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
According to MuleSoft, which system integration term describes the method, format, and protocol used for communication between two system?

A. Component
B. interaction
C. Message
D. Interface

Answer: D

NEW QUESTION 3
An insurance provider is implementing Anypoint platform to manage its application infrastructure and is using the customer hosted runtime for its business due to certain financial requirements it must meet. It has built a number of synchronous API's and is currently hosting these on a mule runtime on one server
These applications make use of a number of components including heavy use of object stores and VM queues.
Business has grown rapidly in the last year and the insurance provider is starting to receive reports of reliability issues from its applications.
The DevOps team indicates that the API's are currently handling too many requests and this is over loading the server. The team has also mentioned that there is a significant downtime when the server is down for maintenance.
As an integration architect, which option would you suggest to mitigate these issues?

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

Answer: B

NEW QUESTION 4
Which Anypoint Platform component helps integration developers discovers and share reusable APIs, connectors, and templates?

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

Answer: A

NEW QUESTION 5
An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL https://api.acme.com
As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub
What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

A. Deploy the Mule applications to CloudhubUpdate the CNAME record for an api.acme.com in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map URLs to their corresponding Mule applications
B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB)Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications
C. Deploy the Mule applications to CloudhubCreate CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancerApply mapping rules in the SLB to map URLs to their corresponding Mule applications
D. Deploy the Mule applications to CloudhubUpdate the CNAME record for api.acme.com in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB)Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

Answer: A

NEW QUESTION 6
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 7
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 8
Which Salesforce API is invoked to deploy, retrieve, create or delete customization information such as custom object definitions using a Mule Salesforce connector in a Mule application?

A. Metadata API
B. REST API
C. SOAP API
D. Bulk API

Answer: A

NEW QUESTION 9
A mule application uses an HTTP request operation to involve an external API. The external API follows the HTTP specification for proper status code usage.
What is possible cause when a 3xx status code is returned to the HTTP Request operation from the external API?

A. The request was not accepted by the external API
B. The request was Redirected to a different URL by the external API
C. The request was NOT RECEIVED by the external API
D. The request was ACCEPTED by the external API

Answer: B

NEW QUESTION 10
Which of the below requirements prevent the usage of Anypoint MQ in a company's network? (Choose two answers)

A. single message payload can be up to 15 MB
B. payloads must be encrypted
C. the message broker must be hosted on premises
D. support for point-to-point messaging
E. ability for a third party outside the company's network to consume events from the queue

Answer: CE

NEW QUESTION 10
According to MuleSoft's API development best practices, which type of API development approach starts with writing and approving an API contract?

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

Answer: D

NEW QUESTION 15
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 18
An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/'CD pipeline. Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?

A. Use runtime manager rest apis for API management and mavenforAPI deployment
B. Use Maven with a custom configuration required for the API lifecycle

C. Use Anypoint CLI or Anypoint Platform REST apis with scripting language such as groovy
D. Use Exchange rest api's for API management and MavenforAPI deployment

Answer: C

NEW QUESTION 22
As an enterprise architect, what are the two reasons for which you would use a canonical data model in the new integration project using Mulesoft Anypoint platform (choose two answers)

A. To have consistent data structure aligned in processes
B. To isolate areas within a bounded context
C. To incorporate industry standard data formats
D. There are multiple canonical definitions of each data type
E. Because the model isolates the back and systems and support mule applications from change

Answer: AE

NEW QUESTION 24
A global organization operates datacenters in many countries. There are private network links between these datacenters because all business data (but NOT metadata) must be exchanged over these private network connections.
The organization does not currently use AWS in any way.
The strategic decision has Just been made to rigorously minimize IT operations effort and investment going forward.
What combination of deployment options of the Anypoint Platform control plane and runtime plane(s) best serves this organization at the start of this strategic journey?

A. MuleSoft-hosted Anypoint Platform control plane CloudHub Shared Worker Cloud in multiple AWS regions
B. Anypoint Platform - Private Cloud Edition Customer-hosted runtime plane in each datacenter
C. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in multiple AWS regions
D. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in each datacenter

Answer: D

NEW QUESTION 26
Mule application muleA deployed in cloudhub uses Object Store v2 to share data across instances. As a part of new requirement , application muleB which is deployed in same region wants to access this Object Store.
Which of the following option you would suggest which will have minimum latency in this scenario?

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

Answer: A

NEW QUESTION 29
An organization is designing a mule application to support an all or nothing transaction between serval database operations and some other connectors so that they all roll back if there is a problem with any of the connectors
Besides the database connector , what other connector can be used in the transaction.

A. VM
B. Anypoint MQ
C. SFTP
D. ObjectStore

Answer: A

NEW QUESTION 31
Which component of Anypoint platform belongs to the platform control plane?

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

Answer: D

NEW QUESTION 34
Organization wants to achieve high availability goal for Mule applications in customer hosted runtime plane. Due to the complexity involved, data cannot be shared among of different instances of same Mule application. What option best suits to this requirement considering high availability is very much critical to the organization?

A. The cluster can be configured
B. Use third party product to implement load balancer
C. High availability can be achieved only in CloudHub
D. Use persistent object store

Answer: B

NEW QUESTION 36
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 39
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 42
A leading eCommerce giant will use MuleSoft APIs on Runtime Fabric (RTF) to process customer orders. Some customer-sensitive information, such as credit card information, is required in request payloads or is included in response payloads in some of the APIs. Other API requests and responses are not authorized to access some of this customer- sensitive information but have been implemented to validate and transform based on the structure and format of this customer-sensitive information (such as account IDs, phone numbers, and postal codes).
What approach configures an API gateway to hide sensitive data exchanged between API consumers and API implementations, but can convert tokenized fields back to their original value for other API requests or responses, without having to recode the API implementations?
Later, the project team requires all API specifications to be augmented with an additional non-functional requirement (NFR) to protect the backend services from a high rate of requests, according to defined service-level agreements (SLAs). The NFR's SLAs are based on a new tiered subscription level "Gold", "Silver", or "Platinum" that must be tied to a new parameter that is being added to the Accounts object in their enterprise data model.
Following MuleSoft's recommended best practices, how should the project team now convey the necessary non-functional requirement to stakeholders?

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

Answer: C

NEW QUESTION 44
Refer to the exhibit.

