySQL database. The architect must design a caching strategy to query the database only when there Is an update to the Employees table or else return a cached response in order to minimize the number of redundant transactions being handled by the database.

A. Use an On Table Row operation configured with the Employees table, call invalidate cache, and hardcode the new Employees data to cach
B. Use an object-store-caching- strategy and set the expiration interval to 1 hour.
C. Use an On Table Row operation configured with the Employees table and cail invalidate cach
D. Use an object-store-caching-strategy and the default expiration interval.
E. Use a Scheduler with a fixed frequency set to every hour to trigger an invalidate cache flo
F. Use an object-store-caching-strategy and the default expiration interval.
G. Use a Scheduler with a fixed frequency set to every hour, triggering an invalidate cache flo
H. Use an object-store-caching-strategy and set the expiration interval to 1 hour.

Answer: B

NEW QUESTION 64
An organization is designing multiple new applications to run on CloudHub in a single Anypoint VPC and that must share data using a common persistent Anypoint object store V2 (OSv2).
Which design gives these mule applications access to the same object store instance?

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

Answer: C

NEW QUESTION 68
In a Mule Application, a flow contains two (2) JMS consume operations that are used to connect to a JMS broker and consume messages from two(2) JMS destination. The Mule application then joins the two JMS messages together.
The JMS broker does not implement high availability (HA) and periodically experiences scheduled outages of upto 10 mins for routine maintenance.
What is the most idiomatic (used for its intented purpose) way to build the mule flow so it can best recover from the expected outages?

A. Configure a reconnection strategy for the JMS connector
B. Enclose the two(2) JMS operation in an Until Successful scope
C. Consider a transaction for the JMS connector
D. Enclose the two(2) JMS operations in a Try scope with an Error Continue error handler

Answer: A

NEW QUESTION 71
What Mule application can have API policies applied by Anypoint Platform to the endpoint exposed by that Mule application?

A. A Mule application that accepts requests over HTTP/1x
B. A Mule application that accepts JSON requests over TCP but is NOT required to provide a response.
C. A Mule application that accepts JSON requests over WebSocket
D. A Mule application that accepts gRPC requests over HTTP/2

Answer: A

NEW QUESTION 72
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 74
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 75
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 80
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 84
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 85
An organization is implementing a Quote of the Day API that caches today's quote. What scenario can use the CloudHub Object Store connector to persist the cache's state?

A. When there is one deployment of the API implementation to CloudHub and another one to customer hosted mule runtime that must share the cache state.
B. When there are two CloudHub deployments of the API implementation by two Anypoint Platform business groups to the same CloudHub region that must share the cache state.
C. When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.
D. When there are three CloudHub deployments of the API implementation to three separate CloudHub regions that must share the cache state.

Answer: C

NEW QUESTION 87
A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.
The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.
What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

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

Answer: D

NEW QUESTION 91
A DevOps team has adequate observability of individual system behavior and performance, but it struggles to track the entire lifecycle of each request across different microservices.
Which additional observability approach should this team consider adopting?

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

Answer: C

NEW QUESTION 93
A bank is implementing a REST API in a Mule application to receive an array of accounts from an online banking platform user interface (UI),retrieve account balances for those accounts from a backend Finance system, and then return the account balances so they can be displayed inthe online banking platform UI. As part of the processing, the MuleSoft API also needs to insert the retrieved account data into an AuditDatabase for auditing purposes. the auditing process should not add latency to the account balance retrieval response back to the onlinebanking platform UI.
The retrieveBalances flow in the Mule application is designed to use an operation in a connector to the Finance system (the Finance operation) that can only look up one account record at a time, and a operation from a different connector to the Audit system (the Audit operation) that can only insert one account record at a time.
To best meet the performance-related requirements, what scope or scopes should be used and how should they be used to incorporate the Finance operation and Audit operation into the retrieveBalances flow?

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

Answer: A

NEW QUESTION 98
A finance giant is planning to migrate all its Mule applications to Runtime fabric (RTF). Currently all Mule applications are deployed cloud hub using automated CI/CD scripts.
As an integration architect, which of the below step would you suggest to ensure that the applications from cloudhub are migrated properly to Runtime Fabric (RTF) with an assumption that organization is keen on keeping the same deployment strategy.

A. No changes need to be made to POM.xml file and CI/CD script should be modified as per the RTF configurations
B. runtimeFabric dependency should be added as a mule plug-in to POM.xml file and CI/CD script should be modified as per the RTF configurations
C. runtimeFabric deployment should be added to POM.xml file in all the mule applications and CI/CD script should be modified as per the RTF configurations
D. runtimeFabric profile should be added mule configuration files in the mule applications and CI/CD script should be modified as per the RTF configurations

Answer: C

NEW QUESTION 103
An organization plans to migrate all its Mule applications to Runtime Fabric (RTF). Currently, all Mule applications have been deployed to CloudHub using automated CI/CD scripts.
What steps should be taken to properly migrate the applications from CloudHub to RTF, while keeping the same automated CI/CD deployment strategy?

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

Answer: D

NEW QUESTION 107

Which productivity advantage does Anypoint Platform have to both implement and manage an AP?

A. Automatic API proxy generation
B. Automatic API specification generation
C. Automatic API semantic versioning
D. Automatic API governance

Answer: A

NEW QUESTION 109
A Mule application is synchronizing customer data between two different database systems.
What is the main benefit of using eXtended Architecture (XA) transactions over local transactions to synchronize these two different database systems?

A. An XA transaction synchronizes the database systems with the least amount of Mule configuration or coding
B. An XA transaction handles the largest number of requests in the shortest time
C. An XA transaction automatically rolls back operations against both database systems if any operation falls
D. An XA transaction writes to both database systems as fast as possible

Answer: B

NEW QUESTION 110
An organization has implemented the cluster with two customer hosted Mule runtimes is hosting an application.
This application has a flow with a JMS listener configured to consume messages from a queue destination. As an integration architect can you advise which JMS listener configuration must be used to receive messages in all the nodes of the cluster?

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

Answer: B

NEW QUESTION 112
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 113
According to MuleSoft, which deployment characteristic applies to a microservices application architecture?

A. Services exist as independent deployment artifacts and can be scaled -independently of other services
B. All services of an application can be deployed together as single Java WAR file
C. A deployment to enhance one capability requires a redeployment of all capabilities
D. Core business capabilities are encapsulated in a single, deployable application

Answer: A

NEW QUESTION 114
As part of a growth strategy, a supplier signs a trading agreement with a large customer. The customer sends purchase orders to the supplier according to the ANSI X12 EDI standard, and the supplier creates the orders in its ERP system using the information in the EDI document.
The agreement also requires that the supplier provide a new RESTful API to process request from the customer for current product inventory level from the supplier?? s ERP system.
Which two fundamental integration use cases does the supplier need to deliver to provide an end-to-end solution for this business scenario? (Choose two.)

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

Answer: AB

NEW QUESTION 115
An organization will deploy Mule applications to Cloudhub, Business requirements mandate that all application logs be stored ONLY in an external splunk consolidated logging service and NOT in Cloudhub.
In order to most easily store Mule application logs ONLY in Splunk, how must Mule application logging be configured in Runtime Manager, and where should the log4j2 splunk appender be defined?

A. Keep the default logging configuration in RuntimeManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support at Mule application deployments.
B. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in EACH Mule application??s log4j2.xml file
C. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manger to support at Mule application deployments.
D. Keep the default logging configuration in Runtime Manager Define the Splunk appender in EACH Mule application log4j2.xml file

Answer: B

NEW QUESTION 116
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 120
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 121
An IT integration delivery team begins a project by gathering all of the requirements, and proceeds to execute the remaining project activities as sequential, non-repeating phases.
Which IT project delivery methodology is this team following?

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

Answer: C

NEW QUESTION 122
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 125
A new Mule application has been deployed through Runtime Manager to CloudHub 1.0 using a CI/CD pipeline with sensitive properties set as cleartext. The Runtime Manager Administrator opened a high priority incident ticket about this violation of their security requirements indicating these sensitive properties values must not be stored or visible in Runtime Manager but should be changeable in Runtime Manager by Administrators with proper permissions.
How can the Mule application be deployed while safely hiding the sensitive properties?

A. Add an ArrayList of all the sensitive properties?? names in the mule-artifact.json file of the application
B. Add encrypted versions of the sensitive properties as global configuration properties in the Mule application
C. Add a new wrapper.java.additional.xx parameter for each sensitive property in the wrapper.conf file used by the CI/CD pipeline scripts
D. Create a variable for each sensitive property and declare them as hidden in the CI/CD pipeline scripts

Answer: B

NEW QUESTION 130
Refer to the exhibit.

