Number: MCIA - Level 1
Passing Score: 800
Time Limit: 120
File Version: 21.0

Exam Code: MCIA - Level 1
Exam Name: MuleSoft Certified Integration Architect - Level 1

Exam A

QUESTION 1
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

Correct Answer: B
Section:
Explanation:
Correct answer is Untill Successful scope can be implemented while calling backend API's The Until Successful scope repeatedly triggers the scope's components (including flow references) until they all succeed or until a maximum number of retries is exceeded The scope provides option to control the max number of retries and the interval between retries The scope can execute any sequence of processors that may fail for whatever reason and may succeed upon retry

QUESTION 2
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 event For each Batch Step scope, all 25 records within a block are processed in parallel All 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 event For each Batch Step scope, all 25 records within a block are processed sequentially, one at a time All 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 event For each Batch Step scope, all 25 records within a block are processed sequentially, one record at a time All the records in a block must be completed before the block of 25 records is available to the next
Batch Step scope
D. The Batch Job scope processes multiple record blocks in parallel Each Batch Step scope is invoked with a batch of 25 records in the payload of the received Mule event For each Batch Step scope, all 4000 records are processed in parallel Individual records can jump ahead to the next Batch Step scope before the rest of the records finish processing in the current Batch Step scope

Correct Answer: A
Section:
Explanation:
Reference: https://docs.mulesoft.com/mule-runtime/4.4/batch-processing-concept

QUESTION 3
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

Correct Answer: D
Section:
Explanation:
Reference: https://docs.mulesoft.com/jms-connector/1.8/jms-performance

QUESTION 4
A stock broking company makes use of CloudHub VPC to deploy Mule applications. Mule application needs to connect to a database application in the customers on-premises corporate data center and also to a Kafka cluster running in
AWS VPC.
How is access enabled for the API to connect to the database application and Kafka cluster securely?

A. Set up a transit gateway to the customers on-premises corporate datacenter to AWS VPC
B. Setup AnyPoint VPN to the customer's on-premise corporate data center and VPC peering with AWS VPC
C. Setup VPC peering with AWS VPC and the customers devices corporate data center
D. Setup VPC peering with the customers onto my service corporate data center and Anypoint VPN to AWS VPC

Correct Answer: B
Section:

QUESTION 5
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

Correct Answer: D
Section:

QUESTION 6
An insurance company is using a CloudHub 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 to detect the policy deletion and send an email to operations team the SMTP connector

Correct Answer: D
Section:

QUESTION 7
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

Correct Answer: C
Section:

QUESTION 8
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.

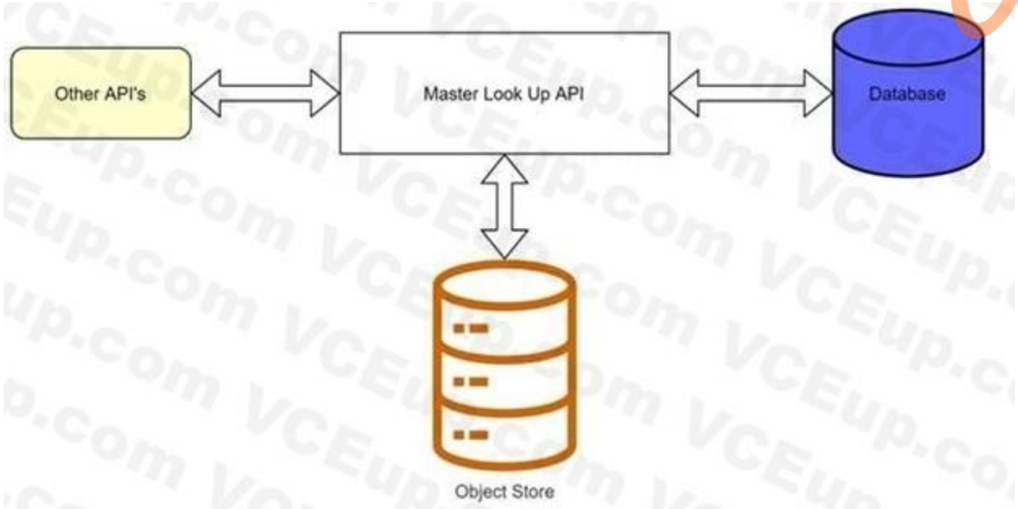

Master lookup API is deployed on two cloudhub workers of 0.1 vCore each because there is a lot of master data to be cached. Master lookup data is stored as a key value pair. The cache gets refreshed if they key is not found in the cache.
Doing performance testing it was observed that the Master lookup 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 implementlocking 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

Correct Answer: A
Section:

QUESTION 9
An XA transaction Is being configured that involves a JMS connector listening for Incoming JMS messages. What is the meaning of the timeout attribute of the XA transaction, and what happens after the timeout expires?

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

Correct Answer: C
Section:
Explanation:
* Setting a transaction timeout for the Bitronix transaction manager
ï Set the transaction timeout either
ñ In wrapper.conf
ñ In CloudHub in the Properties tab of the Mule application deployment ?
ï The default is 60 secs. It is defined as
mule.bitronix.transactiontimeout = 120
* This property defines the timeout for each transaction created for this manager.
If the transaction has not terminated before the timeout expires it will be automatically rolled back.
---------- Additional Info around Transaction Management:
Bitronix is available as the XA transaction manager for Mule applications ?
ï To use Bitronix, declare it as a global configuration element in the Mule application
<bti:transaction-manager />
ï Each Mule runtime can have only one instance of a Bitronix transaction manager, which is shared by all Mule applications
ï For customer-hosted deployments, define the XA transaction manager in a Mule domain
ñ Then share this global element among all Mule applications in the Mule runtime

Transaction Management

Characteristics	Local Transactions	Extended Architecture (XA) Transactions
Connector Requisite 1	All operations inside the transaction must belong to same Connector.	Operations inside the transaction may belong to different Connectors
Connector Requisite 2	Connectors may not be XA enabled	Connectors must be XA enabled
Connector Requisite 3	Connectors should use single config reference	Connectors may use multiple config references
Available resources	JMS, VM, JDBC	JMS, VM, JDBC
Uses Two Phase Commit (2PC)	No	Yes
DB Operations	Perform database operation to only one database resource	Perform database operation to one or more transactional resource
Supports Nested Transactions	Does not support nested transactions.	Supports nested transactions.
Bitronix is available	No	Yes
A.C.I.D Properties	No	Yes
Performance	Better than XA	Latency Increases
Thread Pooling	BLOCKING_IO	BLOCKING_IO
Recovery is case of system failure	No	Using Bitronix

QUESTION 10
Refer to the exhibit.