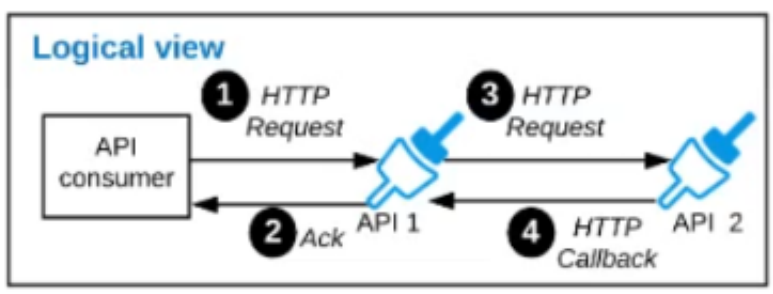

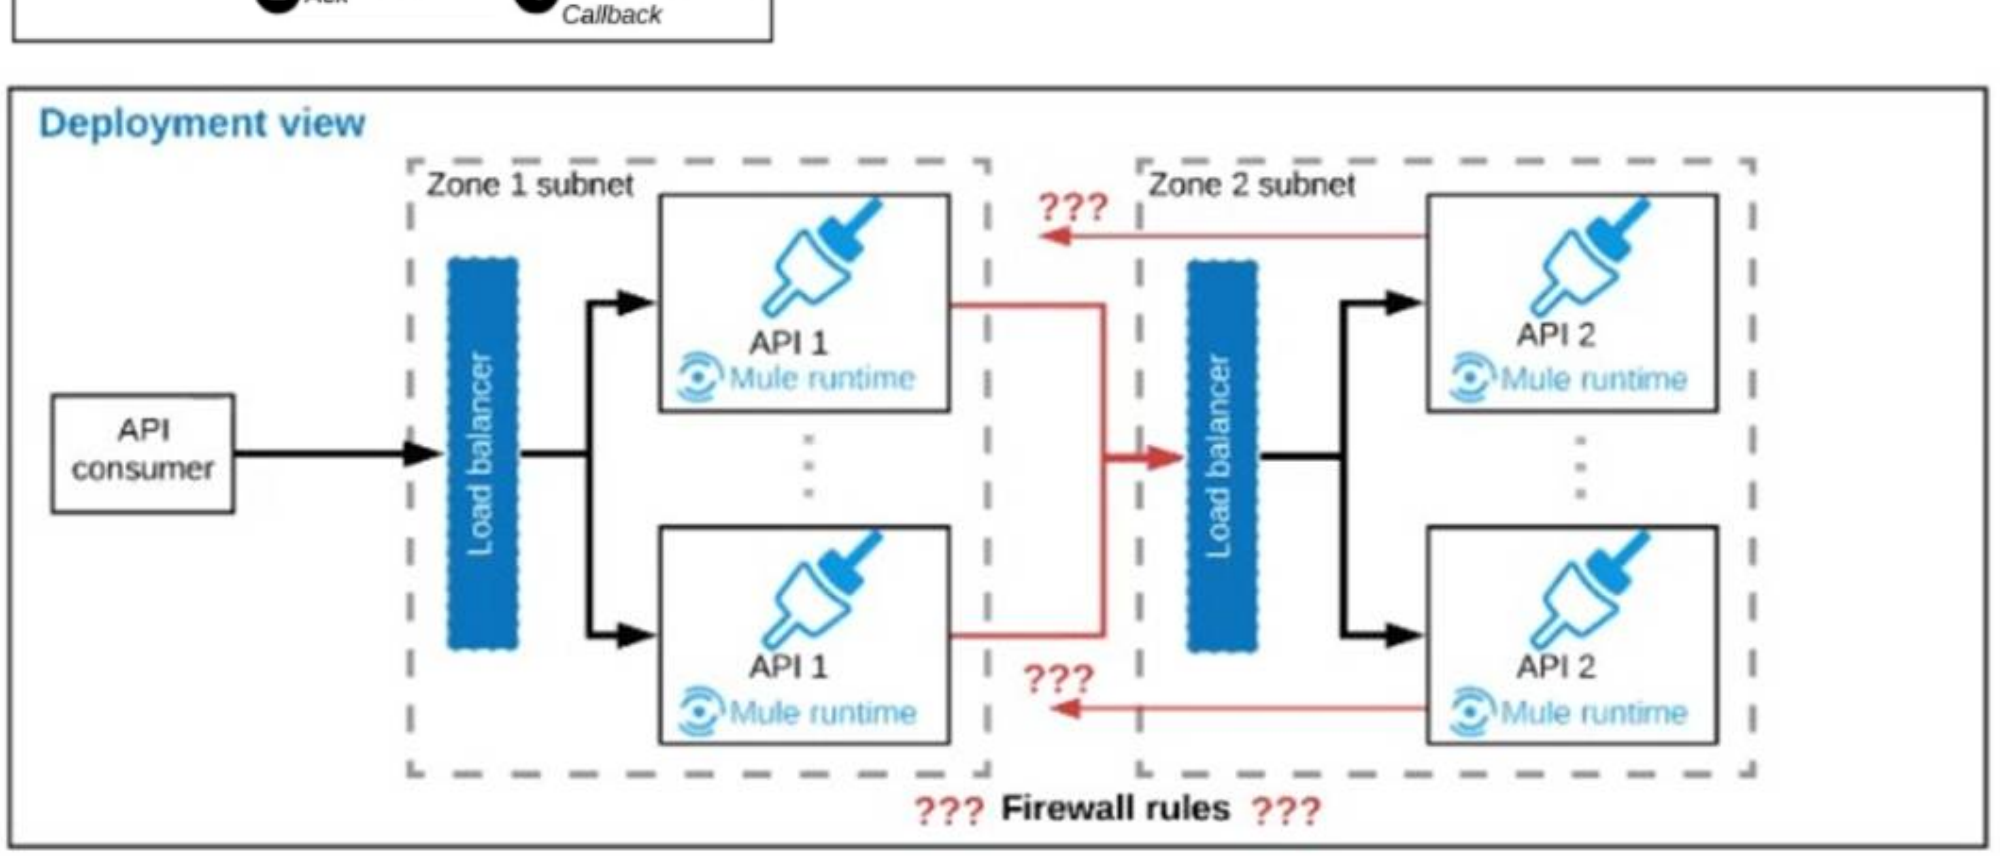

A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the

initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API l to notify about completion of the asynchronous process.
Each API Is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.
In the network architecture, how must the firewall rules be configured to enable the above Interaction between API 1 and API 2?

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

Answer: B

NEW QUESTION 45
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 48
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 49
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 54
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 58
An organization has several APIs that accept JSON data over HTTP POST. The APIs are all publicly available and are associated with several mobile applications and web applications. The organization does NOT want to use any authentication or compliance policies for these APIs, but at the same time, is worried that some bad actor could send payloads that could somehow compromise the applications or servers running the API implementations. What out-of-the-box Anypoint Platform policy can address exposure to this threat?

A. Apply a Header injection and removal policy that detects the malicious data before it is used
B. Apply an IP blacklist policy to all APIs; the blacklist will Include all bad actors
C. Shut out bad actors by using HTTPS mutual authentication for all API invocations
D. Apply a JSON threat protection policy to all APIs to detect potential threat vectors

Answer: D

NEW QUESTION 62
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 63
What is an example of data confidentiality?

A. Signing a file digitally and sending it using a file transfer mechanism
B. Encrypting a file containing personally identifiable information (PV)
C. Providing a server's private key to a client for secure decryption of data during a two- way SSL handshake
D. De-masking a person's Social Security number while inserting it into a database

Answer: B

NEW QUESTION 64
An Order microservice and a Fulfillment microservice are being designed to communicate with their dients through message-based integration (and NOT through API invocations).
The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mute application, the Fulfillment microservice.
The Fulfilment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilted message (a kind of event message). Each OrderFulfilted message can be consumed by any interested Mule application, and the Order microservice is one such Mute application.
What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

A. Order messages are sent to an Anypoint MQ exchange OrderFulfilled messages are sent to an Anypoint MQ queue Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices
B. Order messages are sent to a JMS queu
C. OrderFulfilled messages are sent to a JMS topic Both microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices
D. Order messages are sent directly to the Fulfillment microservice
E. OrderFulfilled messages are sent directly to the Order microservice The Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load of each microservice
F. Order messages are sent to a JMS queu
G. OrderFulfilled messages are sent to a JMS topic The Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice

Answer: B

NEW QUESTION 68
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 71
What are two reasons why a typical MuleSoft customer favors a MuleSoft-hosted Anypoint Platform runtime plane over a customer-hosted runtime for its Mule application deployments? (Choose two.)

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

Answer: CE

NEW QUESTION 75
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 79
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 80
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 83
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 87
Refer to the exhibit.