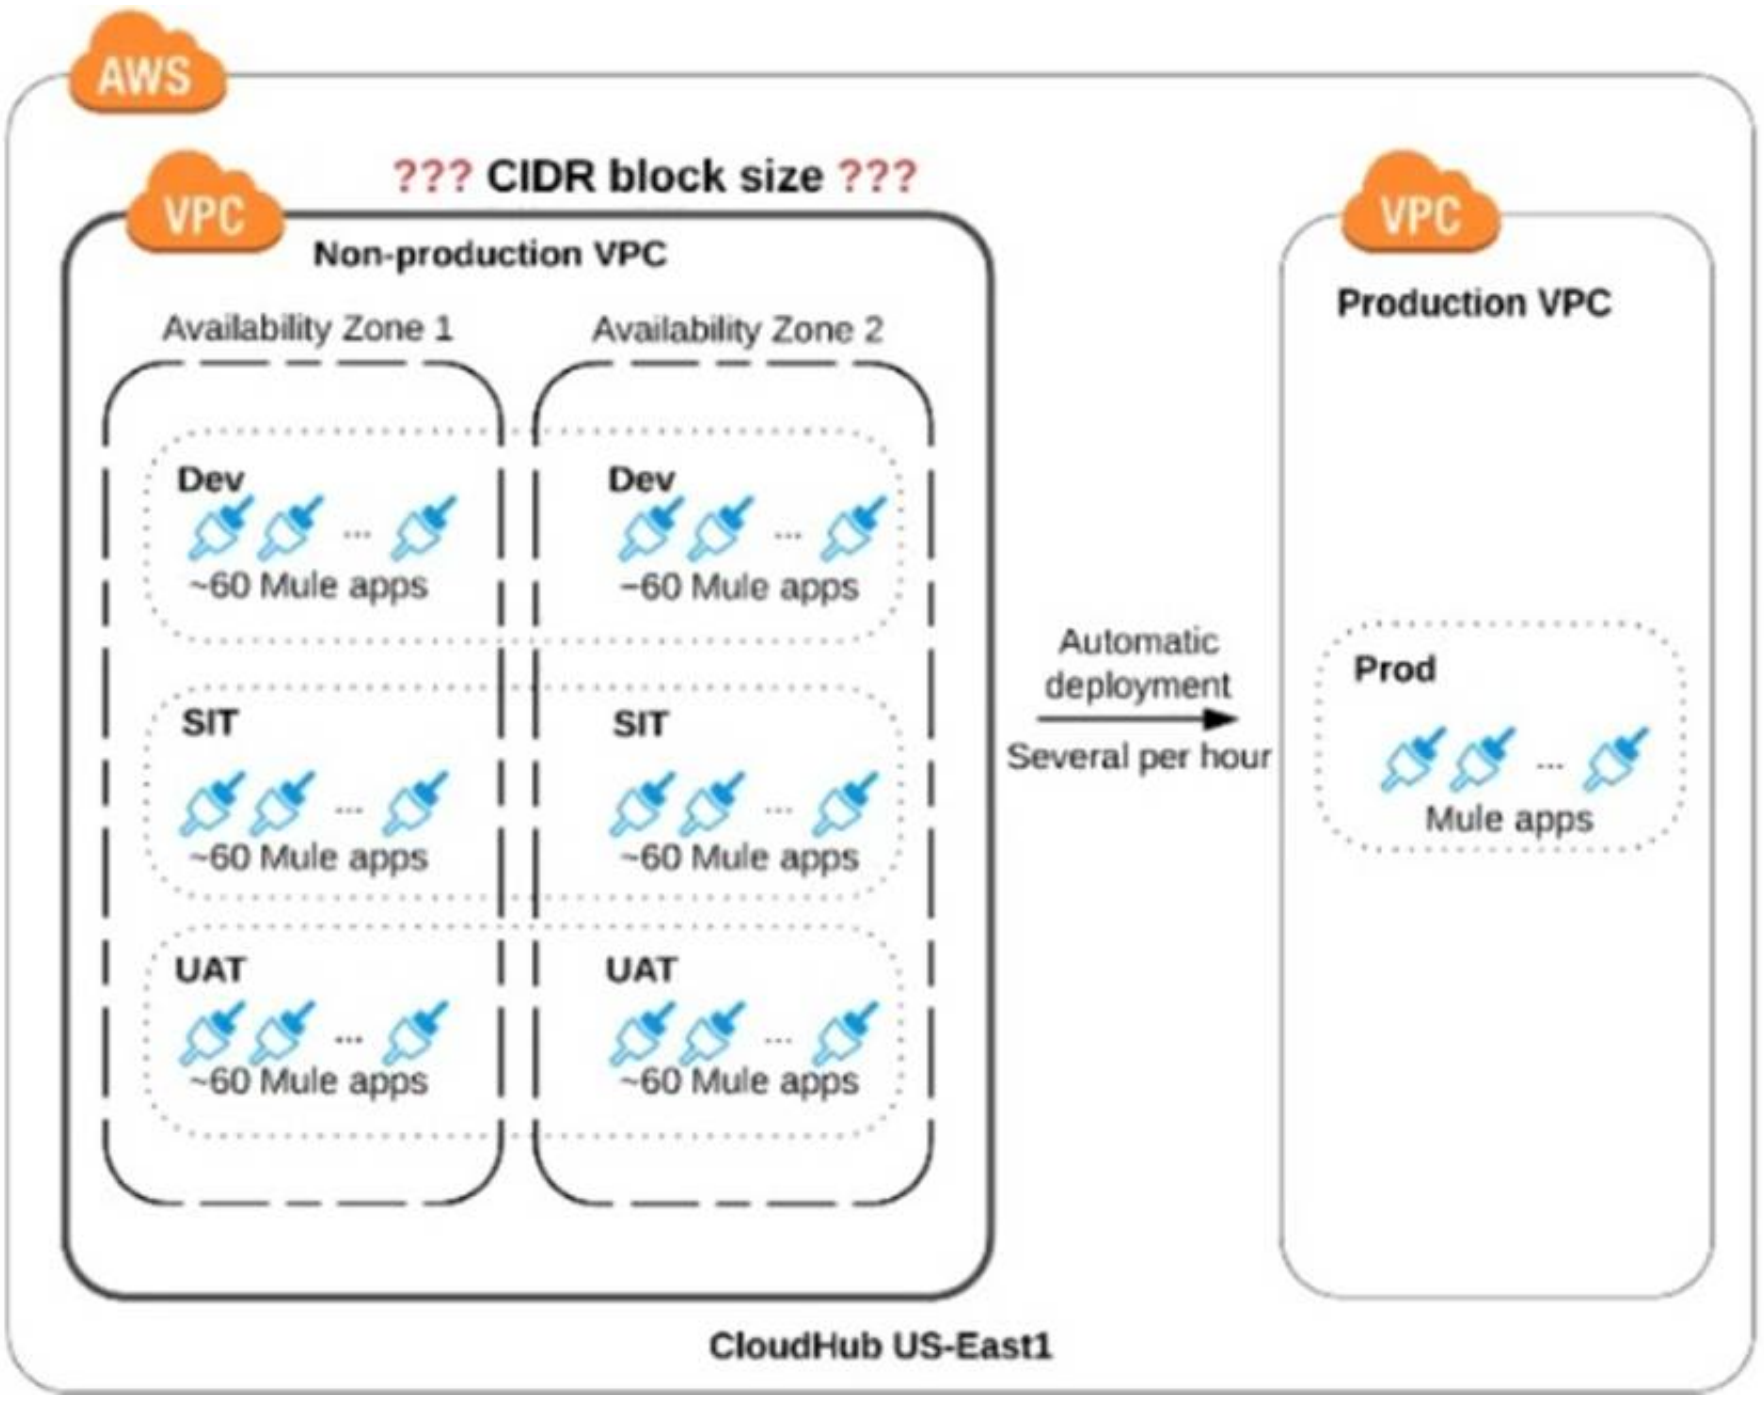

An organization is sizing an Anypoint VPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization currently has three non-production environments (DEV, SIT and UAT) that share this VPC. The AWS region of the VPC has two AZs.
The organization has a very mature DevOps approach which automatically progresses each application through all non-production environments before automatically deploying to production. This process results in several Mule application deployments per hour, using CloudHub's normal zero-downtime deployment feature.
What is a CIDR block for this VPC that results in the smallest usable private IP address range?

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

Answer: D

NEW QUESTION 134
Why would an Enterprise Architect use a single enterprise-wide canonical data model (CDM) when designing an integration solution using Anypoint Platform?

A. To reduce dependencies when integrating multiple systems that use different data formats
B. To automate AI-enabled API implementation generation based on normalized backend databases from separate vendors
C. To leverage a data abstraction layer that shields existing Mule applications from nonbackward compatible changes to the model's data structure
D. To remove the need to perform data transformation when processing message payloads in Mule applications

Answer: A

NEW QUESTION 135
Refer to the exhibit.

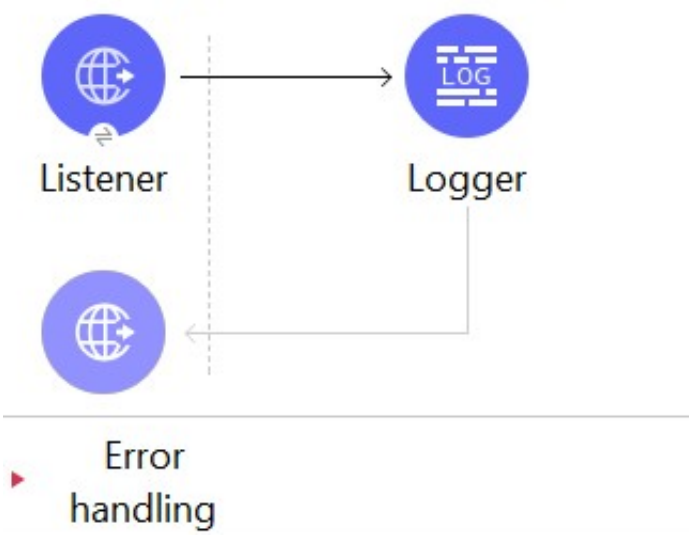

The HTTP Listener and the Logger are being handled from which thread pools respectively?

A. CPU_INTENSIVE and Dedicated Selector pool
B. UBER and NONBLOCKING
C. Shared Selector Pool and CPU LITE
D. BLOCKING _IO and UBER

Answer: C

NEW QUESTION 138
A Mule application contains a Batch Job scope with several Batch Step scopes. The Batch Job scope is configured with a batch block size of 25.
A payload with 4,000 records is received by the Batch Job scope.
When there are no errors, how does the Batch Job scope process records within and between the Batch Step scopes?

A. The Batch Job scope processes multiple record blocks in parallel, and a block of 25 records can jump ahead to the next Batch Step scope over an earlier block of records Each Batch Step scope is invoked with one record in the payload of the received Mule eventFor each Batch Step scope, all 25 records within a block are processed in parallelAll the records in a block must be completed before the block of 25 records is available to the next Batch Step scope
B. The Batch Job scope processes each record block sequentially, one at a time Each Batch Step scope is invoked with one record in the payload of the received Mule eventFor each Batch Step scope, all 25 records within a block are processed sequentially, one at a timeAll 4000 records must be completed before the blocks of records are available to the next Batch Step scope
C. The Batch Job scope processes multiple record blocks in parallel, and a block of 25 records can jump ahead to the next Batch Step scope over an earlier block of records Each Batch Step scope is invoked with one record in the payload of the received Mule eventFor each Batch Step scope, all 25 records within a block are processed sequentially, one record at a timeAll the records in a block must be completed before the block of 25 records is available tothe next Batch Step scope
D. The Batch Job scope processes multiple record blocks in parallelEach Batch Step scope is invoked with a batch of 25 records in the payload of the received Mule eventFor each Batch Step scope, all 4000 records are processed in parallelIndividual records can jump ahead to the next Batch Step scope before the rest of the records finish processing in the current Batch Step scope

Answer: A

NEW QUESTION 141
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 144
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 147
What requirement prevents using Anypoint MQ as the messaging broker for a Mule application?

A. When the payload sent through the message broker must use XML format
B. When the payload sent through the message broker must be encrypted
C. When the messaging broker must support point-to-point messaging

D. When the messaging broker must be deployed on-premises

Answer: D

NEW QUESTION 152
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 153
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 154
Refer to the exhibit.