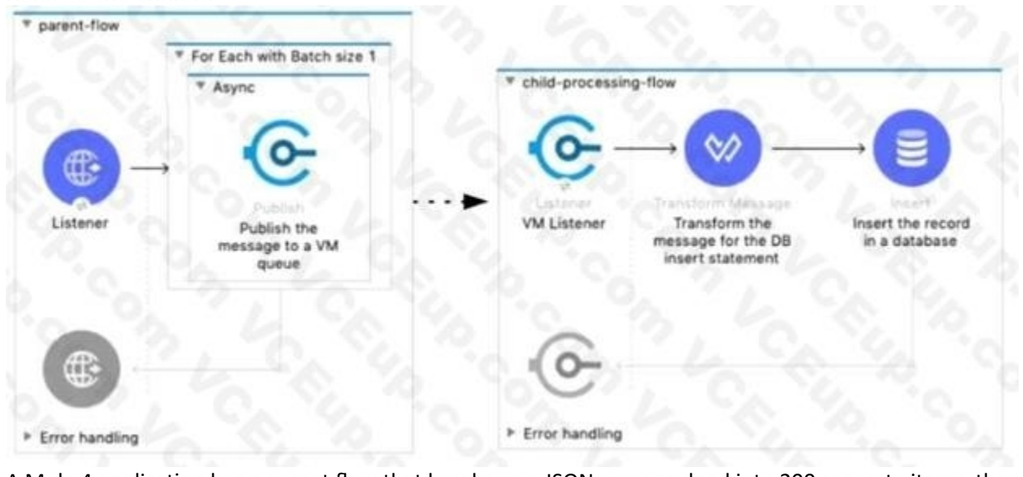

A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items, then sends each item one at a time inside an Async scope to a VM queue.
A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing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

A. EACH item VM message is processed AT MOST ONCE by ONE CloudHub worker, with workers chosen in a deterministic round-robin fashion Each of the four CloudHub workers can be expected to process 1/4 of the Item VM messages (about 50 items)

B. EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker Each of the four CloudHub workers can be expected to process some item VM messages
C. ALL Item VM messages are processed AT LEAST ONCE by the SAME CloudHub worker where the parent flow was invoked This one CloudHub worker processes ALL 200 item VM messages
D. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub worker This one CloudHub worker processes ALL 200 item VM messages

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

QUESTION 11
Refer to the exhibit.

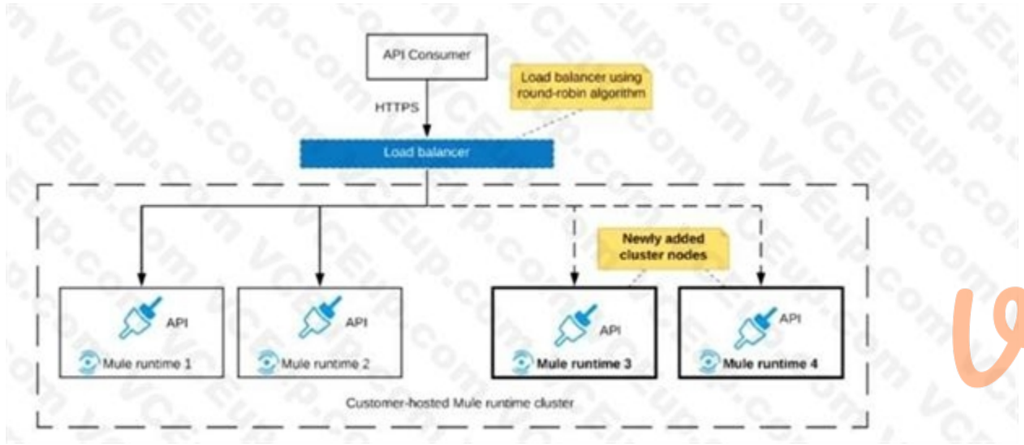

An organization uses a 2-node Mute runtime cluster to host one stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution.
Two additional nodes have been added to the cluster and the load balancer has been configured to recognize the 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

Correct Answer: D
Section:

QUESTION 12
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

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

QUESTION 13
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

Correct Answer: D
Section:
Explanation:
Correct answer is IBM: CICS CRM: SOAP
* Within Anypoint Exchange, MuleSoft offers the IBM CICS connector. Anypoint Connector for IBM CICS Transaction Gateway (IBM CTG Connector) provides integration with back-end CICS apps using the CICS Transaction Gateway.
* Anypoint Connector for Salesforce Marketing Cloud (Marketing Cloud Connector) enables you to connect to the Marketing Cloud API web services (now known as the Marketing Cloud API), which is also known as the Salesforce Marketing
Cloud. This connector exposes convenient operations via SOAP for exploiting the capabilities of Salesforce Marketing Cloud.

QUESTION 14
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

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

QUESTION 15

Refer to the exhibit.

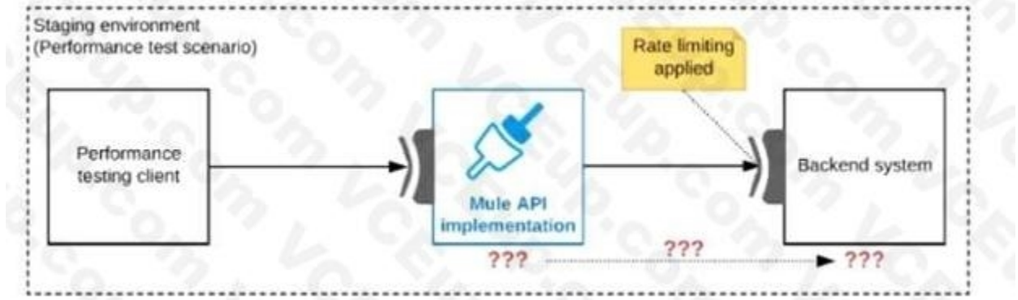

One of the backend systems invoked by an API implementation enforces rate limits on the number of requests a particular client can make. Both the backend system and the API implementation are deployed to several non-production environments in addition to production.
Rate limiting of the backend system 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 characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests
B. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
C. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situation. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
D. Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale

Correct Answer: A
Section:
Explanation:
Correct answer is Create a mocking service that replicates the backend system's production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests
* MUnit is for only Unit and integration testing for APIs and Mule apps. Not for performance Testing, even if it has the ability to Mock the backend.
* Bypassing the backend invocation defeats the whole purpose of performance testing. Hence it is not a valid answer.
* Scaled down performance tests cant be relied upon as performance of API's is not linear against load.

QUESTION 16
An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments.
What must the testing team do before they can start integration testing the API in the Staging environment?