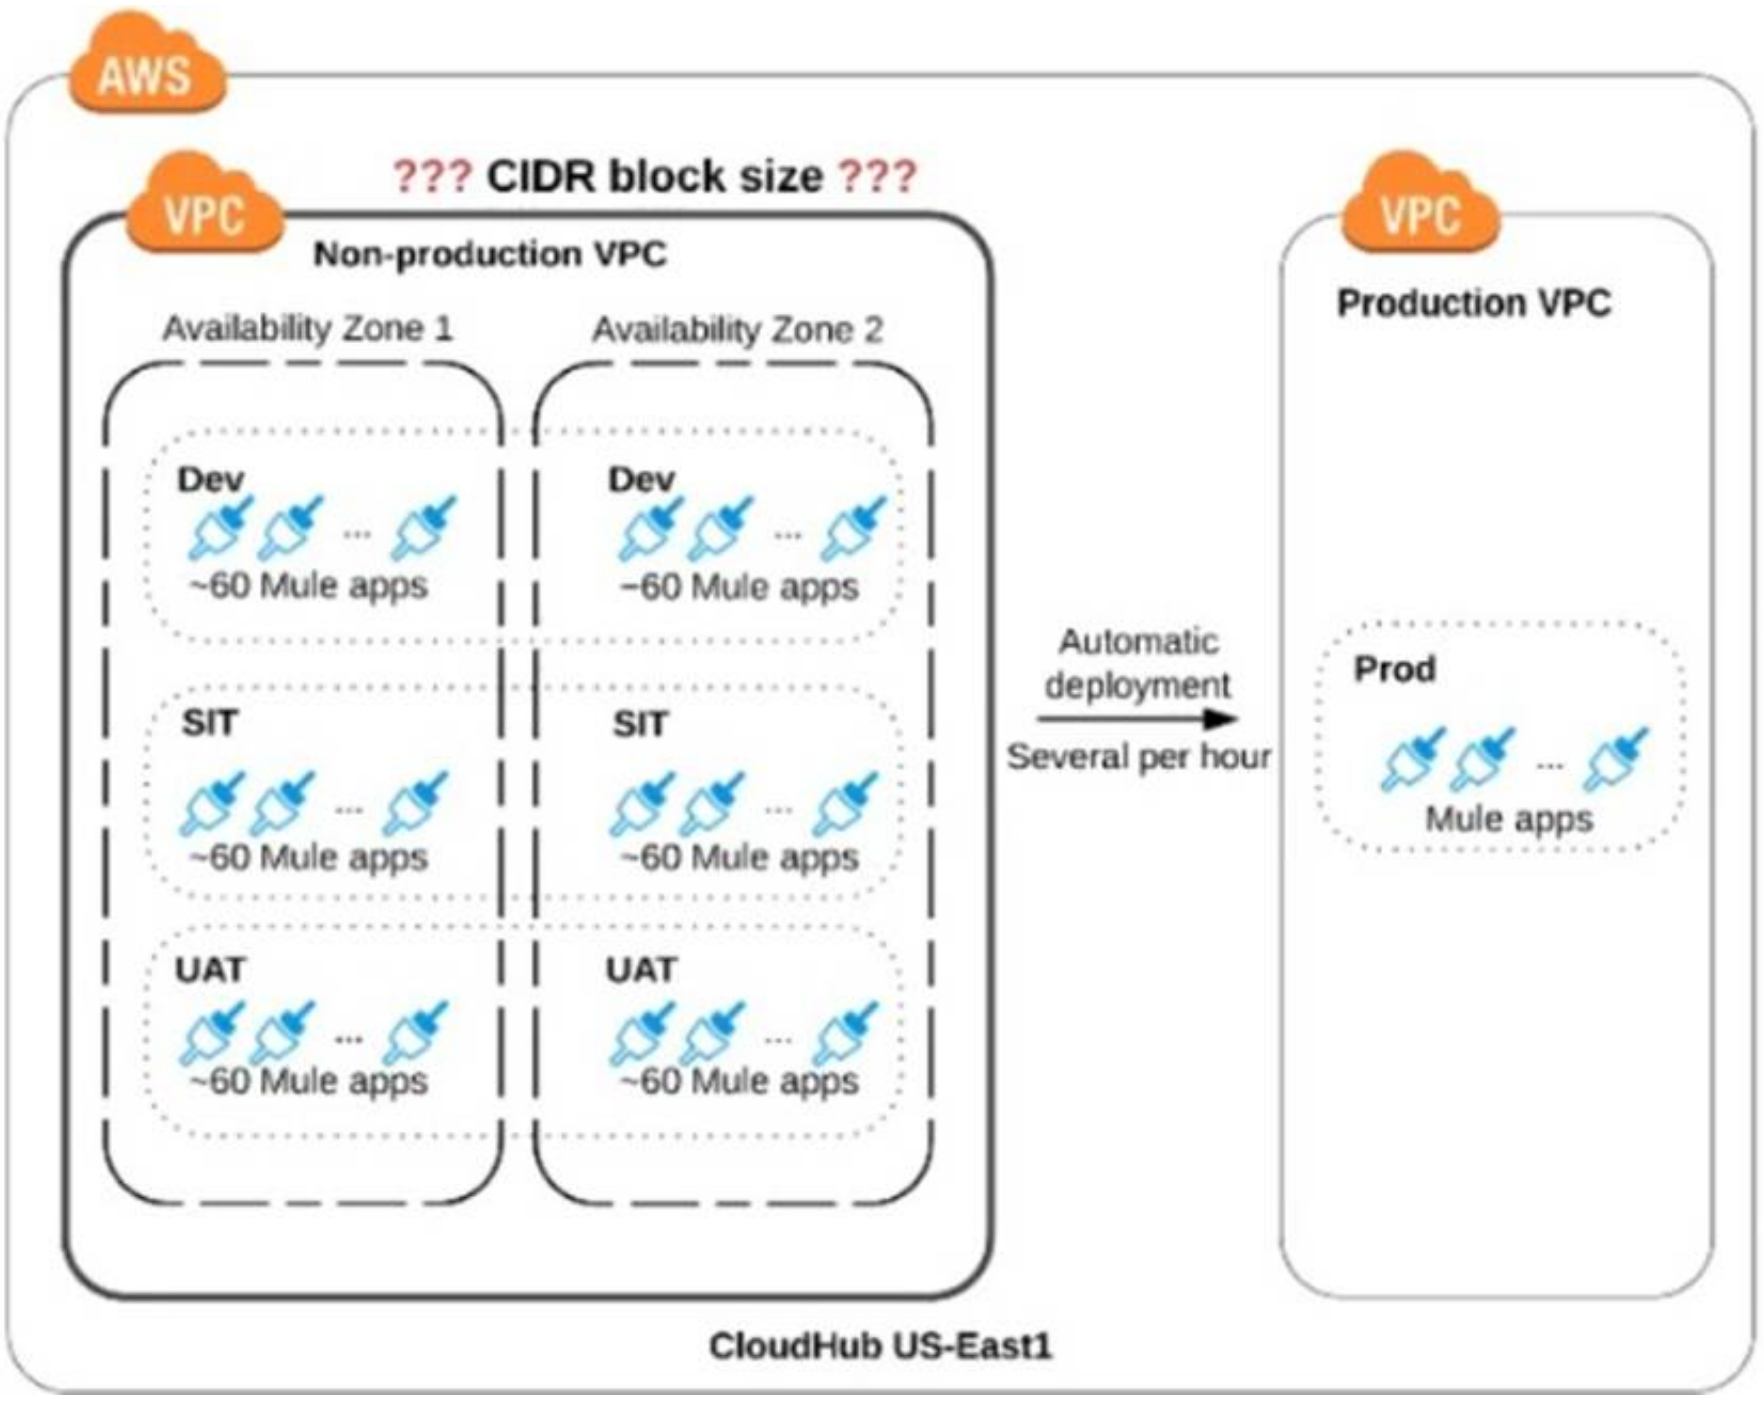

An organization is sizing an Anypoint VPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization currently has three non-production 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 90
A system administrator needs to determine when permissions were last changed for an Anypoint Platform user.
Which Anypoint Platform component should the administrator use to obtain this information?

A. Audit Logging
B. Anypoint Monitoring
C. Anypiont Studio
D. Mule Stack Traces

Answer: A

NEW QUESTION 92
In which order are the API Client, API Implementation, and API interface components called in a typical REST request?

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

Answer: C

NEW QUESTION 95
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 100
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 101
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 106
Which role is primarily responsible for building API implementation as part of a typical MuleSoft integration project?

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

Answer: A

NEW QUESTION 110
An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger of flow for processing a high-volume set of records from the source system and synchronize with the SaaS system using the Batch job scope
After processing 1000 records in a periodic synchronization of 1 lakh records, the replicas in which batch job instance was started went down due to unexpected failure in the runtime fabric environment
What is the consequence of losing the replicas that run the Batch job instance?

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

Answer: D

NEW QUESTION 112
Additional nodes are being added to an existing customer-hosted Mule runtime cluster to improve performance. Mule applications deployed to this cluster are invoked by API clients through a load balancer.
What is also required to carry out this change?

A. A new load balancer must be provisioned to allow traffic to the new nodes in a round- robin fashion
B. External monitoring tools or log aggregators must be configured to recognize the new nodes
C. API implementations using an object store must be adjusted to recognize the new nodes and persist to them
D. New firewall rules must be configured to accommodate communication between API clients and the new nodes

Answer: B

NEW QUESTION 115
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 117
A marketing organization is designing a Mule application to process campaign data. The Mule application will periodically check for a file in a SFTP location and process the records in the file. The size of the file can vary from 10MB to 5GB. Due to the limited availabiltty of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.
The application must transform and send different formats of this file to three different downstream SFTP locations.
What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

A. Use an in-memory repeatable stream
B. Use a file-stored non-repeatable stream
C. Use an in-memory non-repeatable stream
D. Use a file-stored repeatable stream

Answer: A

NEW QUESTION 118
A set of integration Mule applications, some of which expose APIs, are being created to enable a new business process. Various stakeholders may be impacted by this. These stakeholders are a combination of semi-technical users (who understand basic integration terminology and concepts such as JSON and XML) and technically skilled potential consumers of the Mule applications and APIs.
What Is an effective way for the project team responsible for the Mule applications and APIs being built to communicate with these stakeholders using Anypoint Platform and its supplied toolset?

A. Use Anypoint Design Center to implement the Mule applications and APIs and give the various stakeholders access to these Design Center projects, so they can collaborate and provide feedback
B. Create Anypoint Exchange entries with pages elaborating the integration design, including API notebooks (where applicable) to help the stakeholders understand and interact with the Mule applications and APIs at various levels of technical depth
C. Use Anypoint Exchange to register the various Mule applications and APIs and share the RAML definitions with the stakeholders, so they can be discovered
D. Capture documentation about the Mule applications and APIs inline within the Mule integration flows and use Anypoint Studio's Export Documentation feature to provide an HTML version of this documentation to the stakeholders

Answer: B

NEW QUESTION 121
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 125
An integration team follows MuleSoft??s recommended approach to full lifecycle API development.
Which activity should this team perform during the API implementation phase?

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

Answer: B

NEW QUESTION 130
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 133
An airline is architecting an API connectivity project to integrate its flight data into an online aggregation website. The interface must allow for secure communication high-performance and asynchronous message exchange.
What are suitable interface technologies for this integration assuming that Mulesoft fully supports these technologies and that Anypoint connectors exist for these interfaces?

A. AsyncAPI over HTTPSAMQP with RabbitMQ JSON/REST over HTTPS
B. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
C. CSV over FTP YAM L over TLS JSON over HTTPS
D. SOAP over HTTPS HOP over TLS gRPC over HTTPS

Answer: A

NEW QUESTION 134
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 138
A Mule application is built to support a local transaction for a series of operations on a single database. The mule application has a Scatter-Gather scope that participates in the local transaction.
What is the behavior of the Scatter-Gather when running within this local transaction?

A. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
B. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back
C. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
D. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back

Answer: A

NEW QUESTION 141
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 144
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 147
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 149
An integration team uses Anypoint Platform and follows MuleSoft's recommended approach to full lifecycle API development.
Which step should the team's API designer take before the API developers implement the AP! Specification?

A. Generate test cases using MUnit so the API developers can observe the results of running the API
B. Use the scaffolding capability of Anypoint Studio to create an API portal based on the API specification
C. Publish the API specification to Exchange and solicit feedback from the API's consumers
D. Use API Manager to version the API specification

Answer: C

NEW QUESTION 150
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 151
Refer to the exhibit.