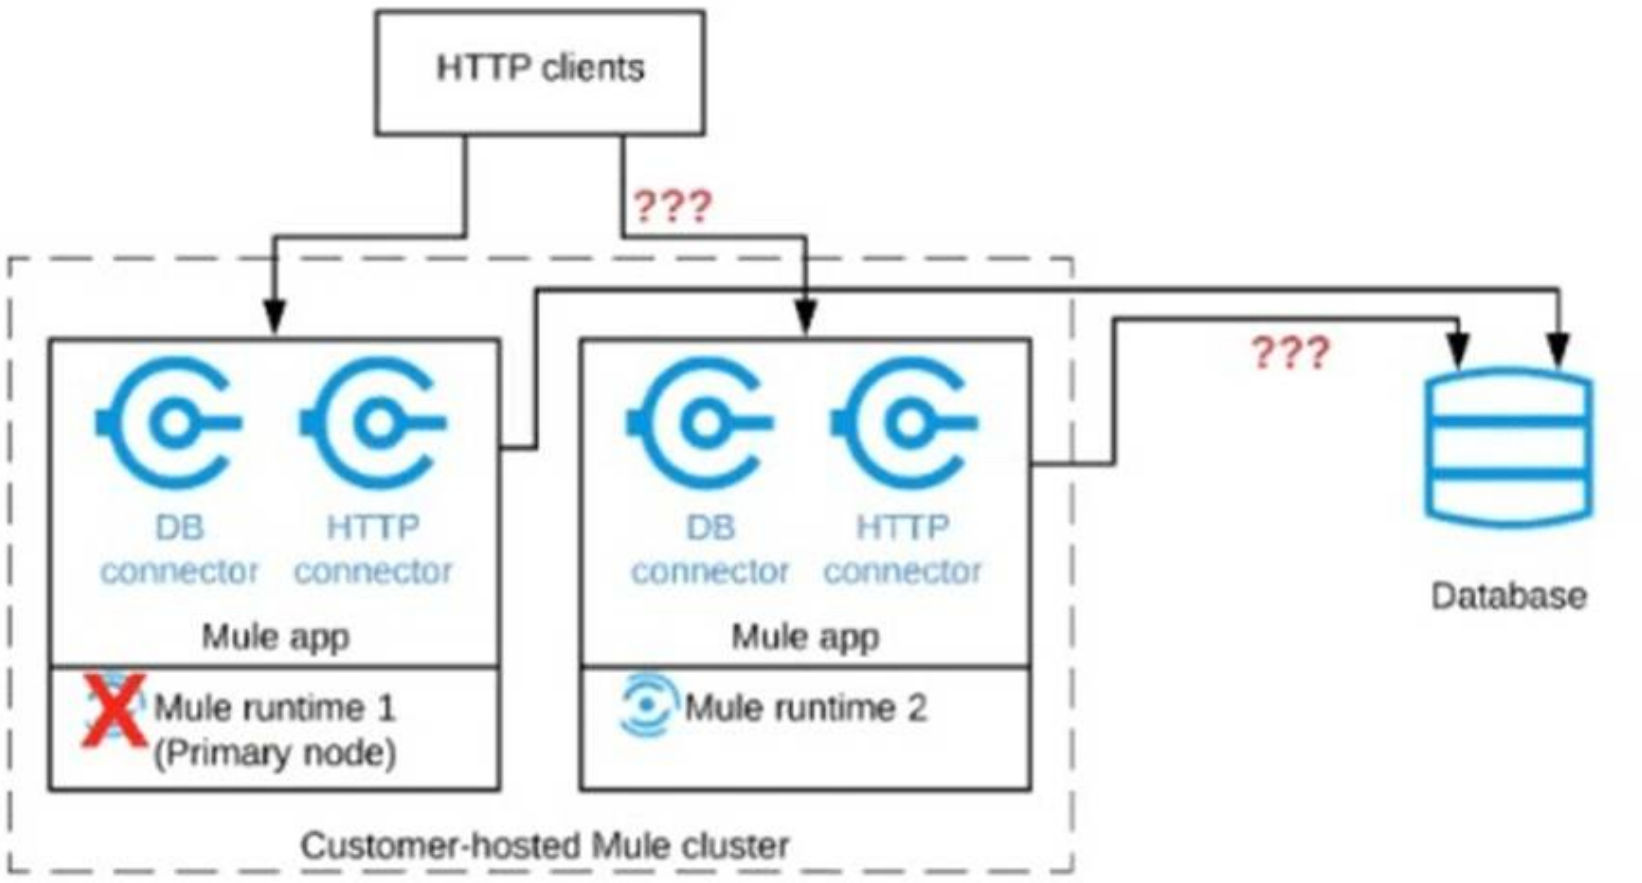

A Mule application is deployed to a cluster of two customer-hosted Mute runtimes. The Mute application has a flow that polls a database and another flow with an HTTP Listener.
HTTP clients send HTTP requests directly to individual cluster nodes.
What happens to database polling and HTTP request handling 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 157
A retailer is designing a data exchange interface to be used by its suppliers. The interface must support secure communication over the public internet. The interface must also work with a wide variety of programming languages and IT systems used by suppliers.
What are suitable interface technologies for this data exchange that are secure, cross- platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

A. EDJFACT XML over SFTP JSON/REST over HTTPS

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

Answer: C

NEW QUESTION 159
A mule application is required to periodically process large data set from a back-end database to Salesforce CRM using batch job scope configured properly process the higher rate of records.
The application is deployed to two cloudhub workers with no persistence queues enabled. What is the consequence if the worker crashes during records processing?

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

Answer: D

NEW QUESTION 163
What is an advantage of using OAuth 2.0 client credentials and access tokens over only API keys for API authentication?

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

Answer: A

NEW QUESTION 166
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 170
An organization is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to rejections from the back-end system will need to be processed manually (outside the back-end system).
The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization??s firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.
What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

A. An Until Successful scope to call the back-end system One or more ActiveMQ long-retry queuesOne or more ActiveMQ dead-letter queues for manual processing
B. One or more On Error scopes to assist calling the back-end system An Until Successful scope containing VM components for long retries A persistent dead-letter VM queue configured in CloudHub
C. One or more On Error scopes to assist calling the back-end system One or more ActiveMQ long-retry queuesA persistent dead-letter object store configured in the CloudHub Object Store service
D. A Batch Job scope to call the back-end systemAn Until Successful scope containing Object Store components for long retries A dead-letter object store configured in the Mule application

Answer: A

NEW QUESTION 173
An organization??s IT team must secure all of the internal APIs within an integration solution by using an API proxy to apply required authentication and authorization policies.
Which integration technology, when used for its intended purpose, should the team choose to meet these requirements if all other relevant factors are equal?

A. API Management (APIM)
B. Robotic Process Automation (RPA)
C. Electronic Data Interchange (EDI)
D. Integration Platform-as-a-service (PaaS)

Answer: A

NEW QUESTION 178
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 183
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 185
A company is designing a mule application to consume batch data from a partner's ftps server The data files have been compressed and then digitally signed using PGP. What inputs are required for the application to securely consumed these files?

A. ATLS context Key Store requiring the private key and certificate for the company PGP public key of partnerPGP private key for the company
B. ATLS context first store containing a public certificate for partner ftps server and the PGP public key of the partnerTLS contact Key Store containing the FTP credentials
C. TLS context trust or containing a public certificate for the ftps server The FTP username and password The PGP public key of the partner
D. The PGP public key of the partner The PGP private key for the company The FTP username and password

Answer: D

NEW QUESTION 189
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 190
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 191
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 192
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 193
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 195
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 197
What operation can be performed through a JMX agent enabled in a Mule application?

A. View object store entries
B. Replay an unsuccessful message
C. Set a particular tog4J2 log level to TRACE
D. Deploy a Mule application

Answer: A

NEW QUESTION 198
Refer to the exhibit.

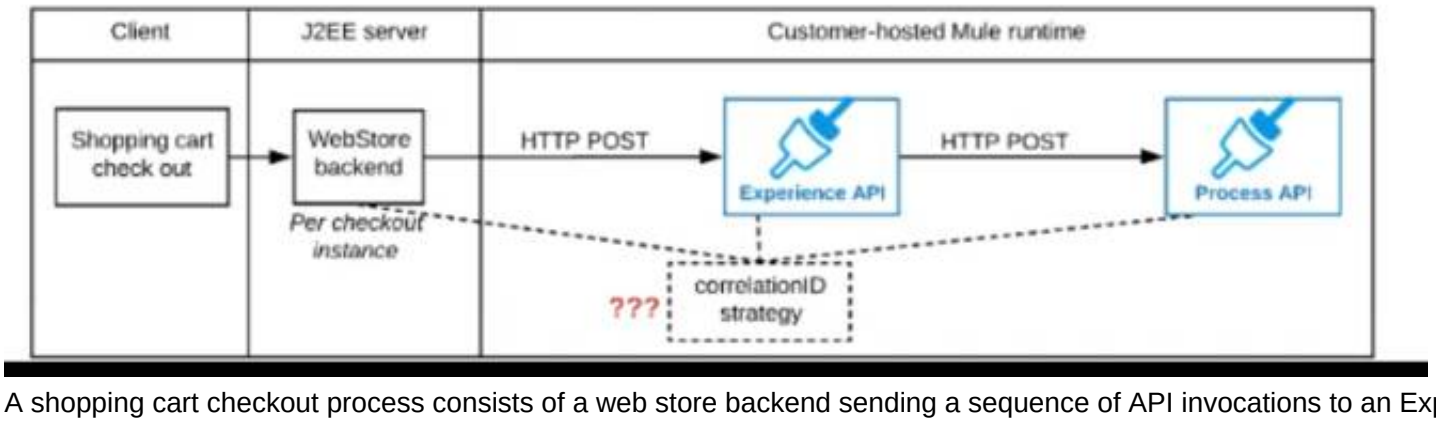

A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

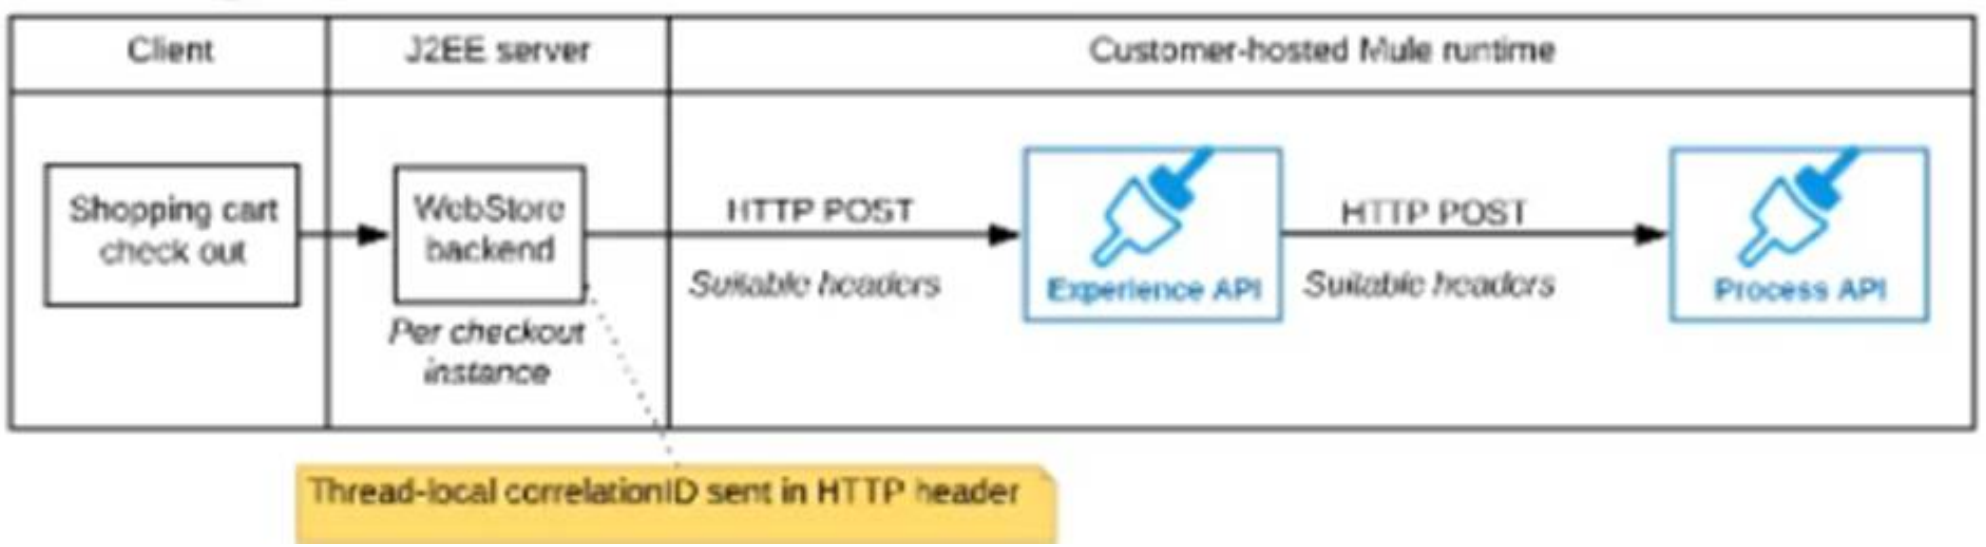

B) The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-It HTTP request header In each API invocation belonging to that checkout
No special code or configuration is included in the Experience API and Process API implementations to generate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