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

Correct Answer: B
Section:
Explanation:
* It's mentioned that the API is governed by a Client ID Enforcement policy in all environments.
* Client ID Enforcement policy allows only authorized applications to access the deployed API implementation.
* Each authorized application is configured with credentials: client_id and client_secret.
* At runtime, authorized applications provide the credentials with each request to the API implementation.
MuleSoft Reference: https://docs.mulesoft.com/api-manager/2.x/policy-mule3-client-id-basedpolicies

QUESTION 17

What requires configuration of both a key store and a trust store for an HTTP Listener?

A. Support for TLS mutual (two-way) authentication with HTTP clients
B. Encryption of requests to both subdomains and API resource endpoints fhttPs://aDi.customer.com/ and https://customer.com/api)
C. Encryption of both HTTP request and HTTP response bodies for all HTTP clients
D. Encryption of both HTTP request header and HTTP request body for all HTTP clients

Correct Answer: A
Section:
Explanation:
1 way SSL : The server presents its certificate to the client and the client adds it to its list of trusted certificate. And so, the client can talk to the server.
2-way SSL: The same principle but both ways. i.e. both the client and the server has to establish trust between themselves using a trusted certificate. In this way of a digital handshake, the server needs to present a certificate to authenticate itself to client and client has to present its certificate to server.
* TLS is a cryptographic protocol that provides communications security for your Mule app.
* TLS offers many different ways of exchanging keys for authentication, encrypting data, and guaranteeing message integrity Keystores and Truststores Truststore and keystore contents differ depending on whether they are used for clients or servers:
For servers: the truststore contains certificates of the trusted clients, the keystore contains the private and public key of the server. For clients: the truststore contains certificates of the trusted servers, the keystore contains the private and public key of the client.
Adding both a keystore and a truststore to the configuration implements two-way TLS authentication also known as mutual authentication.
* in this case, correct answer is Support for TLS mutual (two-way) authentication with HTTP clients.

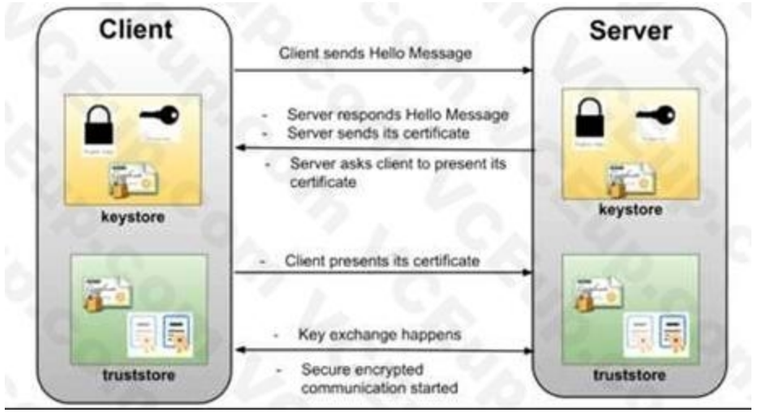

QUESTION 18
A retailer is designing a data exchange interface to be used by its suppliers. The interface must support secure communication over the public internet. The interface must also work with a wide variety of programming languages and IT systems used by suppliers.
What are suitable interface 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

Correct Answer: C
Section:
Explanation:
* As per definition of API by Mulesoft , it is Application Programming Interface using HTTP-basedprotocols. Non-HTTP-based programmatic interfaces are not APIs.
* HTTP-based programmatic interfaces are APIs even if they don't use REST or JSON. Henceimplementation based on Java RMI, CORBA/IIOP, raw TCP/IP interfaces are not API's as they are notusing HTTP.

* One more thing to note is FTP was not built to be secure. It is generally considered to be an insecure protocol because it relies on clear-text usernames and passwords for authentication and does not use encryption.
* Data sent via FTP is vulnerable to sniffing, spoofing, and brute force attacks, among other basic attack methods.
Considering the above points only correct option is
-XML over ActiveMQ
- XML over SFTP
- XML/REST over HTTPS

QUESTION 19
An organization currently uses a multi-node Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications. The organization is planning to transition to a deployment model based on Docker containers in a Kubernetes cluster. The organization has already created a standard Docker image containing a Mule runtime and all required dependencies (including a JVM), but excluding the Mule application itself.
What is an expected outcome of this transition to container-based Mule application deployments?

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

Correct Answer: R, E, Q, U, I, R, E, D, R, E, D, E, S, I, G, N, O, F, M, U, L, E, A, P, P, L, I, C, A, T, I, O, N, S, T, O, F, O, L, L, O, W, M, I, C, R, O, S, E, R, V, I, C, E
Section:
Explanation:
* Organization can continue using existing load balancer even if backend application changes are there. So option A is ruled out.
* As Mule runtime is within their datacenter, this model is RTF and not PCE. So option C is ruled out.
Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications -- This mean PCE or Hybird not RTF - Also mentioned in Question is that - Mule runtime is hosting several Mule Application, so that also rules out RTF and as for hosting multiple Application it will have Domain project which need redesign to make it microservice architecture ----------
---
Answer: Required redesign of Mule applications to follow microservice
Explanation:architecture principles

QUESTION 20
A team would like to create a project skeleton that developers can use as a starting point when creating API Implementations with Anypoint Studio. This skeleton should help drive consistent use of best practices within the team.
What type of Anypoint Exchange artifact(s) should be added to Anypoint Exchange to publish the project skeleton?

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

Correct Answer: D
Section:
Explanation:
* Sharing Mule applications as templates is a great way to share your work with other people who are in your organization in Anypoint Platform. When they need to build a similar application they can create the mule application using the template project from Anypoint studio.
* Anypoint Templates are designed to make it easier and faster to go from a blank canvas to a production application. They're bit for bit Mule applications requiring only Anypoint Studio to build and design, and are deployable both on-premises and in the cloud.
* Anypoint Templates are based on five common data Integration patterns and can be customized and extended to fit your integration needs. So even if your use case involves different endpoints or connectors than those included in the template, they still offer a great starting point.
Some of the best practices while creating the template project: - Define the common error handler as part of template project, either using pom dependency or mule config file - Define common logger/audit framework as

part of the template project - Define the env specific properties and secure properties file as per the requirement - Define global.xml for global configuration - Define the config file for connector configuration like Http,Salesforce,File,FTP etc - Create separate folders to create DWL,Properties,SSL certificates etc - Add the dependency and configure the pom.xml as per the business need - Configure the mule-artifact.json as per the business need

QUESTION 21
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

Correct Answer: D
Section:
Explanation:
* Correct answer is To directly reference one shared and customized log4j2.xml file from multiple Mule applications. Key word to note in the answer is directly.
* By default, CloudHub replaces a Mule application's log4j2.xml file with a CloudHub log4j2.xml file.
This specifies the CloudHub appender to write logs to the CloudHub logging service.
* You cannot modify CloudHub log4j2.xml file to add any custom appender. But there is a process in order to achieve this. You need to raise a request on support portal to disable CloudHub provided Mule application log4j2 file.