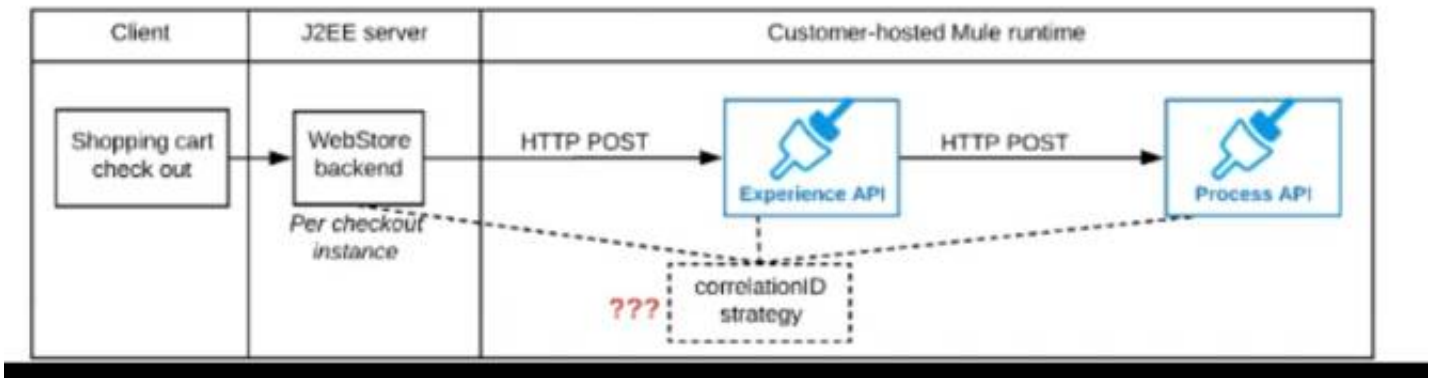

A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

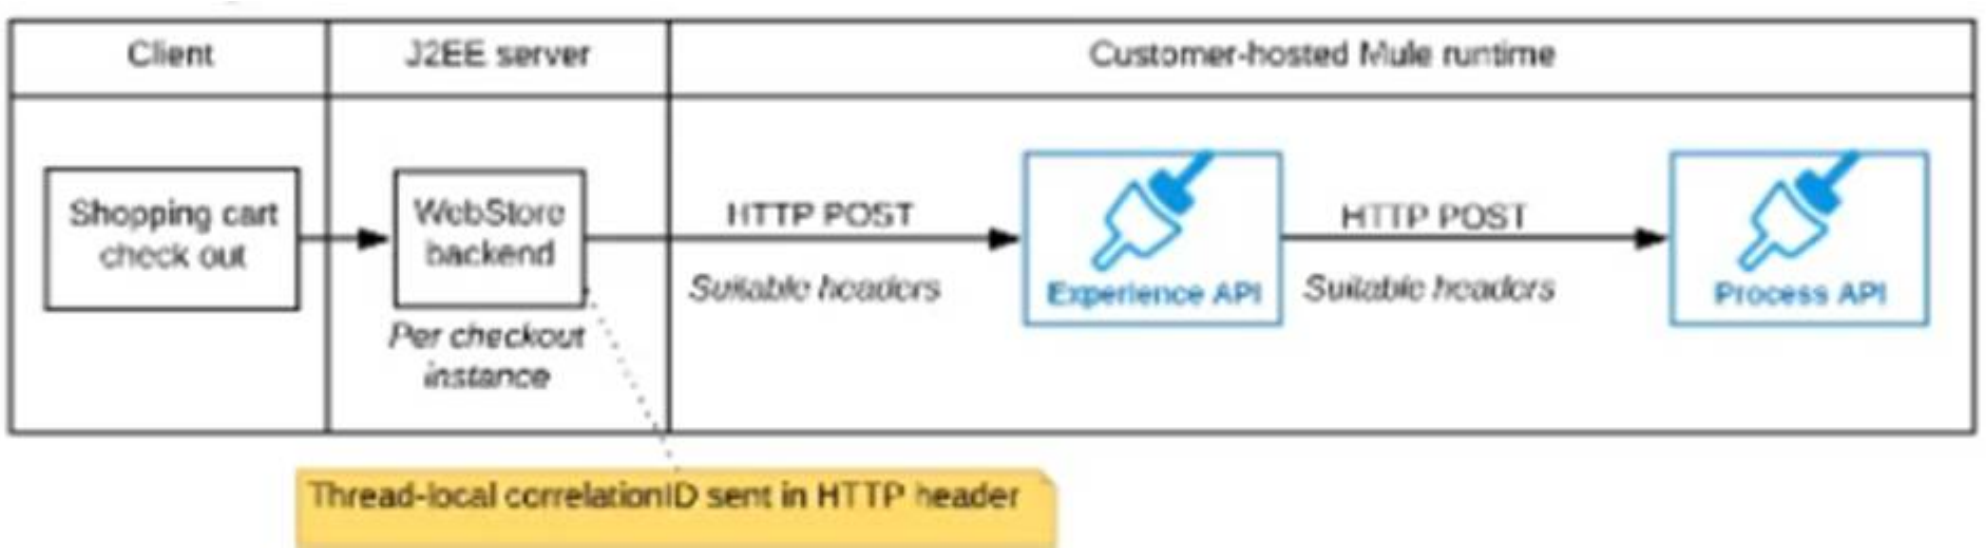

B) The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-It HTTP request header In each API invocation belonging to that checkout
No special code or configuration is included in the Experience API and Process API implementations to generate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

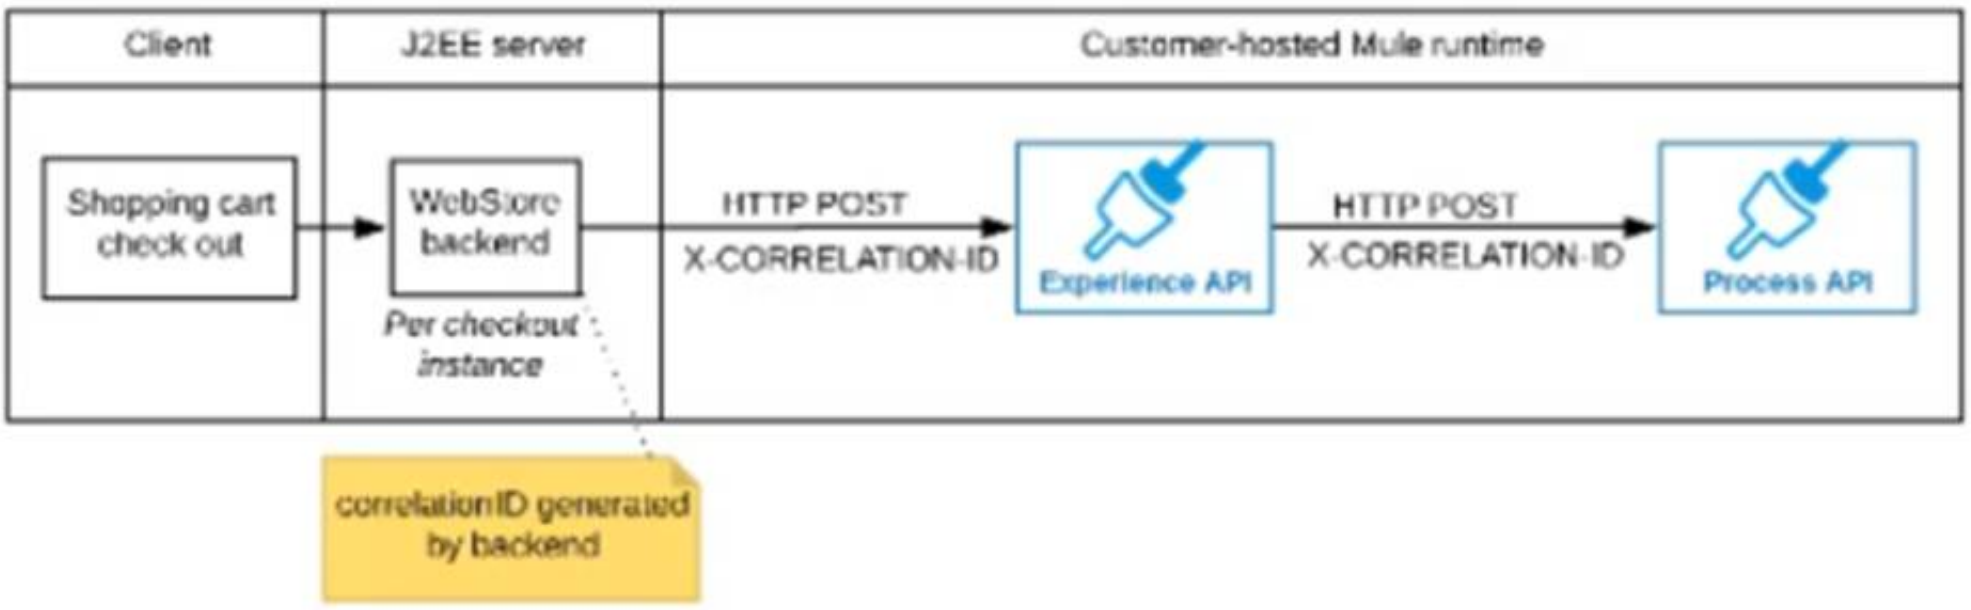

C) The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.
The Experience API implementation must be coded to also propagate the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