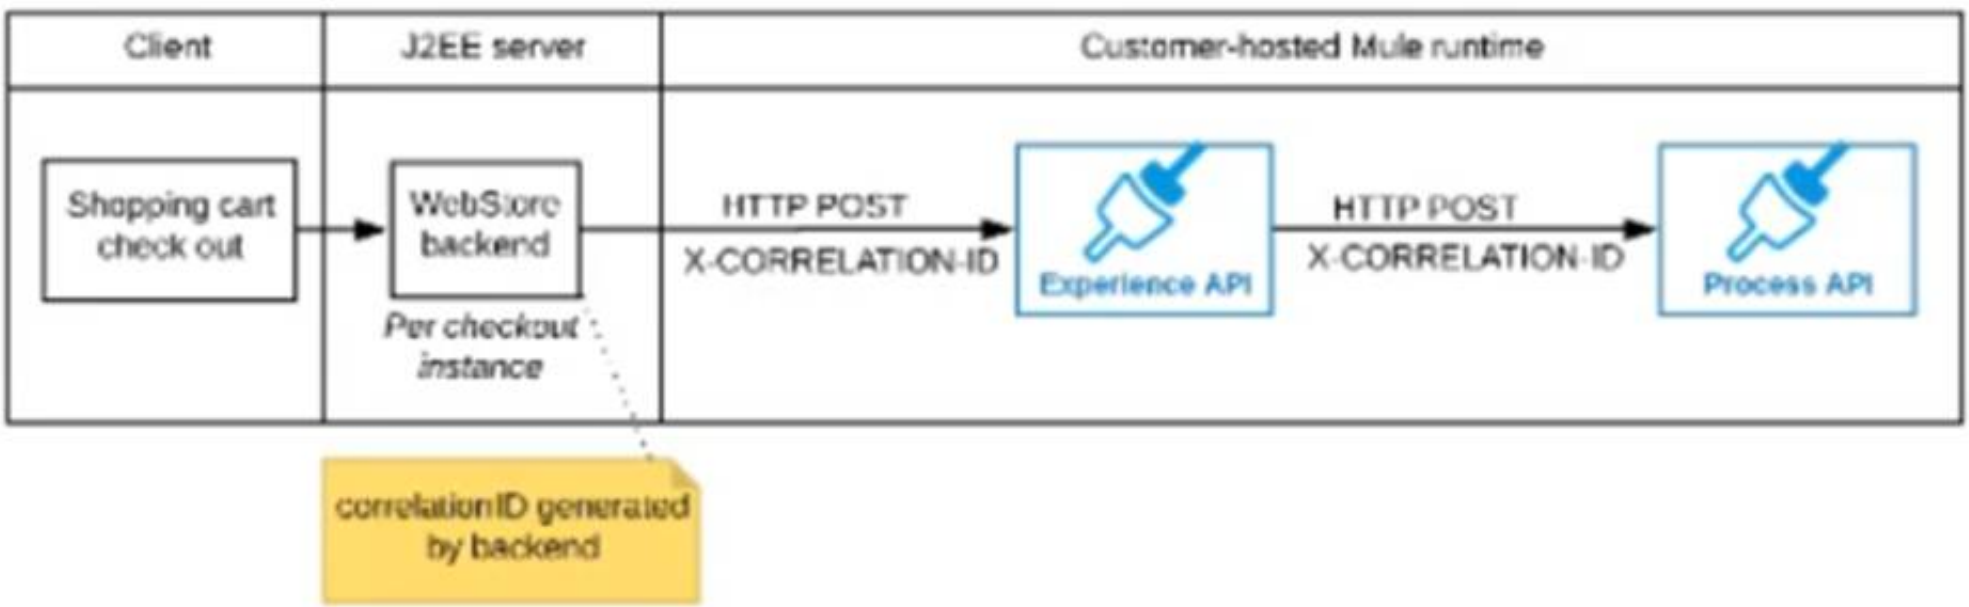

C) The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.
The Experience API implementation must be coded to also propagate the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

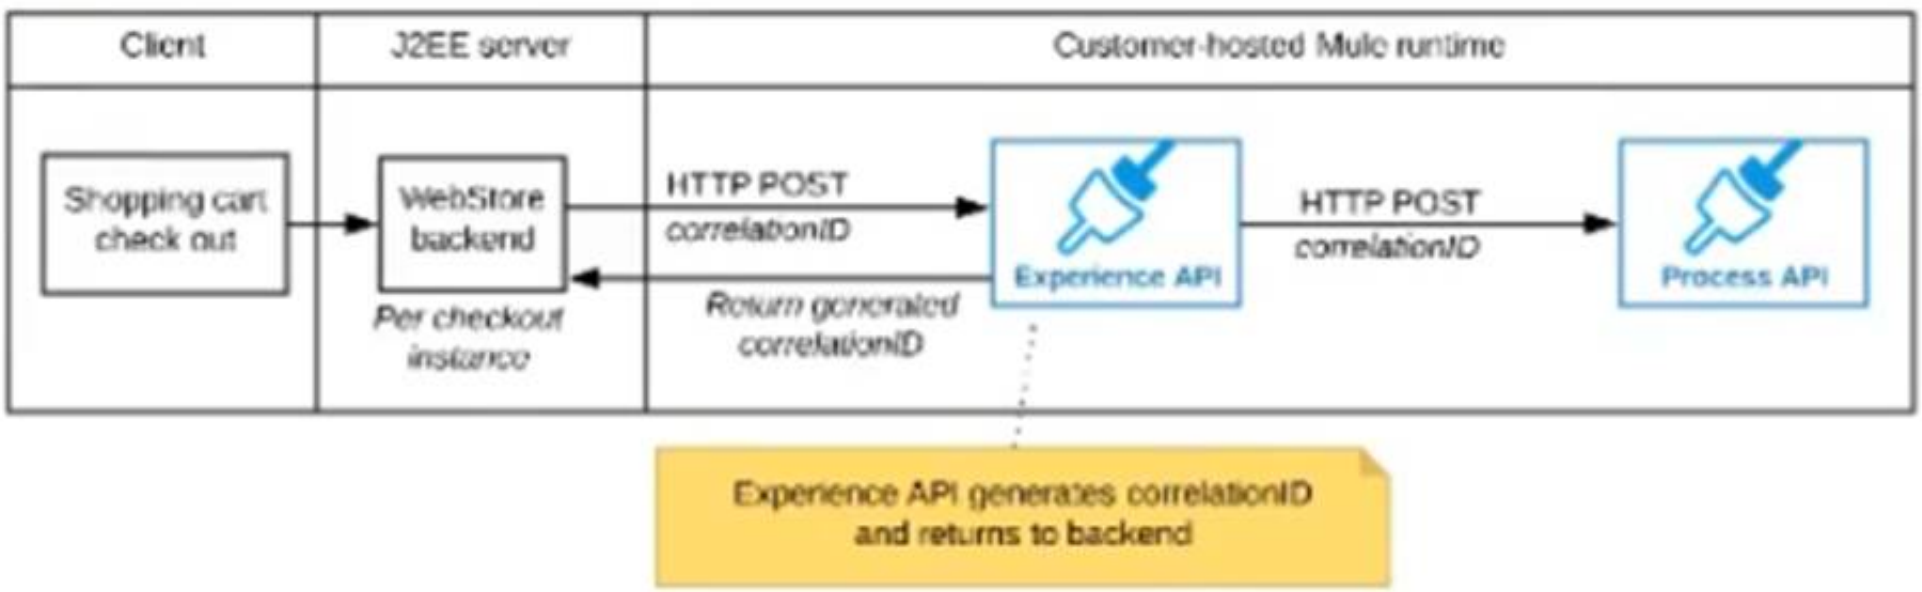

D) The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API
The Experience API and Process API implementations must be coded to receive the custom correlation ID In the HTTP requests and propagate It in suitable HTTP request headers

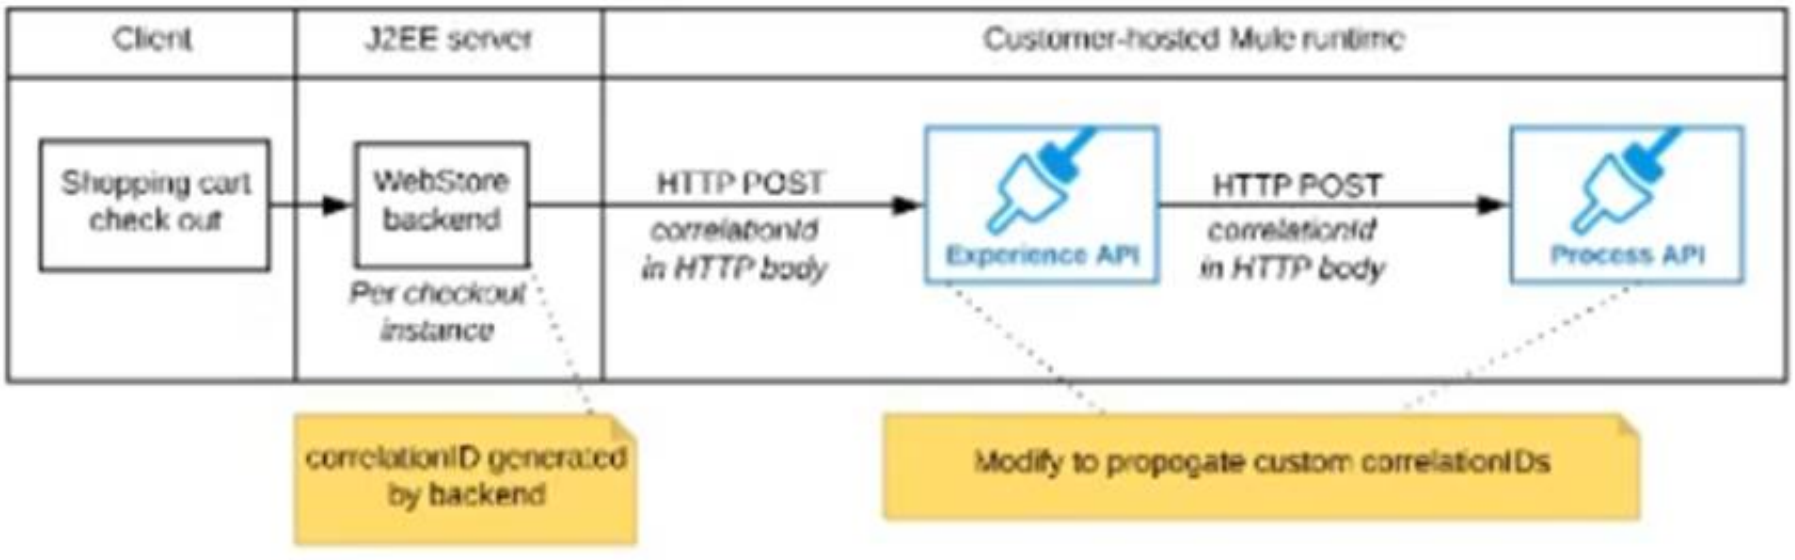

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

NEW QUESTION 200
Refer to the exhibit.