* Once this is done , Mule application's log4j2.xml file is used which you can use to send/export application logs to other log4j2 appenders, such as a custom logging system MuleSoft does not own any responsibility for lost logging data due to misconfiguration of your own log4j appender if it happens by any chance.

* One more difference between customer-hosted Mule runtimes and CloudHub deployed mule instance is that
- CloudHub system log messages cannot be sent to external log management system without installing custom CH logging configuration through support
- where as Customer-hosted runtime can send system and application log to external log management system MuleSoft Reference:
https://docs.mulesoft.com/runtime-manager/viewing-log-data
https://docs.mulesoft.com/runtime-manager/custom-log-appender

QUESTION 22
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

Correct Answer: T, H, E, J, M, S, C, O, N, N, E, C, T, O, R, S, U, P, P, O, R, T, S, B, O, T, H, S, E, N, D, I, N, G, A, N, D, R, E, C, E, I, V, I, N, G, O, F, J, M, S
Section:
Explanation:
* To send message or receive JMS (Java Message Service) message no separate network connection need to be established. So option A, C and D are ruled out.
Answer: The JMS connector supports both sending and receiving of JMS
Explanation:messages over the protocol determined by the JMS provider.
* JMS Connector enables sending and receiving messages to queues and topics for any message service that implements the JMS specification.
* JMS is a widely used API for message-oriented middleware.
* It enables the communication between different components of a distributed application to be loosely coupled, reliable, and asynchronous.
MuleSoft Doc Reference: https://docs.mulesoft.com/jms-connector/1.7/

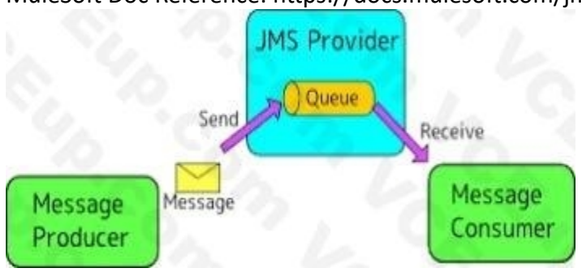

QUESTION 23
Refer to the exhibit.

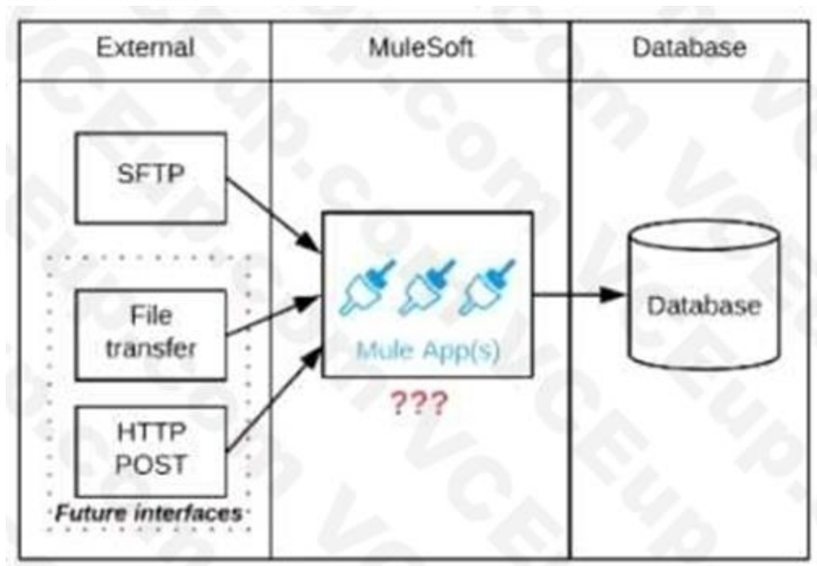

A business process involves the receipt of a file from an external vendor over SFTP. The file needs to be parsed and its content processed, validated, and ultimately persisted to a database. The delivery mechanism is expected to change in the future as more vendors send similar files using other mechanisms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

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

Correct Answer: C
Section:
Explanation:
* Scatter-Gather is used for parallel processing, to improve performance. In this scenario, input files are coming from different vendors so mostly at different times. Goal here is to minimize the impact of future change. So scatter Gather is not the correct choice.
* If we use 1 API to receive all files from different Vendors, any new vendor addition will need changes to that 1 API to accommodate new requirements. So Option A and C are also ruled out.
* Correct answer is Create an API that receives the file and invokes a Process API with the data contained in the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed. Answer to this question lies in the API led connectivity approach.
* API-led connectivity is a methodical way to connect data to applications through a series of reusable and purposeful modern APIs that are each developed to play a specific role ñ unlock data from systems, compose data into processes, or deliver an experience. System API : System API tier, which provides consistent, managed, and secure access to backend systems. Process APIs : Process APIs take core assets and combines them with some business logic to create a higher level of value.
Experience APIs : These are designed specifically for consumption by a specific end-user app or device.
So in case of any future plans , organization can only add experience API on addition of new Vendors, which reuse the already existing process API. It will keep impact minimal.