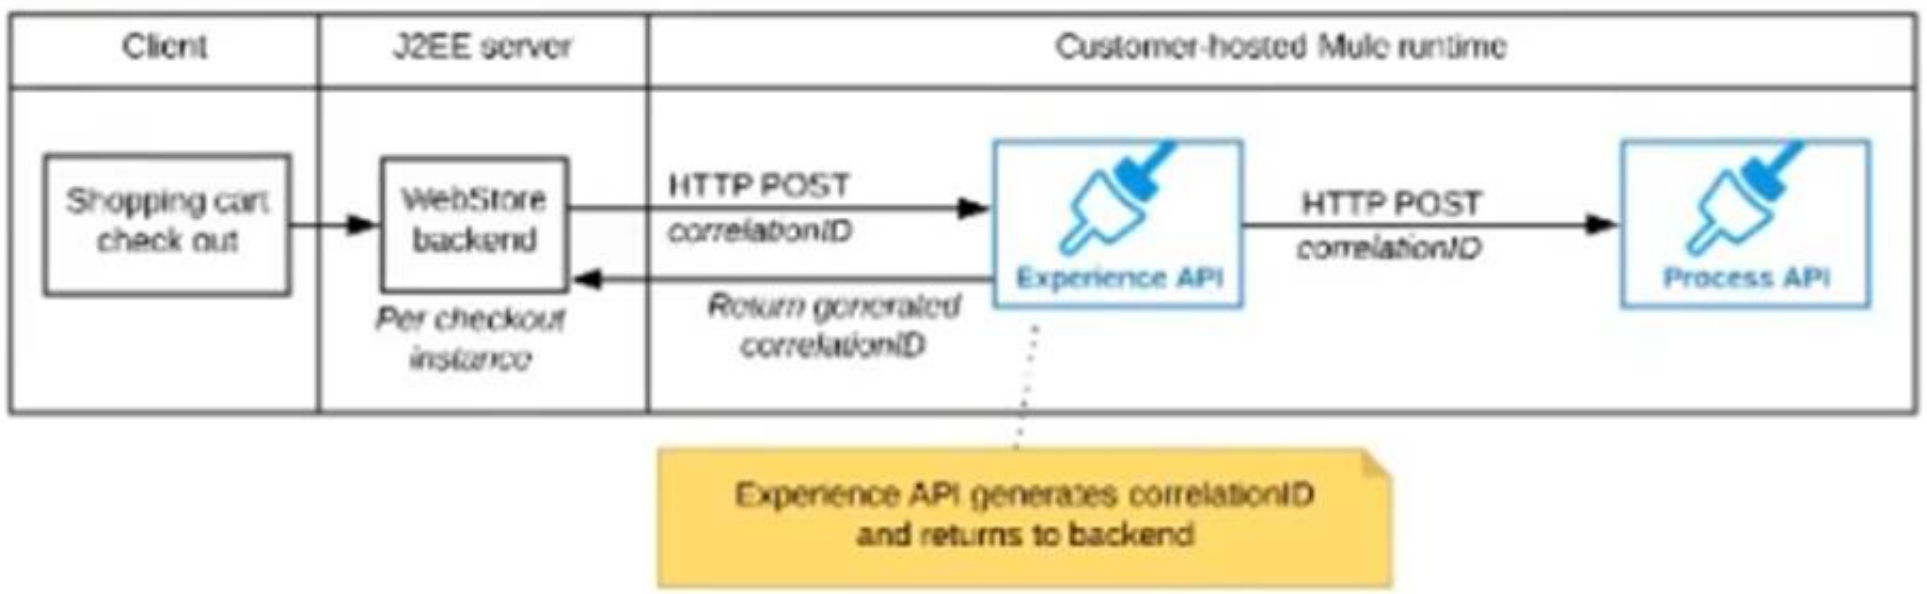

D) The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API
The Experience API and Process API implementations must be coded to receive the custom correlation ID In the HTTP requests and propagate It in suitable HTTP request headers

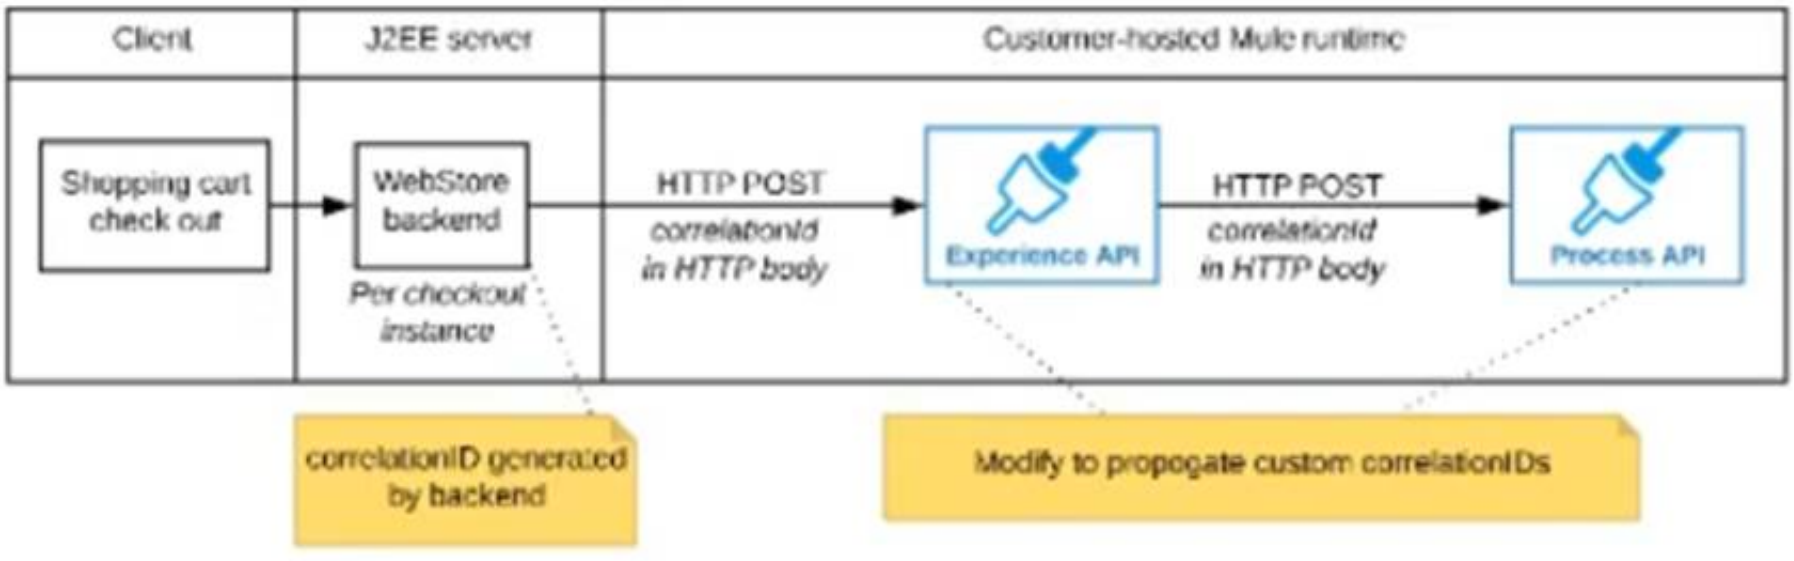

A. Option A
B. Option B
C. Option C

D. Option D

Answer: B

NEW QUESTION 154
Refer to the exhibit.

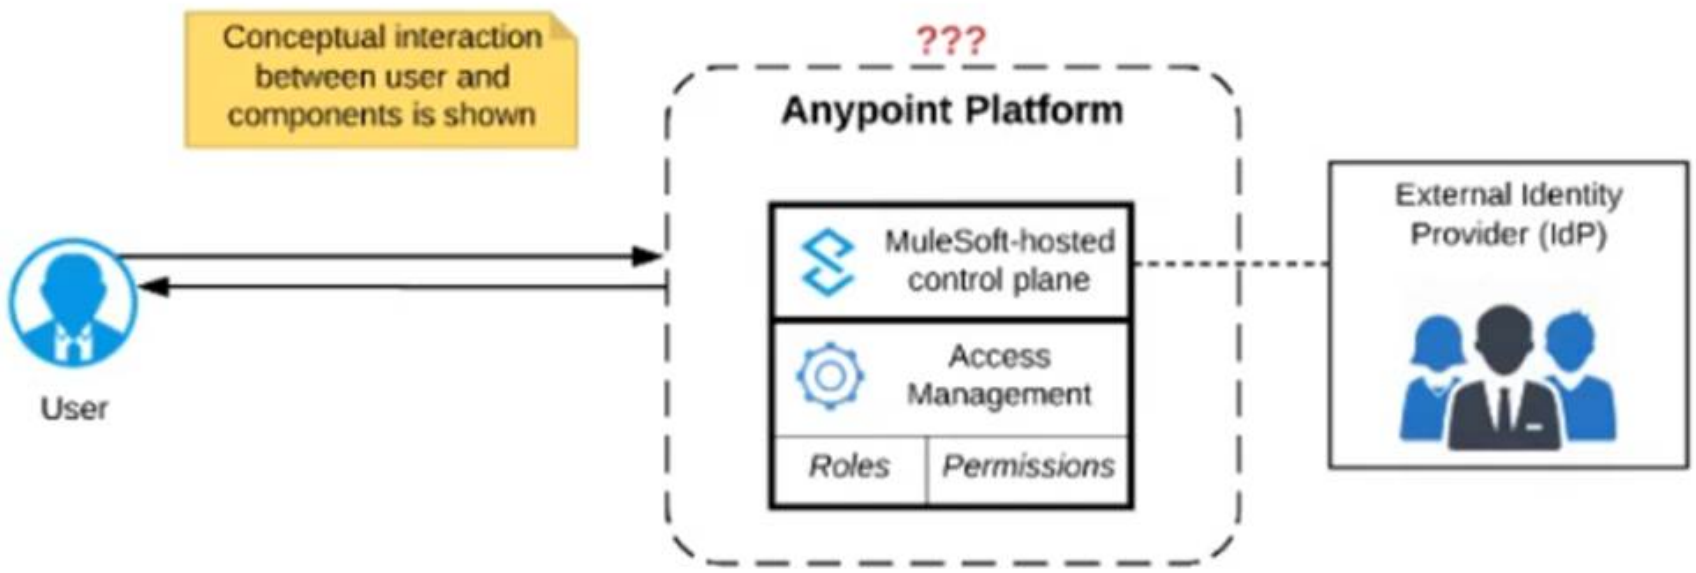

Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.
What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled 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 156
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 157
An integration architect is designing an API that must accept requests from API clients for both XML and JSON content over HTTP/1.1 by default.
Which API architectural style, when used for its intended and typical purposes, should the architect choose to meet these requirements?

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

Answer: C

NEW QUESTION 160
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 164
What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

A. To send Mule application log entries to Splunk
B. To change tog4j2 tog levels in Anypoint Runtime Manager without having to restart the Mule application
C. To log certain messages to a custom log category

D. To directly reference one shared and customized log4j2.xml file from multiple Mule applications

Answer: D

NEW QUESTION 165
A mule application designed to fulfil two requirements
a) Processing files are synchronously from an FTPS server to a back-end database using VM intermediary queues for load balancing VM events
b) Processing a medium rate of records from a source to a target system using batch job scope
Considering the processing reliability requirements for FTPS files, how should VM queues be configured for processing files as well as for the batch job scope if the application is deployed to Cloudhub workers?

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

Answer: A

NEW QUESTION 166
An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?