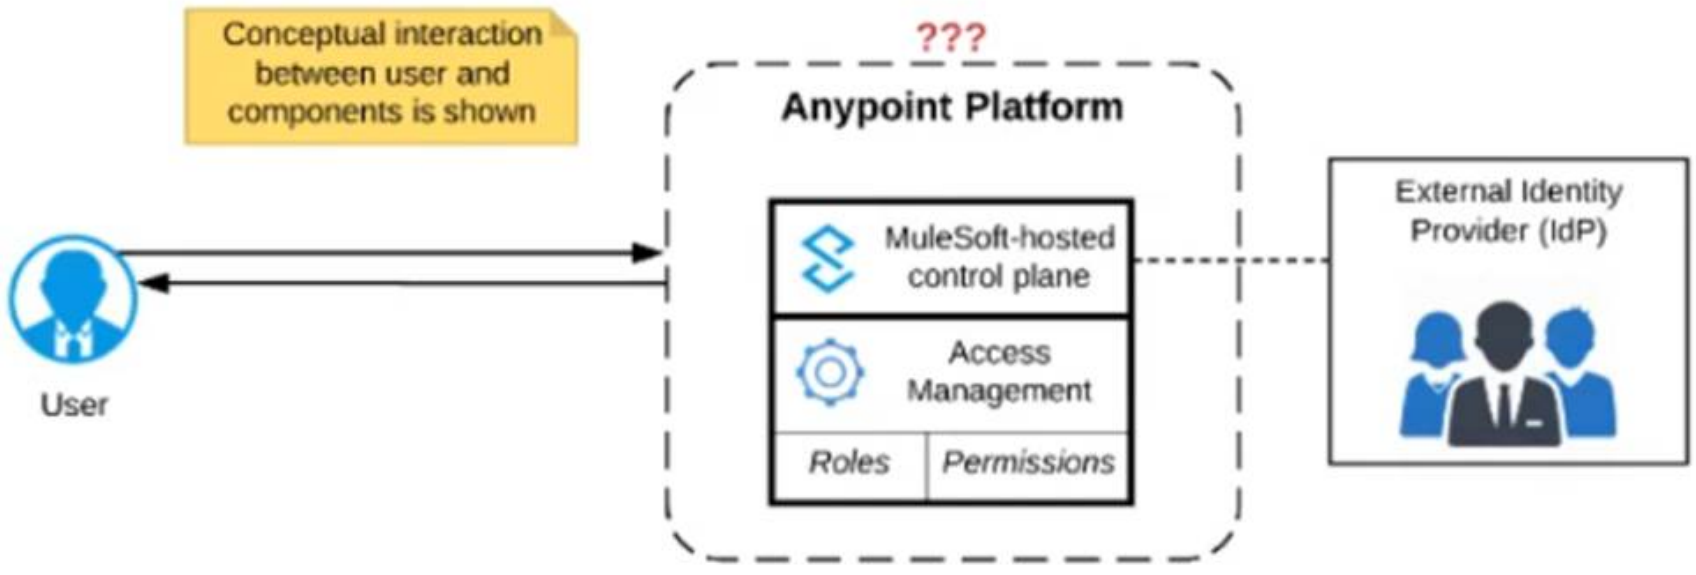

Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.
What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled 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 205
A team has completed the build and test activities for a Mule application that implements a System API for its application network.
Which Anypoint Platform component should the team now use to both deploy and monitor the System AP\ implementation?

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

Answer: D

NEW QUESTION 208
A project team is working on an API implementation using the RAML definition as a starting point. The team has updated the definition to include new operations and has published a new version to exchange. Meanwhile another team is working on a mule application consuming the same API implementation.
During the development what has to be performed by the mule application team to take advantage of the newly added operations?

A. Scaffold the client application with the new definition
B. Scaffold API implementation application with the new definition
C. Update the REST connector from exchange in the client application
D. Update the API connector in the API implementation and publish to exchange

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 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 218
Which Exchange asset type represents configuration modules that extend the functionality of an API and enforce capabilities such as security?

A. Rulesets
B. Policies
C. RESTAPIs
D. Connectors

Answer: B

NEW QUESTION 222
As a part of design , Mule application is required call the Google Maps API to perform a distance computation. The application is deployed to cloudhub.
At the minimum what should be configured in the TLS context of the HTTP request configuration to meet these requirements?

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

Answer: D

NEW QUESTION 226
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 228
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 230
Mule application A receives a request Anypoint MQ message REQU with a payload containing a variable-length list of request objects. Application A uses the For Each scope to split the list into individual objects and sends each object as a message to an Anypoint MQ queue.
Service S listens on that queue, processes each message independently of all other messages, and sends a response message to a response queue.
Application A listens on that response queue and must in turn create and publish a response Anypoint MQ message RESP with a payload containing the list of responses sent by service S in the same order as the request objects originally sent in REQU.
Assume successful response messages are returned by service S for all request messages.
What is required so that application A can ensure that the length and order of the list of objects in RESP and REQU match, while at the same time maximizing message throughput?

A. Use a Scatter-Gather within the For Each scope to ensure response message order Configure the Scatter-Gather with a persistent object store
B. Perform all communication involving service S synchronously from within the For Each scope, so objects in RESP are in the exact same order as request objects in REQU
C. Use an Async scope within the For Each scope and collect response messages in a second For Each scope in the order In which they arrive, then send RESP using this list of responses
D. Keep track of the list length and all object indices in REQU, both in the For Each scope and in all communication involving service Use persistent storage when creating RESP

Answer: D

NEW QUESTION 233
A Kubernetes controller automatically adds another pod replica to the resource pool in response to increased application load.
Which scalability option is the controller implementing?

A. Down
B. Diagonal
C. Vertical
D. Horizontal

Answer: D

NEW QUESTION 235
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 237
As a part of project , existing java implementation is being migrated to Mulesoft. Business is very tight on the budget and wish to complete the project in most economical way possible.
Canonical object model using java is already a part of existing implementation. Same object model is required by mule application for a business use case. What is the best way to achieve this?

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

Answer: A

NEW QUESTION 240
A payment processing company has implemented a Payment Processing API Mule application to process credit card and debit card transactions, Because the Payment Processing API handles highly sensitive information, the payment processing company requires that data must be encrypted both In-transit and at-rest.
To meet these security requirements, consumers of the Payment Processing API must create request message payloads in a JSON format specified by the API, and the message payload values must be encrypted.
How can the Payment Processing API validate requests received from API consumers?

A. A Transport Layer Security (TLS) - Inbound policy can be applied in API Manager to decrypt the message payload and the Mule application implementation can then use the JSON Validation module to validate the JSON data
B. The Mule application implementation can use the APIkit module to decrypt and then validate the JSON data
C. The Mule application implementation can use the Validation module to decrypt and then validate the JSON data
D. The Mule application implementation can use DataWeave to decrypt the message payload and then use the JSON Scheme Validation module to validate the JSON data

Answer: A

NEW QUESTION 241
An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to custom hosted Mule runtimes (on-premises) while others execute in the MuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems.
How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access the backend systems?

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

Answer: D

NEW QUESTION 243
Customer has deployed mule applications to different customer hosted mule run times. Mule applications are managed from Anypoint platform.
What needs to be configured to monitor these Mule applications from Anypoint monitoring and what sends monitoring data to Anypoint monitoring?

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

Answer: A

NEW QUESTION 247
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 251
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 252
When a Mule application using VM queues is deployed to a customer-hosted cluster or multiple CloudHub v1.0 workers/replicas, how are messages consumed across the nodes?

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

Answer: C

NEW QUESTION 254
An organization is creating a set of new services that are critical for their business. The project team prefers using REST for all services but is willing to use SOAP with common WS-" standards if a particular service requires it.
What requirement would drive the team to use SOAP/WS-* for a particular service?

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

Answer: D

NEW QUESTION 259
What Anypoint Connectors support transactions?

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

Answer: A

NEW QUESTION 264
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 268
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 273
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 278
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 281
What API policy would LEAST likely be applied to a Process API?

A. Custom circuit breaker
B. Client ID enforcement
C. Rate limiting
D. JSON threat protection

Answer: D

NEW QUESTION 283
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 286
Refer to the exhibit.

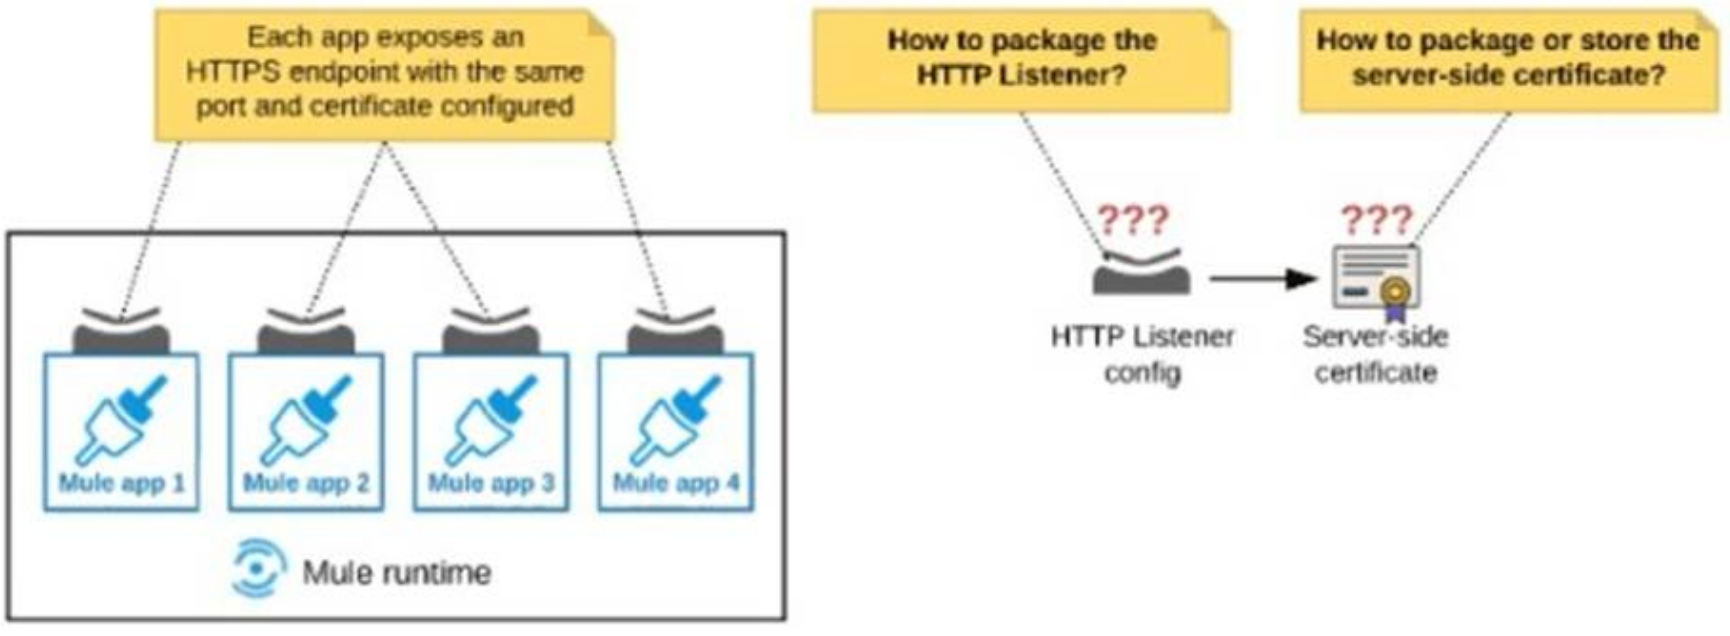

An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a server-side certificate that rotates often.
What is the most effective way to package the HTTP Listener and package or store 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 289
A company wants its users to log in to Anypoint Platform using the company's own internal user credentials. To achieve this, the company needs to integrate an external identity provider (IdP) with the company's Anypoint Platform master organization, but SAML 2.0 CANNOT be used. Besides SAML 2.0, what single-sign-on standard can the company use to integrate the IdP with their Anypoint Platform master organization?