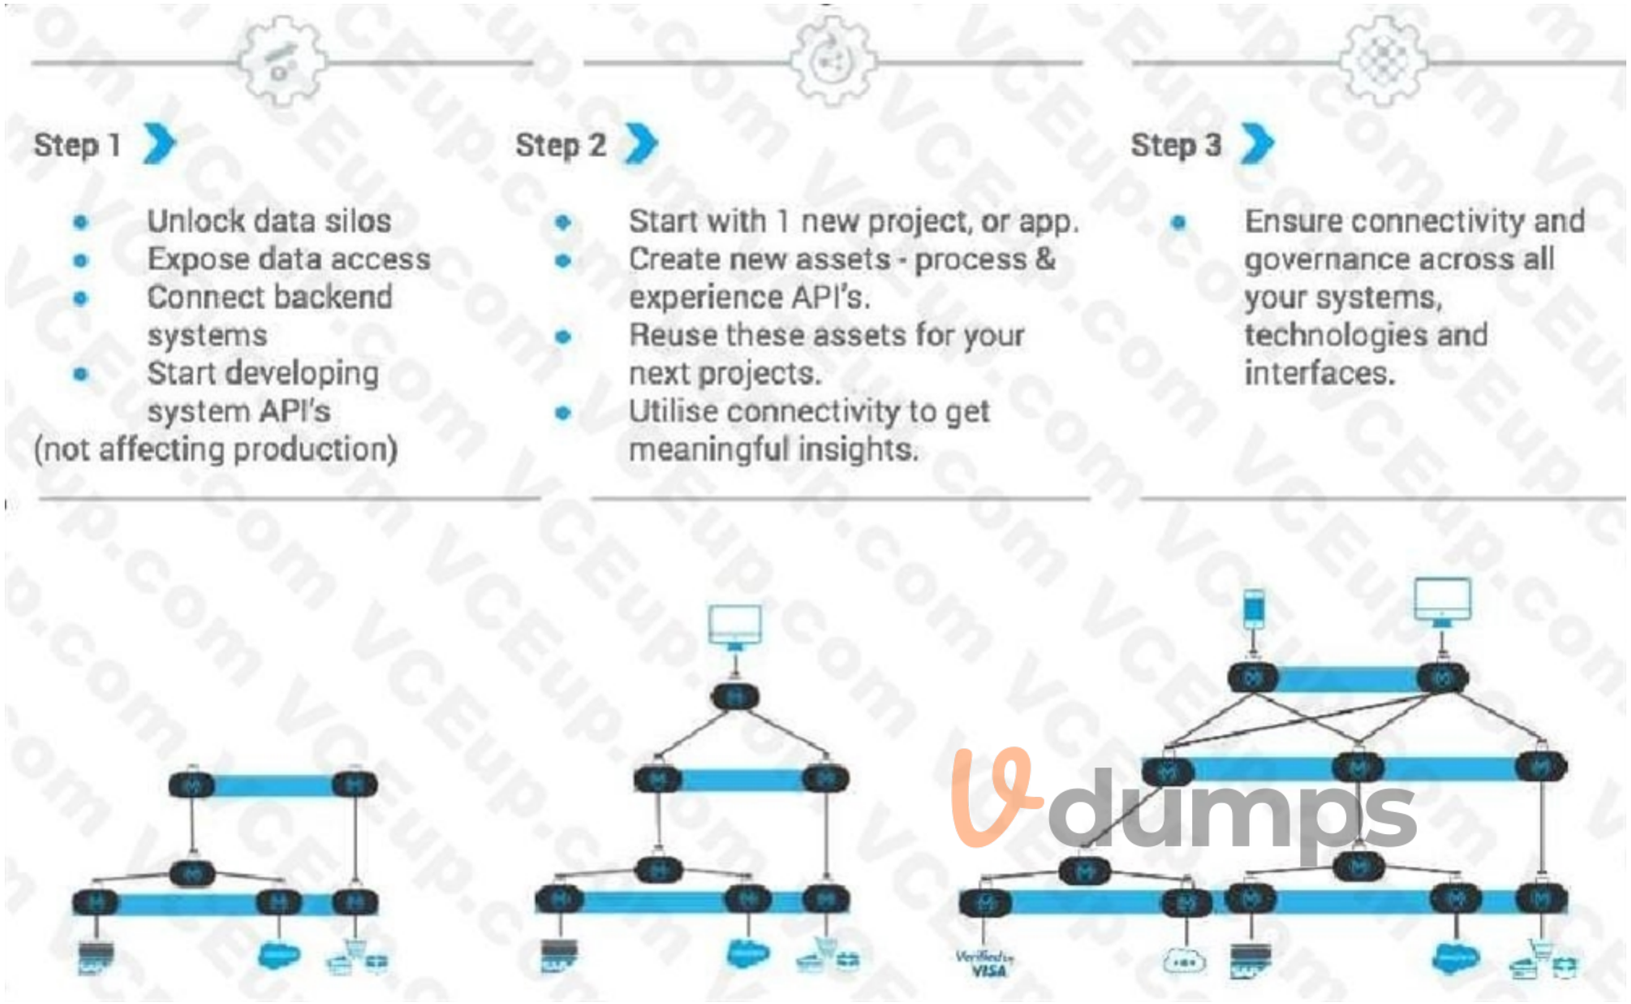

QUESTION 24

Refer to the exhibit.

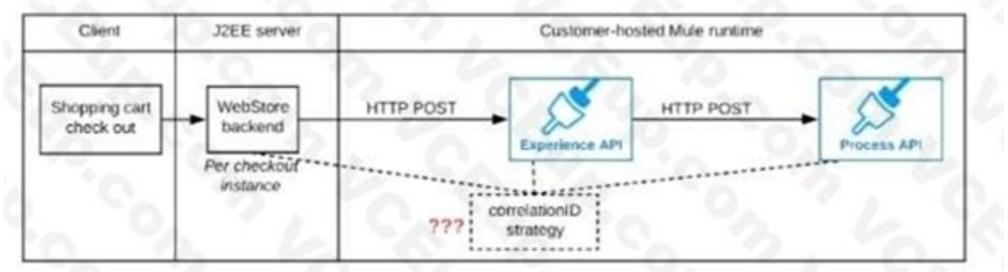

A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all APIimplementations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

End-to-end correlation of all HTTP requests and responses belonging to each individual checkoutInstance is required. This is to be done through a common correlation ID, so that all log entrieswritten by the web store backend, Experience API implementation, and Process API implementationinclude the same correlation ID for all requests and responses belonging to the same checkoutinstance.

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

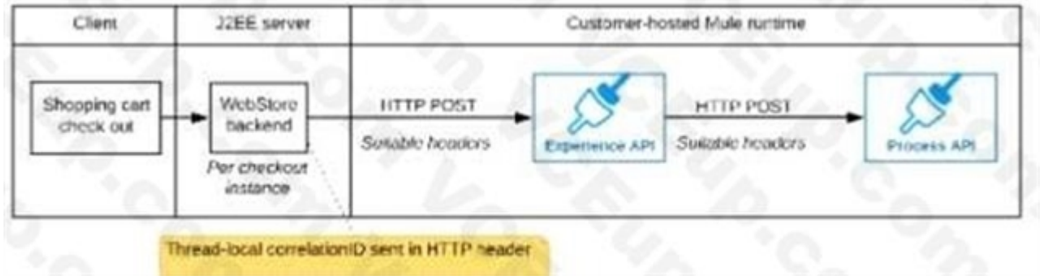

B. The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-It HTTP request header In each API invocation belonging to that checkout No special code or configuration is included in the Experience API and Process API implementations to generate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