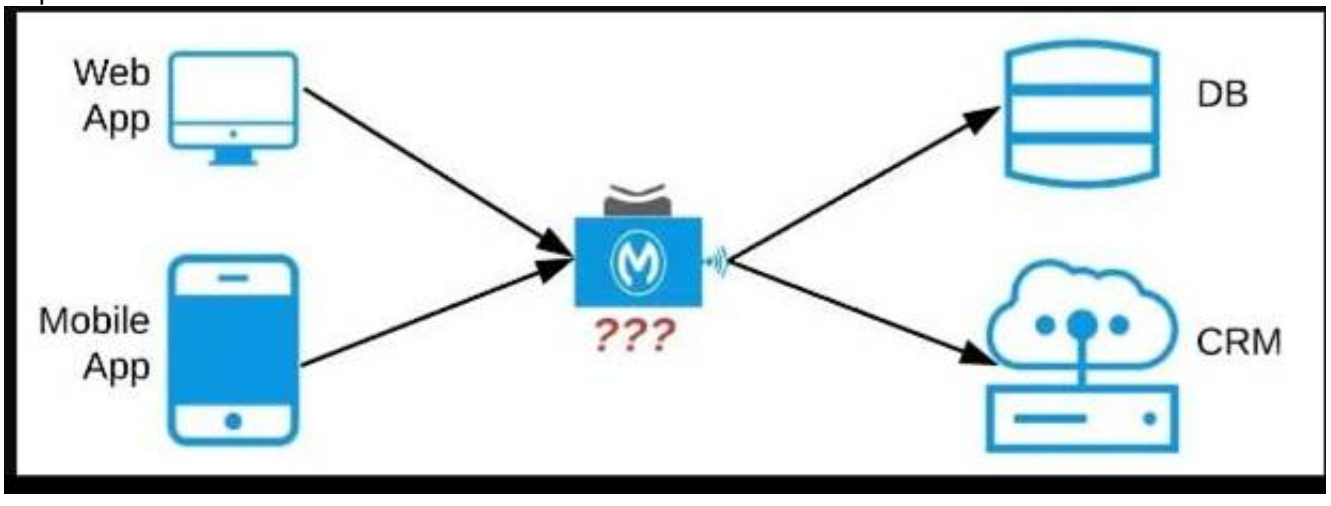

A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
B. One set of APIs (Experience API, Process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

Answer: C

NEW QUESTION 167
What is a core pillar of the MuleSoft Catalyst delivery approach?
A. Business outcomes
B. Technology centralization
C. Process thinking
D. Scope reduction

Answer: A

NEW QUESTION 170
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 172
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 173
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 174
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 178
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 180
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 185
As a part of project requirement, client will send a stream of data to mule application. Payload size can vary between 10mb to 5GB. Mule application is required to transform the data and send across multiple sftp servers. Due to the cost cuttings in the organization, mule application can only be allocated one worker with size of 0.2 vCore.
As an integration architect , which streaming strategy you would suggest to handle this scenario?

A. In-memory non repeatable stream
B. File based non-repeatable stream
C. In-memory repeatable stream
D. File based repeatable storage

Answer: D

NEW QUESTION 186
49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58. https://api.acme.com/partnefs/vl. The API implementation is estimated to require deployment to 5 CloudHub workers.

The retailer has obtained a public X.509 certificate for the name apl.acme.com, signed by a reputable CA, to be used as the server certificate.
Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

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

Answer: C

NEW QUESTION 191
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 196
A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction.
What is the behavior of the Scatter-Gather when running within this local transaction?

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

Answer: A

NEW QUESTION 201
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 204
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 206
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 208
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 209
An external REST client periodically sends an array of records in a single POST request to a Mule application API endpoint.
The Mule application must validate each record of the request against a JSON schema before sending it to a downstream system in the same order that it was received in the array
Record processing will take place inside a router or scope that calls a child flow. The child flow has its own error handling defined. Any validation or communication failures should not prevent further processing of the remaining records.
To best address these requirements what is the most idiomatic(used for it intended purpose) router or scope to used in the parent flow, and what type of error handler should be used in the child flow?

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

Answer: B

NEW QUESTION 212
The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure.
Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).
Given the company's current setup and requirements, what is the most idiomatic (used for its intended purpose) way to send Mule application logs to the external LMT?

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

Answer: A

NEW QUESTION 216
An organization uses one specific CloudHub (AWS) region for all CloudHub deployments. How are CloudHub workers assigned to availability zones (AZs) when the organization's Mule applications are deployed to CloudHub in that region?

A. Workers belonging to a given environment are assigned to the same AZ within that region.
B. AZs are selected as part of the Mule application's deployment configuration.
C. Workers are randomly distributed across available AZs within that region.
D. An AZ is randomly selected for a Mule application, and all the Mule application's CloudHub workers are assigned to that one AZ

Answer: C

NEW QUESTION 219
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 223
A company is modernizing its legal systems lo accelerate access lo applications and data while supporting the adoption of new technologies. The key to achieving this business goal is unlocking the companies' key systems and dala including microservices miming under Docker and kubernetes containers using apis.
Considering the current aggressive backlog and project delivery requirements the company wants to take a strategic approach in the first phase of its transformation projects by quickly deploying API's in mule runtime that are able lo scale, connect to on premises systems and migrate as needed.
Which runtime deployment option supports company's goals?

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

Answer: C

NEW QUESTION 228
Refer to the exhibit.

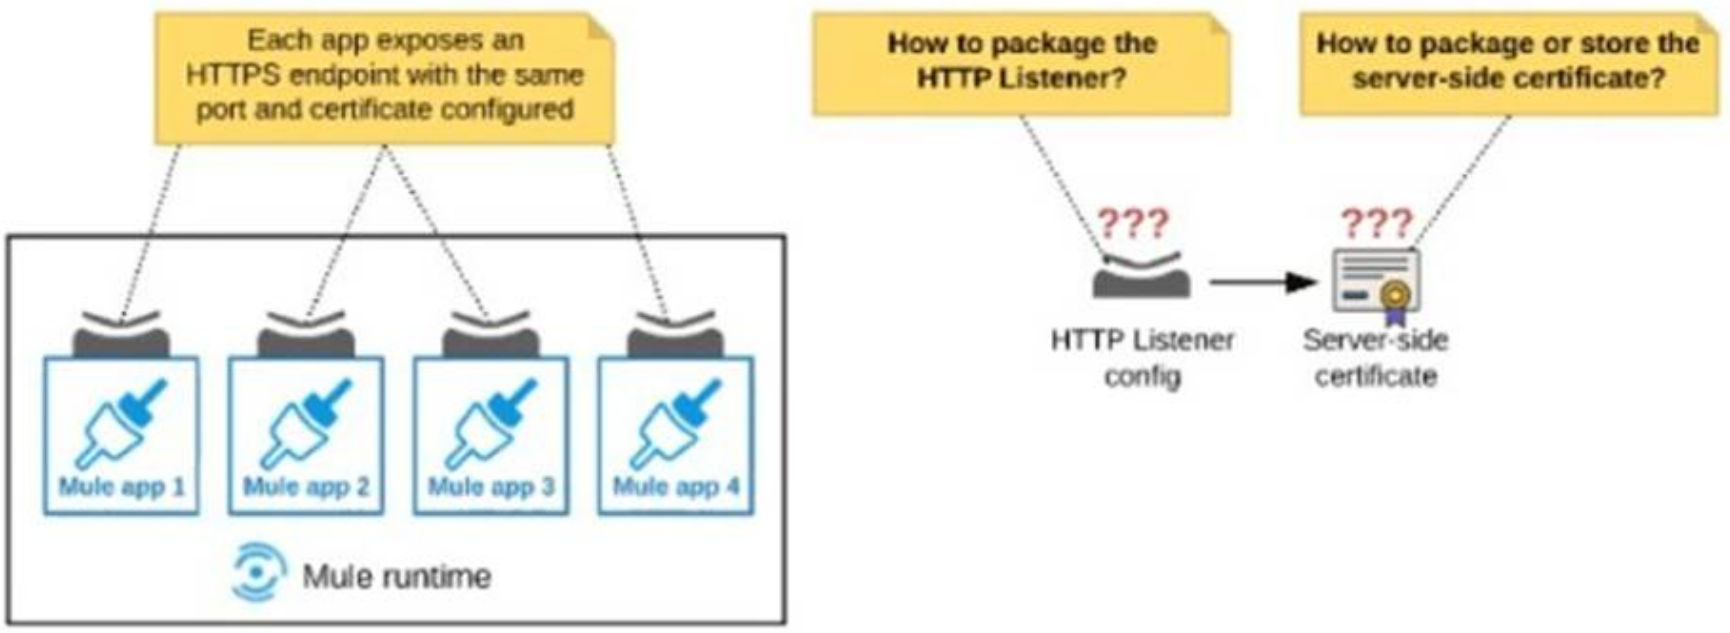

An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a server-side certificate that rotates often.
What is the most effective way to package the HTTP Listener and package or store 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 230
A mule application must periodically process a large dataset which varies from 6 GB lo 8 GB from a back-end database and write transform data lo an FTPS server using a properly configured bad job scope.
The performance requirements of an application are approved to run in the cloud hub 0.2 vCore with 8 GB storage capacity and currency requirements are met.
How can the high rate of records be effectively managed in this application?

A. Use streaming with a file storage repeatable strategy for reading records from the database and batch aggregator with streaming to write to FTPS
B. Use streaming with an in-memory reputable store strategy for reading records from the database and batch aggregator with streaming to write to FTPS
C. Use streaming with a file store repeatable strategy for reading records from the database and batch aggregator with an optimal size
D. Use streaming with a file store repeatable strategy reading records from the database and batch aggregator without any required configuration

Answer: A

NEW QUESTION 231
What Is a recommended practice when designing an integration Mule 4 application that reads a large XML payload as a stream?

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

Answer: C

NEW QUESTION 233
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 238
Refer to the exhibit.