A. SAML 1.0
B. OAuth 2.0
C. Basic Authentication
D. OpenID Connect

Answer: D

NEW QUESTION 290
How does timeout attribute help inform design decisions while using JMS connector listening for incoming messages in an extended architecture (XA) transaction?

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

Answer: D

NEW QUESTION 294
A mule application is deployed to a Single Cloudhub worker and the public URL appears in Runtime Manager as the APP URL.
Requests are sent by external web clients over the public internet to the mule application App url. Each of these requests routed to the HTTPS Listener event source of the running Mule application.
Later, the DevOps team edits some properties of this running Mule application in Runtime Manager.
Immediately after the new property values are applied in runtime manager, how is the current Mule application deployment affected and how will future web client requests to the Mule application be handled?

A. Cloudhub will redeploy the Mule application to the OLD Cloudhub workerNew web client requests will RETURN AN ERROR until the Mule application is redeployed to the OLD Cloudhub worker
B. CloudHub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests will RETURN AN ERROR until the NEW Cloudhub worker is available
C. Cloudhub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker until the NEW Cloudhub worker is available.
D. Cloudhub will redeploy the mule application to the OLD Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.

Answer: C

NEW QUESTION 299
According to the National Institute of Standards and Technology (NIST), which cloud computing deployment model describes a composition of two or more distinct clouds that support data and application portability?

A. Private cloud
B. Hybrid cloud 4
C. Public cloud
D. Community cloud

Answer: B

NEW QUESTION 301
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 305
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 307
Refer to the exhibit.
A Mule application is being designed to expose a SOAP web service to its clients.
What language is typically used inside the web service's interface definition to define the data structures that the web service Is expected to exchange with its clients?

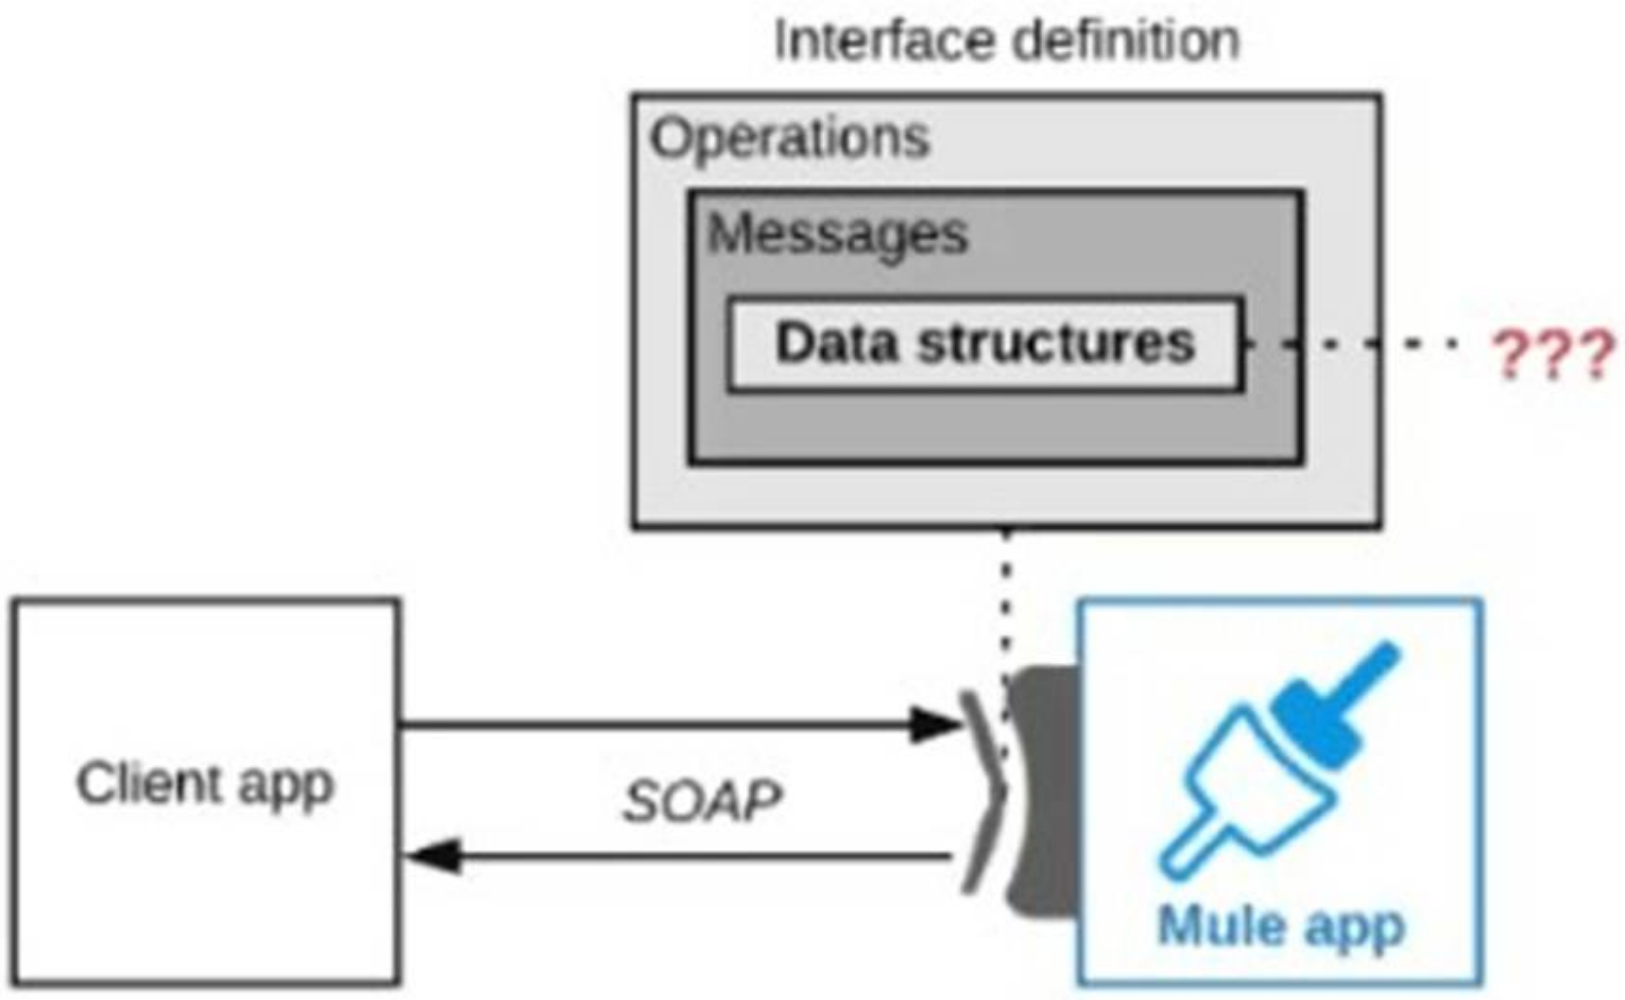

A. WSDL
B. XSD
C. JSON Schema
D. RAMI

Answer: B

NEW QUESTION 309
One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make.
Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments.
The production environment however does not have any rate limiting.
What is the cost-effective approach to conduct performance test of the API implementation in the non-production staging environment?

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

Answer: C

NEW QUESTION 310
An insurance company has an existing API which is currently used by customers. API is deployed to customer hosted Mule runtime cluster. The load balancer that is used to access any APIs on the mule cluster is only configured to point to applications hosted on the server at port 443.
Mule application team of a company attempted to deploy a second API using port 443 but the application will not start and checking logs shows an error indicating the address is already in use.
Which steps must the organization take to resolve this error and allow customers to access both the API's?

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

Answer: C

NEW QUESTION 312
Refer to the exhibit.

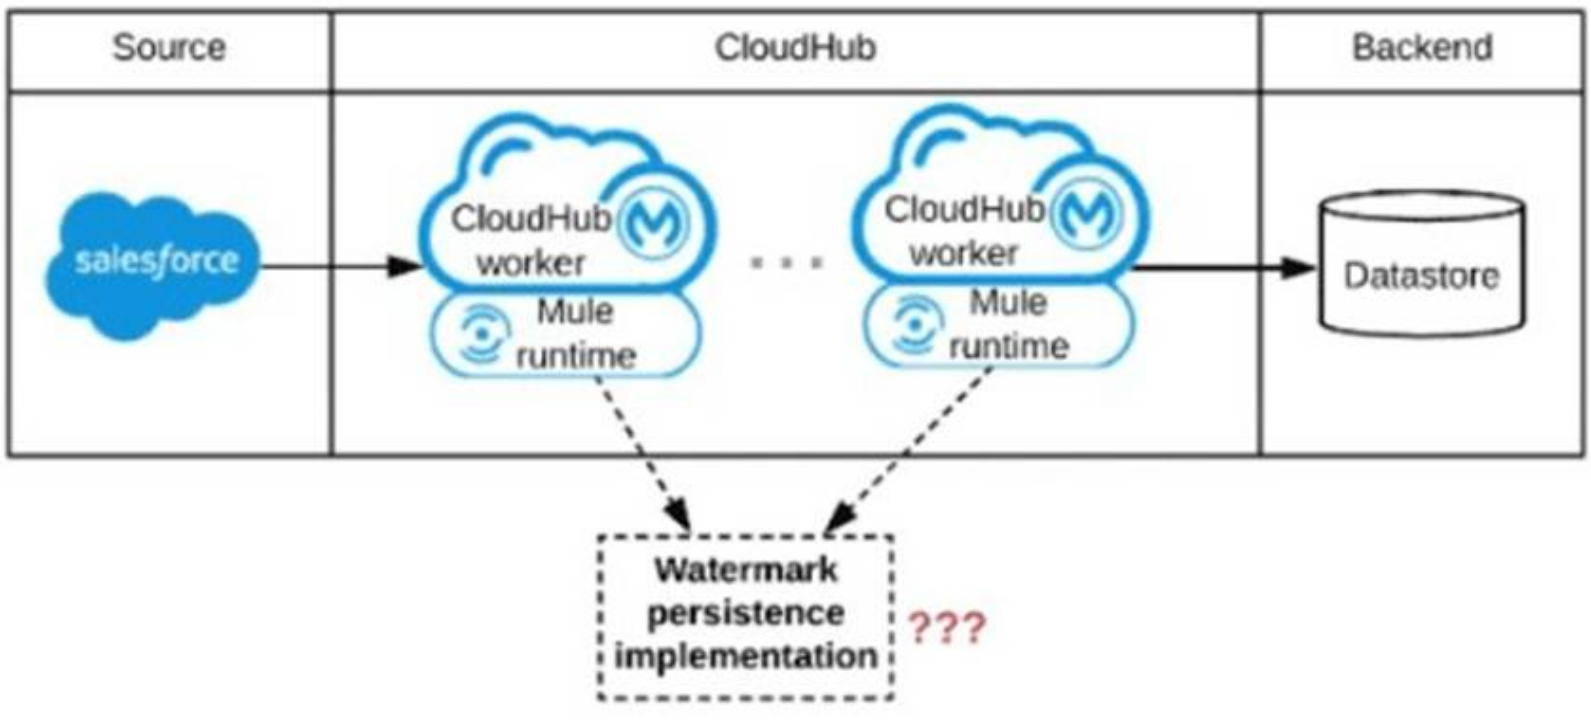

A Mule application is being designed to be deployed to several CIoudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.
A watermark will be used to only retrieve those Satesforce Accounts that have been modified since the last time the integration logic ran.
What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?