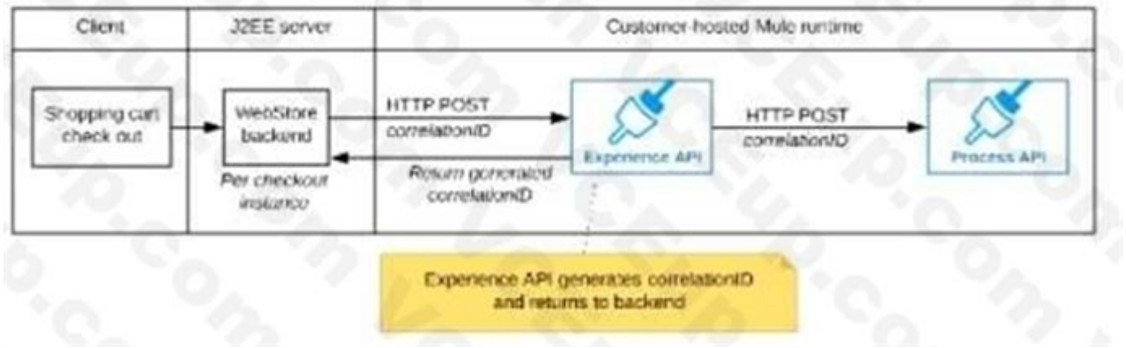

D. The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API The Experience API and Process API implementations must be coded to receive the custom correlation ID In the
HTTP requests and propagate It in suitable HTTP request headers

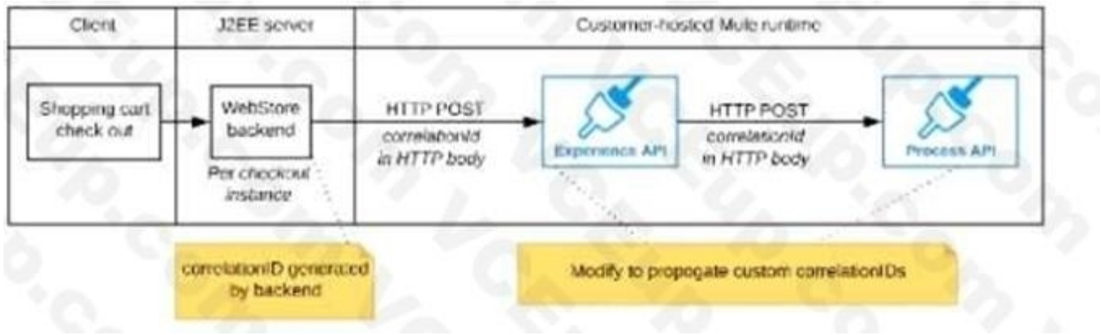

Correct Answer: B
Section:
Explanation:
Correct answer is "The web store backend generates a new correlation ID value at the start of checkout and sets it on the X¨CORRELATION-ID HTTP request header in each API invocation belonging to that checkout No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID" : By design, Correlation Ids cannot be changed within a flow in Mule 4 applications and can be set only at source.
This ID is part of the Event Context and is generated as soon as the message is received by the application. When a HTTP Request is received, the request is inspected for "X-Correlation-Id" header. If "X-Correlation-Id" header is present, HTTP connector uses this as the Correlation Id. If "X-Correlation-Id" header is NOT present, a Correlation Id is randomly generated. For Incoming HTTP Requests: In order to set a custom Correlation Id, the client invoking the HTTP request must set "X-Correlation-Id" header. This will ensure that the Mule Flow uses this Correlation Id. For Outgoing HTTP Requests: You can also propagate the existing Correlation Id to downstream APIs. By default, all outgoing HTTP Requests send "X-Correlation-Id" header. However, you can choose to set a different value to "X-Correlation- Id" header or set "Send Correlation Id" to NEVER.
Mulesoft Reference: https://help.mulesoft.com/s/article/How-to-Set-Custom-Correlation-Id-for-Flows-with-HTTP-Endpoint-in-Mule-4

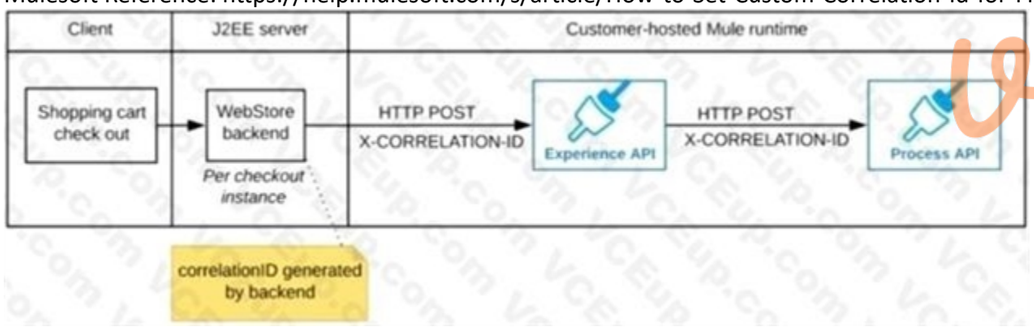

QUESTION 25
What operation can be performed through a JMX agent enabled in a Mule application?

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

Correct Answer: C
Section:
Explanation:
JMX Management Java Management Extensions (JMX) is a simple and standard way to manage applications, devices, services, and other resources. JMX is dynamic, so you can use it to monitor and manage resources as they are created, installed, and implemented. You can also use JMX to monitor and manage the Java Virtual Machine (JVM). Each resource is instrumented by one or more Managed Beans, or MBeans. All MBeans are registered in an MBean Server. The JMX server agent consists of an MBean Server and a set of services for handling Mbeans. There are several agents provided with Mule for JMX support. The easiest way to configure JMX is to use the default JMX support agent.
Log4J Agent The log4j agent exposes the configuration of the Log4J instance used by Mule for JMX management. You enable the Log4J agent using the <jmx-log4j> element. It does not take any additional properties MuleSoft

Reference:
https://docs.mulesoft.com/mule-runtime/3.9/jmxmanagement

QUESTION 26
Refer to the exhibit.