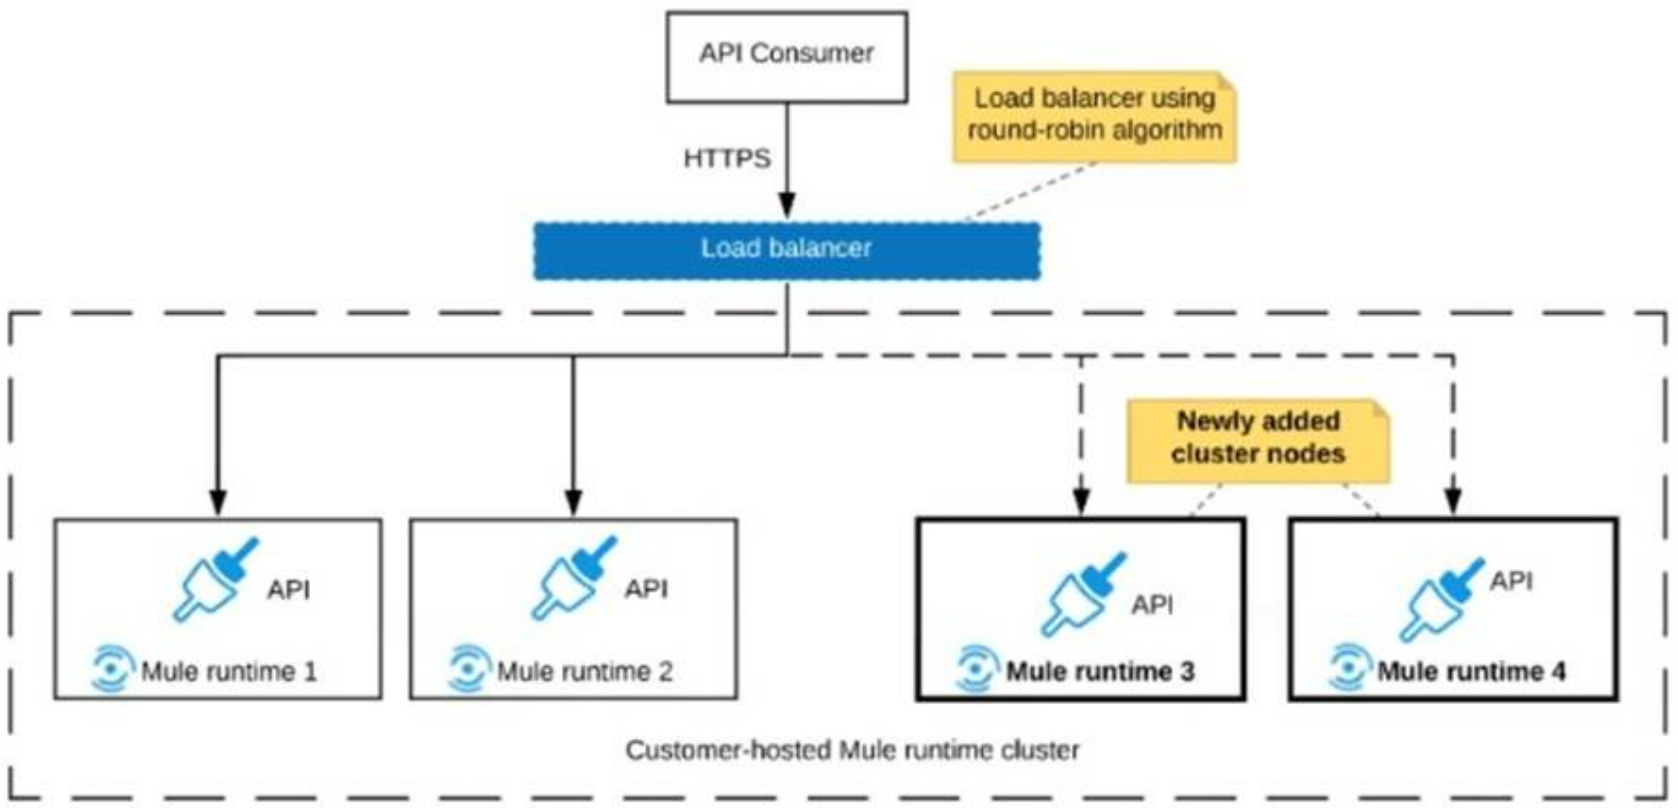

An organization uses a 2-node Mute runtime cluster to host one stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution.
Two additional nodes have been added to the cluster and the load balancer has been configured to recognize the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

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

Answer: D

NEW QUESTION 241
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 246
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 deplpoyed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates.
What type of restrictions exist on the types of certificates that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

A. Only MuleSoft-provided certificates are exposed.
B. Only customer-provided wildcard certificates are exposed.
C. Only customer-provided self-signed certificates are exposed.
D. Only underlying Mule application certificates are exposed (pass-through)

Answer: A

NEW QUESTION 250
A banking company is developing a new set of APIs for its online business. One of the critical API's is a master lookup API which is a system API. This master lookup API uses persistent object store. This API will be used by all other APIs to provide master lookup data.