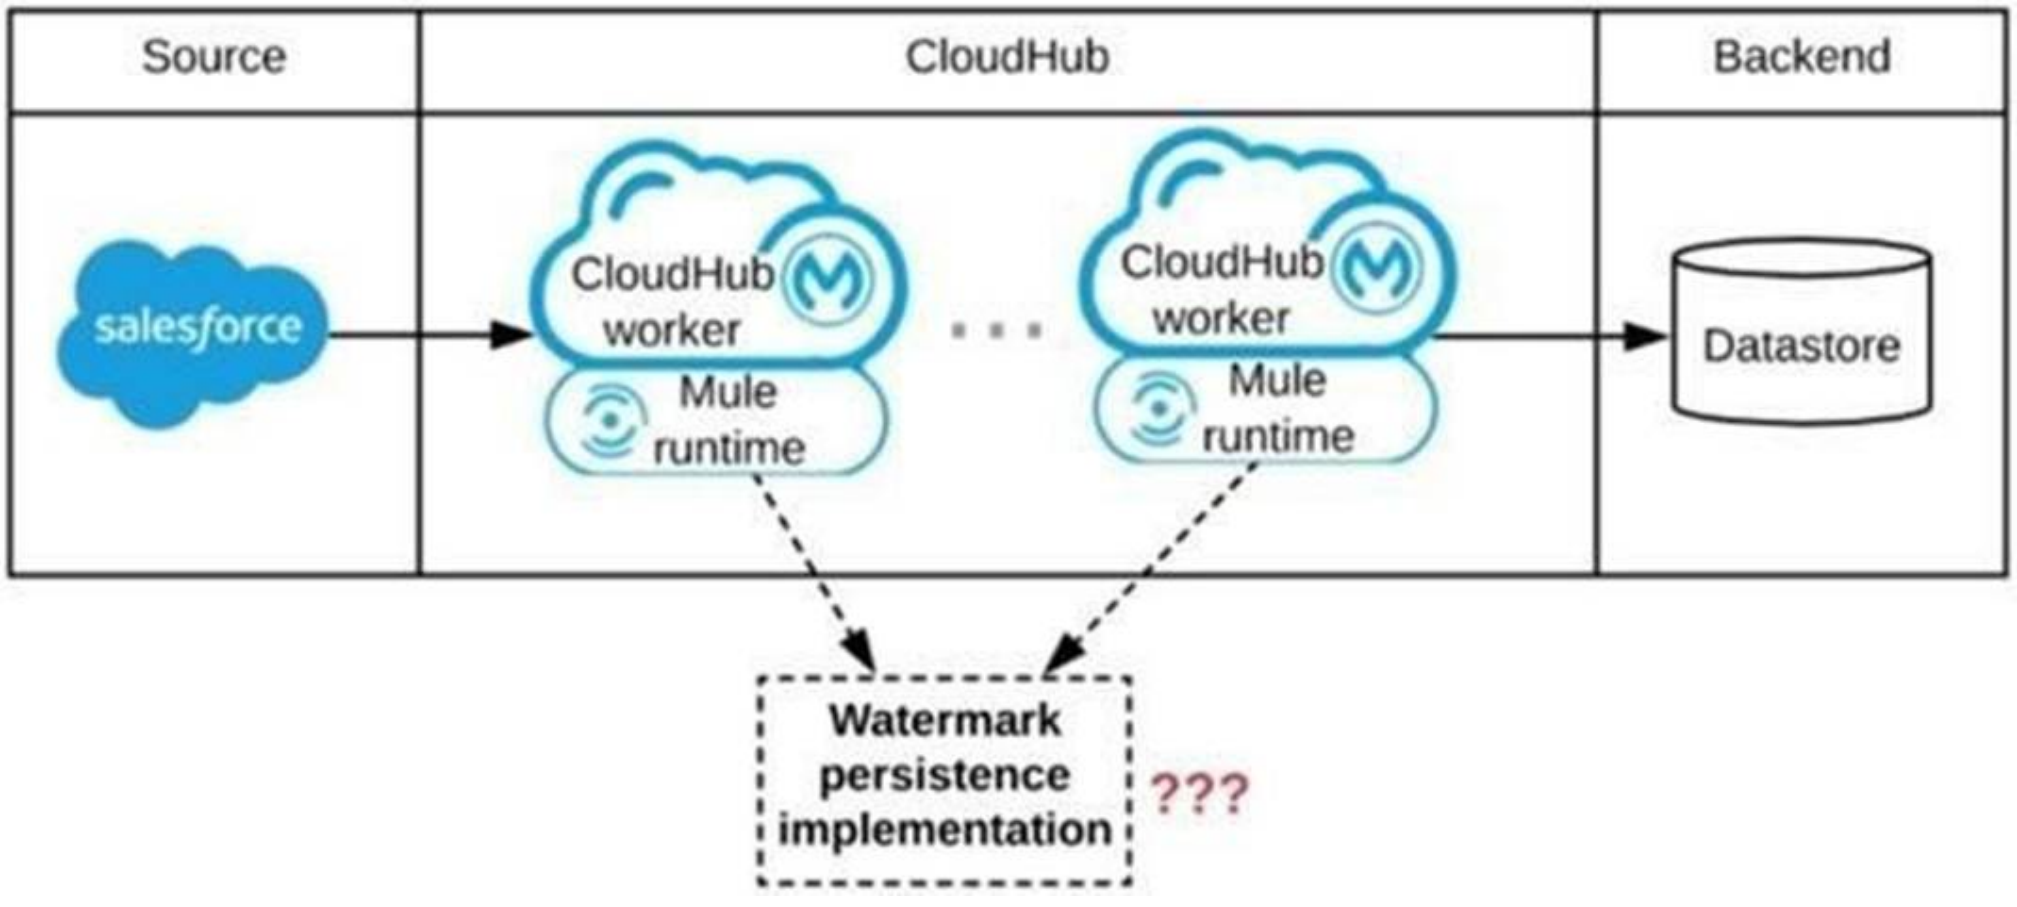

A. Persistent Anypoint MQ Queue
B. Persistent Object Store
C. Persistent Cache Scope
D. Persistent VM Queue

Answer: B

NEW QUESTION 315
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 319
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
An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.
What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

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

Answer: A

NEW QUESTION 326
Refer to the exhibit.

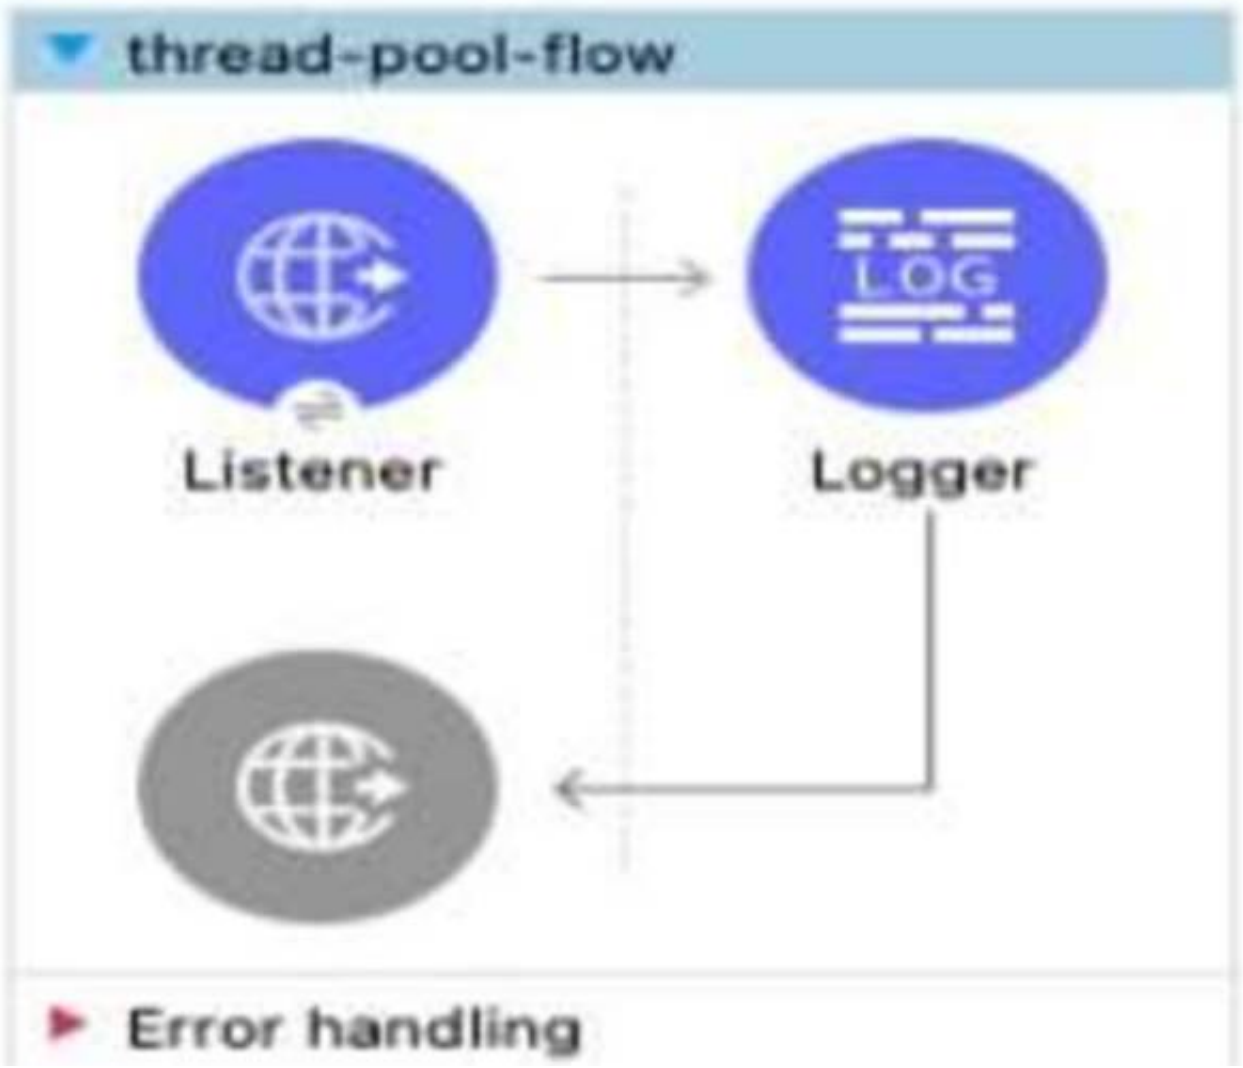

A customer is running Mule applications on Runtime Fabric for Self-Managed Kubernetes (RTF-BYOKS) in a multi-cloud environment.
Based on this configuration, how do Agents and Runtime Manager communicate, and what Is exchanged between them?