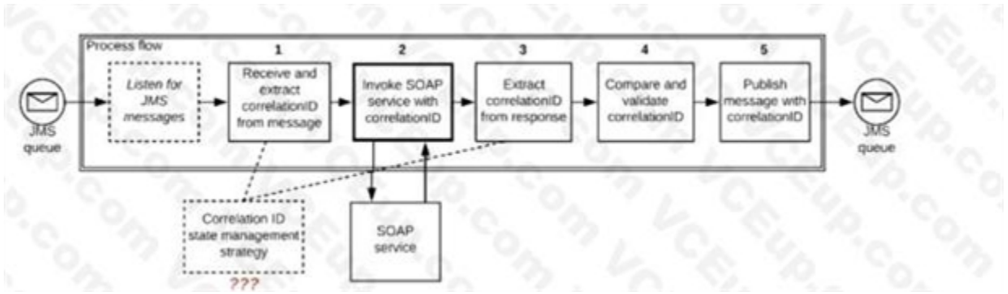

A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to receive JMS messages from a JMS queue.
To process each received JMS message, the following steps are performed in a flow:
Step l: The JMS Correlation ID header is read from the received JMS message.
Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMSCorrelation ID as one parameter in the SOAP request.
Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.
Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.
Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.
Where should the Mule application store the JMS Correlation ID values received in Step 1 and Step 3 so that the validation in Step 4 can be performed, while also making the overall Mule application highly available, fault-tolerant, performant, and maintainable?

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

Correct Answer: C
Section:
Explanation:
* If we store Correlation id value in step 1 as Mule event variables/attributes, the values will be cleared after server restart and we want system to be fault tolerant.
* The Correlation ID value in Step 1 should be stored in a persistent object store.
* We don't need to store Correlation ID value in Step 3 to persistent object store. We can store it but as we also need to make application performant. We can avoid this step of accessing persistent object store.
* Accessing persistent object stores slow down the performance as persistent object stores are by default stored in shared file systems.
* As the SOAP service is idempotent in nature. In case of any failures , using this Correlation ID saved in first step we can make call to SOAP service and validate the Correlation ID.
Additional Information:
* Competing Consumers are multiple consumers that are all created to receive messages from a single Point-to-Point Channel. When the channel delivers a message, any of the consumers could potentially receive it. The messaging system's implementation determines which consumer actually receives the message, but in effect the consumers compete with each other to be the receiver. Once a consumer receives a message, it can delegate to the rest of its application to help process the message.

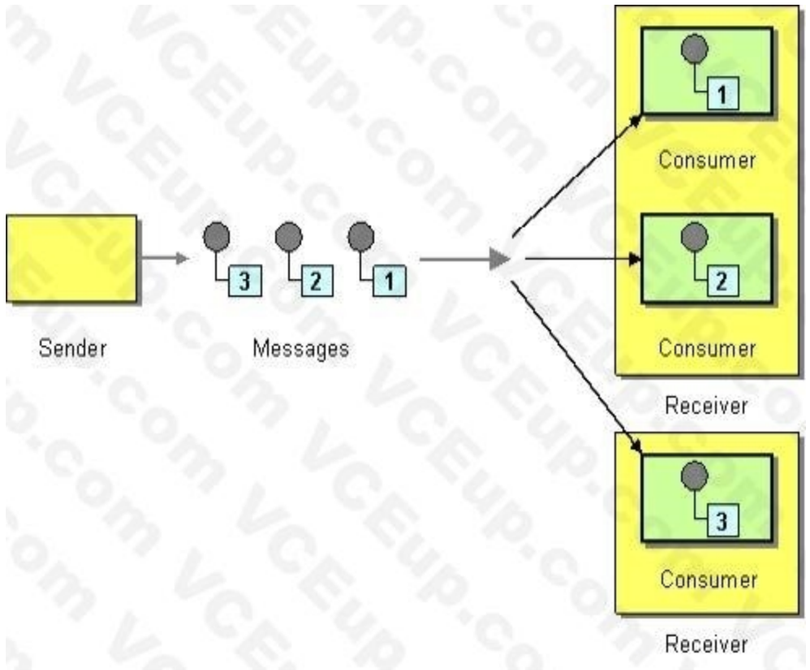

* In case you are unaware about term idempotent re is more info:
Idempotent operations means their result will always same no matter how many times these operations are invoked.