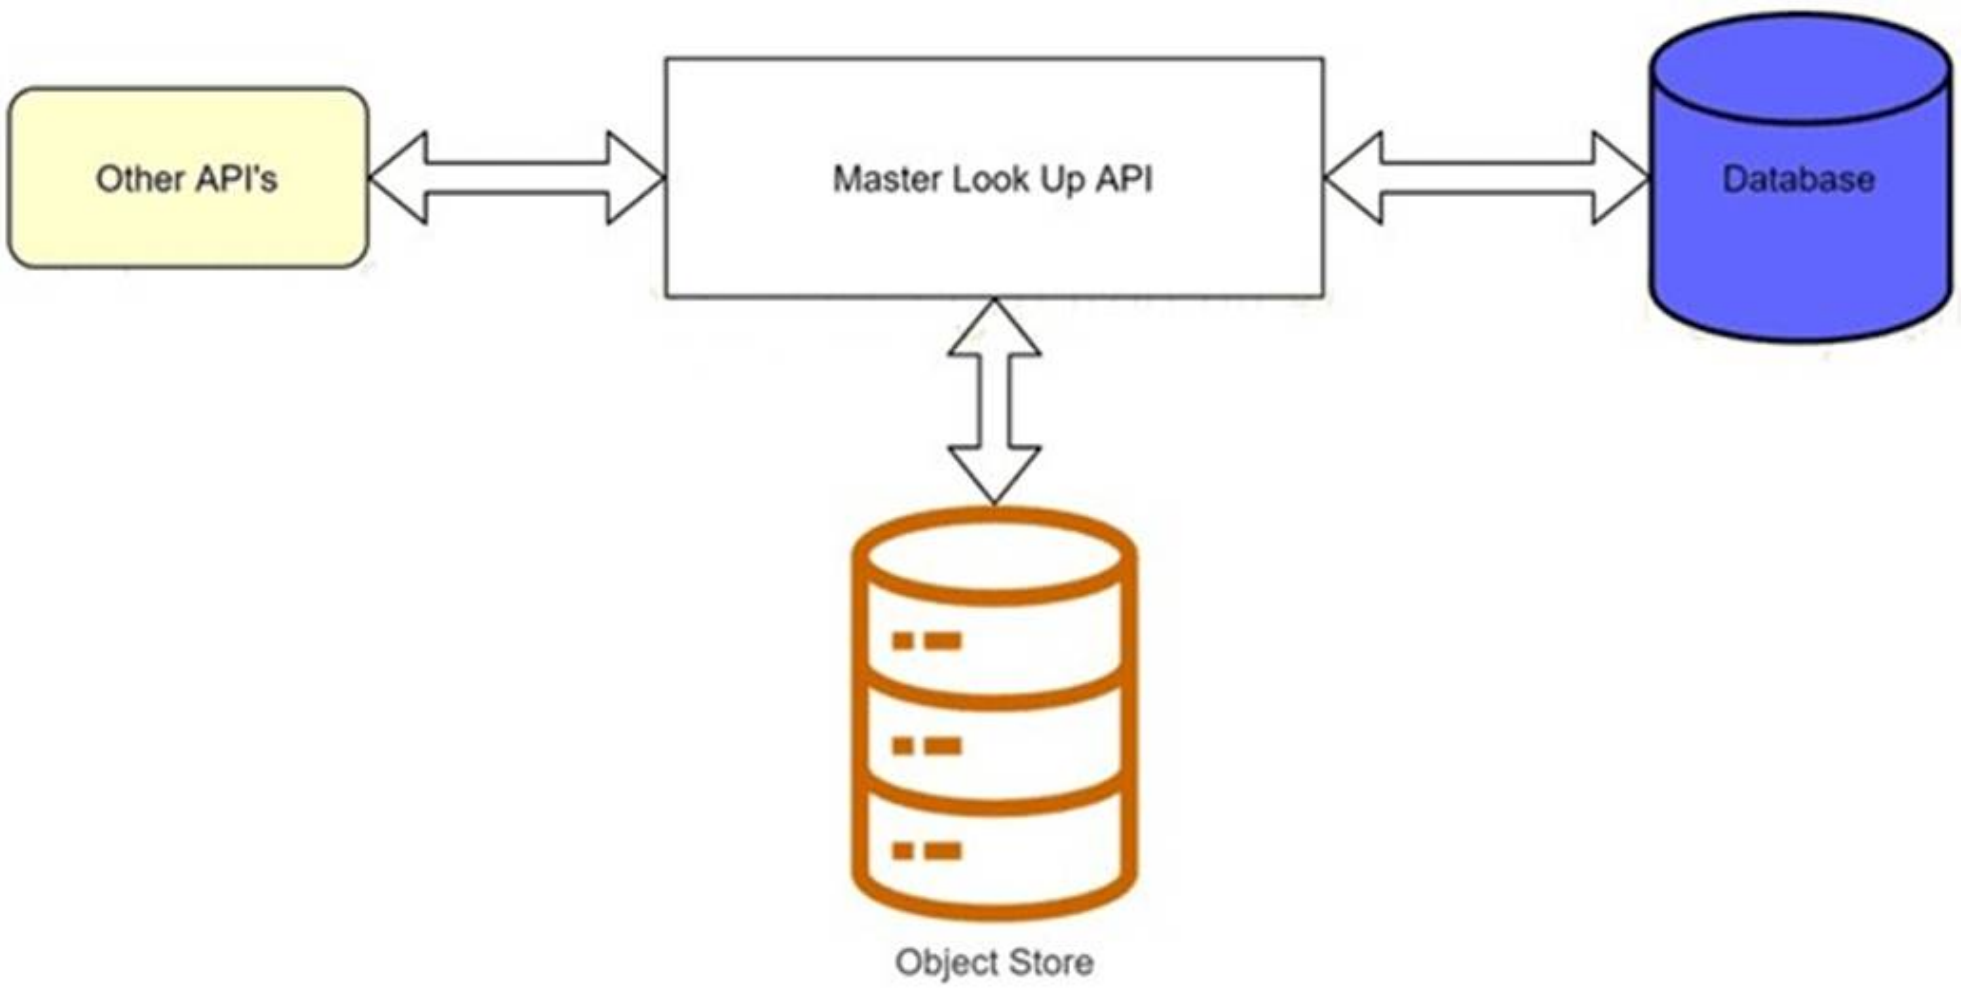

Master lookup API is deployed on two cloudhub workers of 0.1 vCore each because there is a lot of master data to be cached. Master lookup data is stored as a key value pair. The cache gets refreshed if they key is not found in the cache.
Doing performance testing it was observed that the Master lookup API has a higher response time due to database queries execution to fetch the master lookup data.
Due to this performance issue, go-live of the online business is on hold which could cause potential financial loss to Bank.
As an integration architect, which of the below option you would suggest to resolve performance issue?

A. Implement HTTP caching policy for all GET endpoints for the master lookup API and implement locking to synchronize access to object store
B. Upgrade vCore size from 0.1 vCore to 0,2 vCore
C. Implement HTTP caching policy for all GET endpoints for master lookup API
D. Add an additional Cloudhub worker to provide additional capacity

Answer: A

NEW QUESTION 255
An organization designing a hybrid, load balanced, single cluster production environment. Due to performance service level agreement goals, it is looking into running the Mule applications in an active-active multi node cluster configuration.
What should be considered when running its Mule applications in this type of environment?

A. All event sources, regardless of time , can be configured as the target source by the primary node in the cluster
B. An external load balancer is required to distribute incoming requests throughout the cluster nodes
C. A Mule application deployed to multiple nodes runs in an isolation from the other nodes in the cluster
D. Although the cluster environment is fully installed configured and running, it will not process any requests until an outage condition is detected by the primary node in the cluster.

Answer: B

NEW QUESTION 260
A developer is examining the responses from a RESTful web service that is compliant with the Mypertext Transfer Protocol (HTTP/1.1) a8 defined by the Internet Engineering Task Force (IETF).
In this HTTP/1.1-compliant web service, which class of HTTP response status codes should be specified to indicate when client requests are successfully received, understood, and accepted by the web service?

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

Answer: B

NEW QUESTION 264
Refer to the exhibit.

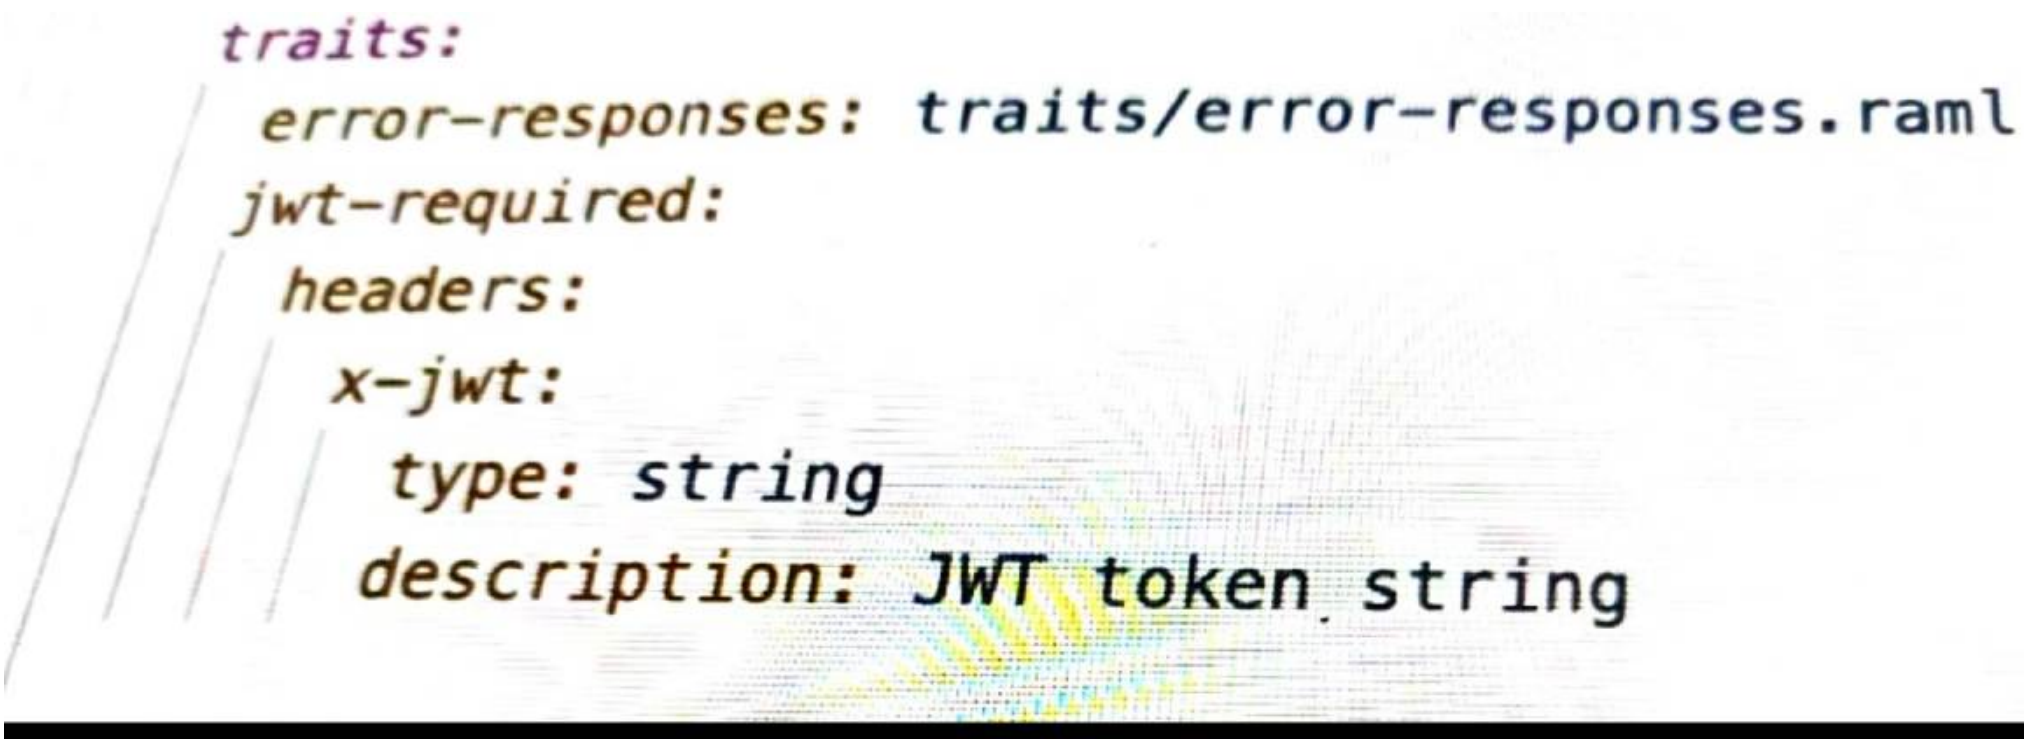

What is the type data format shown in the exhibit?

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

Answer: C

NEW QUESTION 268
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 273
What metrics about API invocations are available for visualization in custom charts using Anypoint Analytics?

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

Answer: C

NEW QUESTION 277
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 282
A MuteSoft developer must implement an API as a Mule application, run the application locally, and execute unit tests against the Running application.
Which Anypoint Platform component can the developer use to full all of these requirements?

A. API Manager
B. API Designer
C. Anypoint CLI
D. Anypoint Studio

Answer: D

NEW QUESTION 287
What is true about automating interactions with Anypoint Platform using tools such as Anypoint Platform REST API's, Anypoint CLI or the Mule Maven plugin?

A. By default, the Anypoint CLI and Mule Maven plugin are not included in the Mule runtime
B. Access to Anypoint Platform API;s and Anypoint CLI can be controlled separately thruough the roles and permissions in Anypoint platform, so that specific users can get access to Anypoint CLI while others get access to the platform API's
C. Anypoint Platform API's can only automate interactions with CloudHub while the Mule maven plugin is required for deployment to customer hosted Mule runtimes
D. API policies can be applied to the Anypoint platform API's so that only certain LOS's has access to specific functions

Answer: A

NEW QUESTION 292
An organization is designing Mule application which connects to a legacy backend. It has been reported that backend services are not highly available and experience downtime quite often. As an integration architect which of the below approach you would propose to achieve high reliability goals?

A. Alerts can be configured in Mule runtime so that backend team can be communicated when services are down
B. Until Successful scope can be implemented while calling backend API's
C. On Error Continue scope to be used to call in case of error again
D. Create a batch job with all requests being sent to backend using that job as per the availability of backend API's

Answer: B

NEW QUESTION 294
.....

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