A. BLOCKING_IO, UBER
B. UBER, Dedicated NIO Selector Pool
C. CPU_LITE, CPU_INTENSIVE
D. Shared NIO Selector Pool, CPU_LITE

Answer: D

NEW QUESTION 330
Refer to the exhibit.

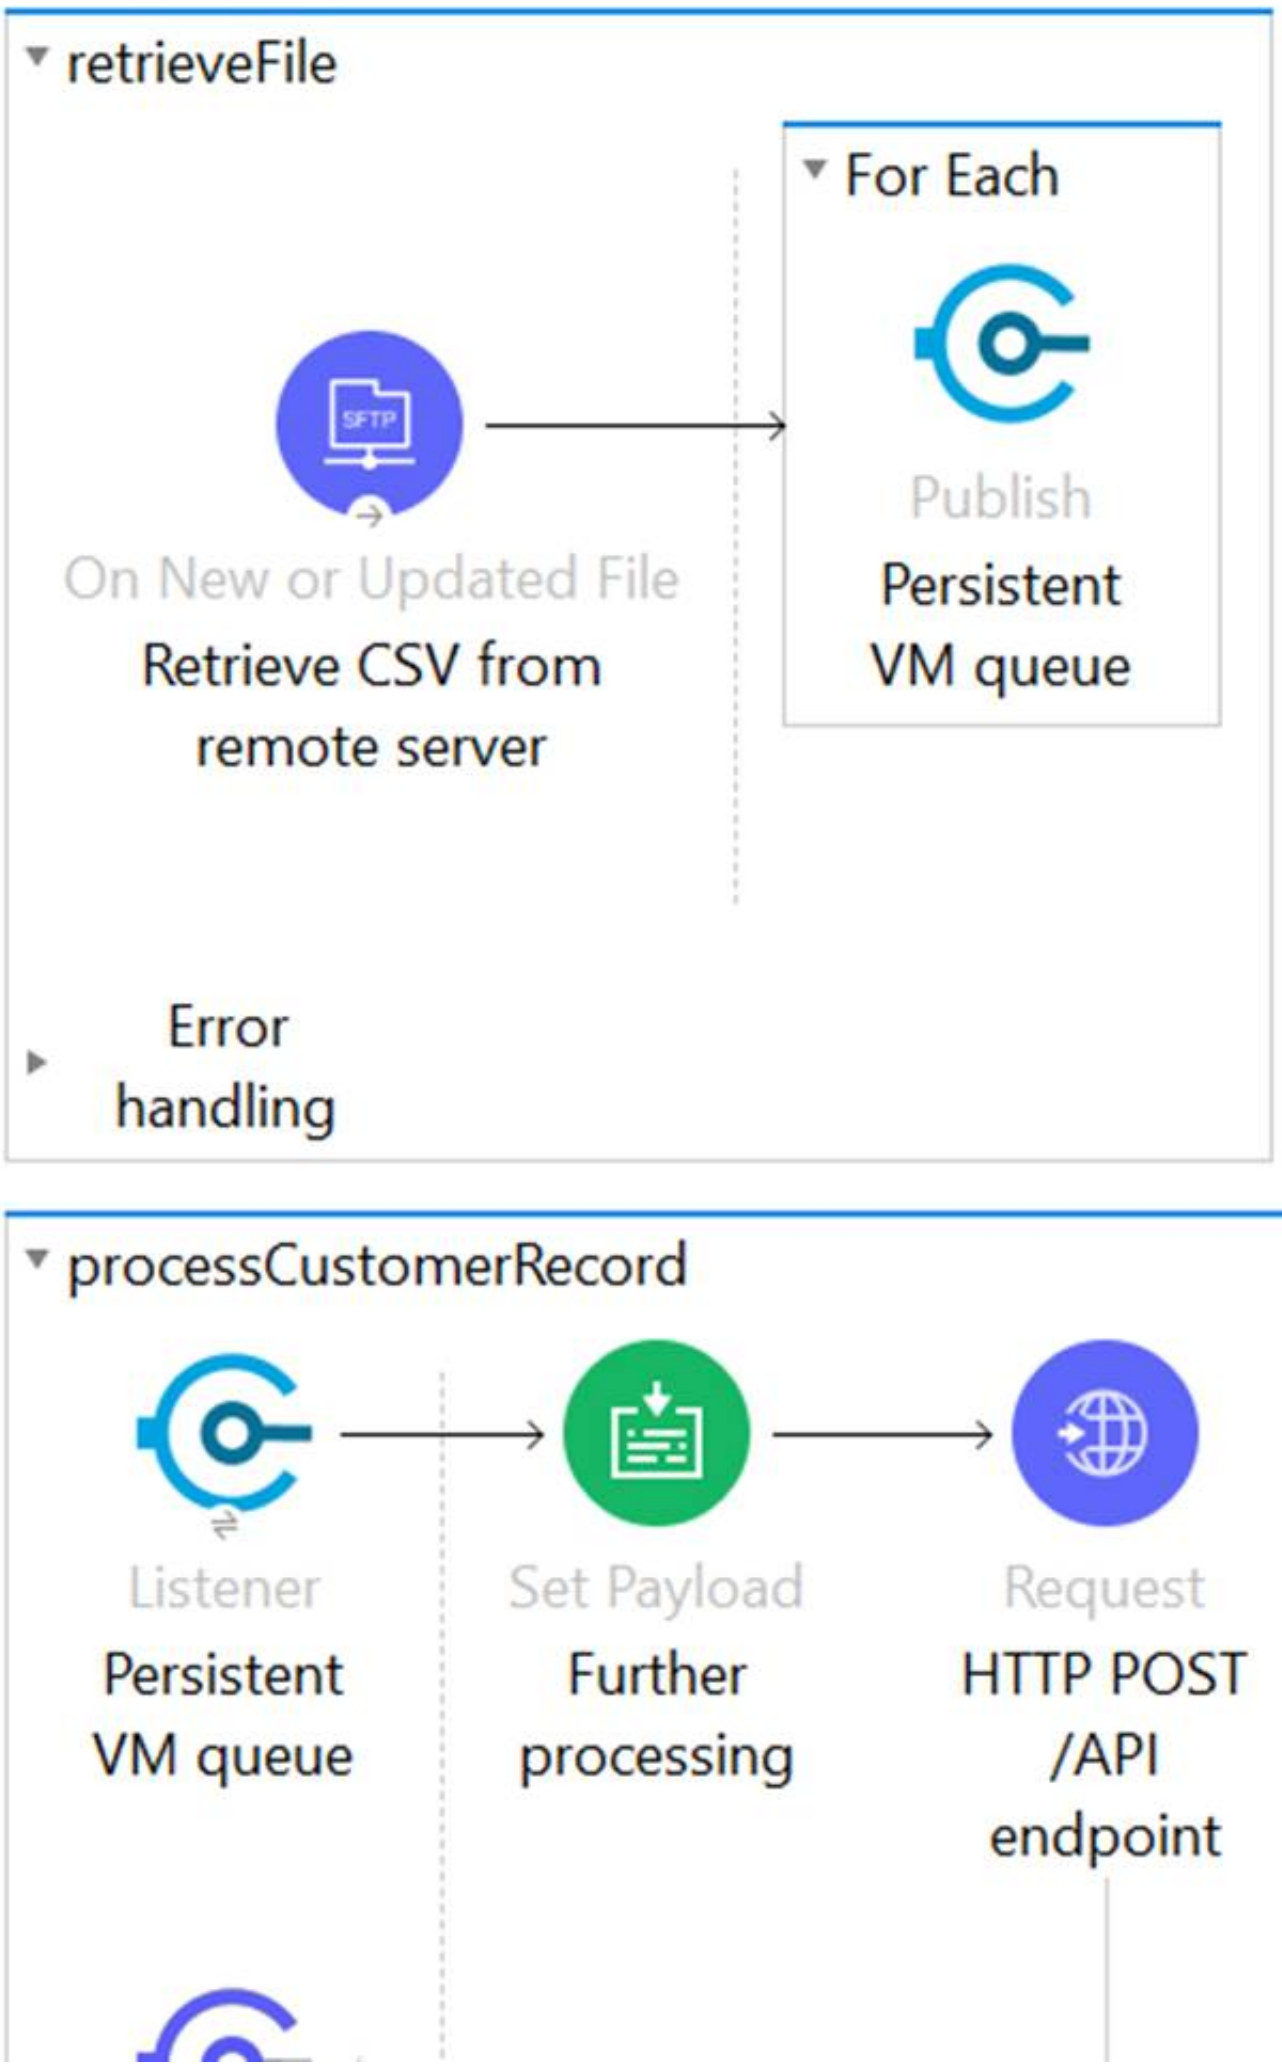

This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow??s VM Listner receives messages from 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 333
An organization uses a four(4) node customer hosted Mule runtime cluster to host one(1)
stateless api implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution. Each node in the cluster has been sized to be able to accept four(4) times the current number of requests.
Two(2) nodes in the cluster experience a power outage and are no longer available. The load balancer directs the outage and blocks the two unavailable the nodes from receiving further HTTP requests.
What performance-related consequence is guaranteed to happen to average, assuming the remaining cluster nodes are fully operational?

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

Answer: A

NEW QUESTION 337
According to MuleSoft, what Action should an IT organization take regarding its technology assets in order to close the IT delivery.

A. Make assets easily discoverable via a central repository
B. Focus project delivery efforts on custom assets that meet the specific requirements of each individual line of business
C. Create weekly meetings that all members of IT attend to present justification and request approval to use existing assets
D. Hire additional staff to meet the demand for asset creation required for approved projects and timelines

Answer: A

NEW QUESTION 339
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 340
What is a defining characteristic of an integration-Platform-as-a-Service (iPaaS)?

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

Answer: A

NEW QUESTION 341
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 344
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 348
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 353
What is required before an API implemented using the components of Anypoint Platform can be managed and governed (by applying API policies) on Anypoint Platform?

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

Answer: A

NEW QUESTION 354
What condition requires using a CloudHub Dedicated Load Balancer?

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

Answer: D

NEW QUESTION 355
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
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

MuleSoft-Integration-Architect-I Practice Exam Features:

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

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

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

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