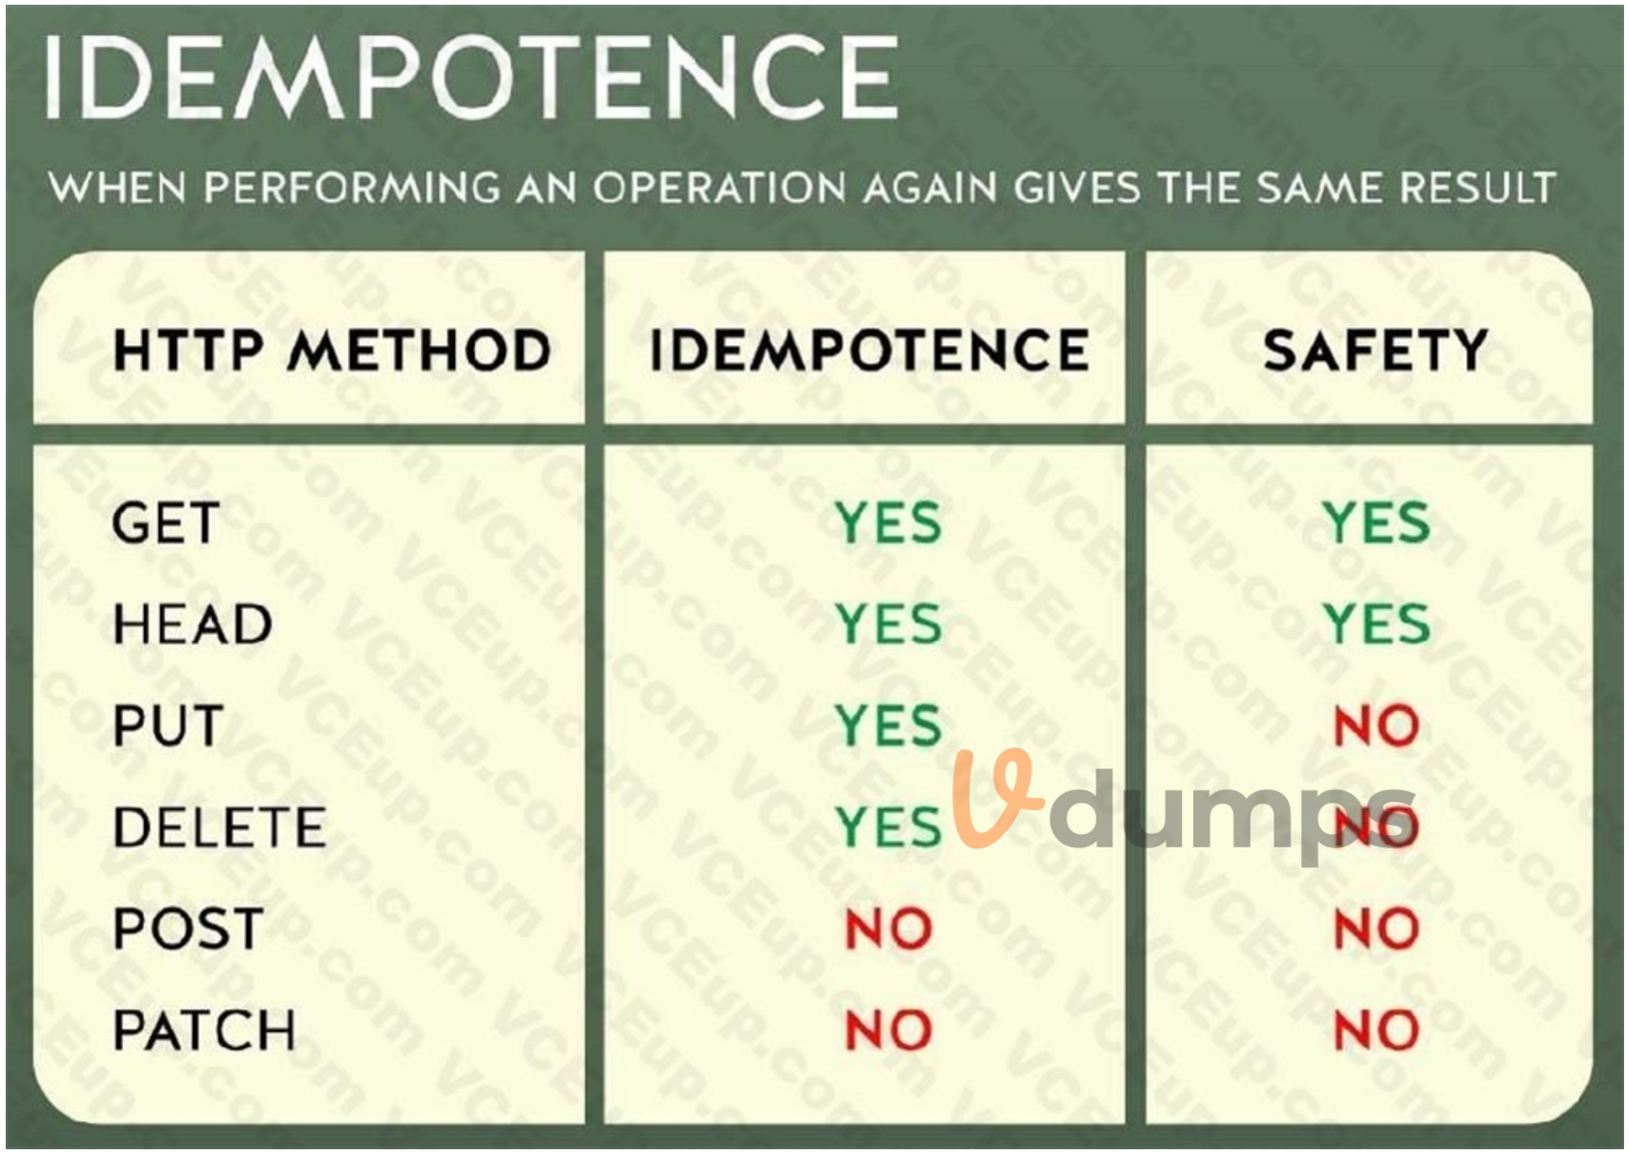

QUESTION 27
An integration Mute application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mute application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system. Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system).
The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed.
The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages.
What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

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

D. Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing

Correct Answer: D
Section:
Explanation:
Correct answer is using below set of activities Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing We will see why this is correct answer but before that lets understand few of the concepts which we need to know. Until Successful Scope The Until Successful scope processes messages through its processors until the entire operation succeeds. Until Successful repeatedly retries to process a message that is attempting to complete an activity such as: - Dispatching to outbound endpoints, for example, when calling a remote web service that may have availability issues. - Executing a component method, for example, when executing on a Spring bean that may depend on unreliable resources. - A sub-flow execution, to keep re-executing several actions until they all succeed, - Any other message processor execution, to allow more complex scenarios. How this will help requirement :
Using Until Successful Scope we can retry sending the order to backend systems in case of error to avoid manual processing later. Retry values can be configured in Until Successful Scope Apache ActiveMQ It is an open source message broker written in Java together with a full Java Message Service client ActiveMQ has the ability to deliver messages with delays thanks to its scheduler. This functionality is the base for the broker redelivery plug-in. The redelivery plug-in can intercept dead letter processing and reschedule the failing messages for redelivery. Rather than being delivered to a DLQ, a failing message is scheduled to go to the tail of the original queue and redelivered to a message consumer.
How this will help requirement : If backend application is down for a longer duration where Until Successful Scope wont work, then we can make use of ActiveMQ long retry Queue. The redelivery plug-in can intercept dead letter processing and reschedule the failing messages for redelivery. Mule Reference:
https://docs.mulesoft.com/mule-runtime/4.3/migration-core-until-successful

QUESTION 28
What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

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

Correct Answer: A
Section:
Explanation:
* Shared load balancers don't allow you to configure custom SSL certificates or proxy rules * Dedicated Load Balancer are optional but you need to purchase them additionally if needed.
* TLS is a cryptographic protocol that provides communications security for your Mule app. TLS offers many different ways of exchanging keys for authentication, encrypting data, and guaranteeing message integrity.
* The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
* Only a DLB allows the configuration of a custom TLS server certificate * DLB enables you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
* To use a DLB in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.
* MuleSoft Reference: https://docs.mulesoft.com/runtime-manager/dedicated-load-balancertutorialAdditional Info on SLB Vs DLB:

	Shared Load Balancer	Dedicated Load Balancer
VPC	Shared VPC (Mulesoft)	VPC (Customer)
Default Load Balancer	Cloudhub provides Deault Shared Load Balancer available in All Environment	Need to Purchase
Organization Use	Multiple Oragnization	Specific to Organization
Certificate	Mulesoft Certificate	Organization Certificate
TLS Support	Yes	Yes
URL Mapping	Fixed URL Mapping	Customer URL Mapping
Timeout	30 Sec Session Timeout	Custom Timeout
Ports	Public Port {80 : 8081, 443 : 8082}	Private Port {80 : 8091, 443 : 8092}
Fashion	Round Robin	Round Robin
Supports HTTPS Protocol	Yes	Yes
Worker Assignment	No	Yes
IP Blacklisting/ Whitelisting	No https://docs.mulesoft.com/runtime-manager/lb-whitelists	Yes
Configure Custom Domain	No	Yes
Custom Certificate	No	Yes
Rate Limit	Lower Rate Limit and applied According to Region	Higher Rate Limit Threshold
VPC	Anypoint VPC optional	Can't Use DLB without Anypoint VPC

QUESTION 29
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

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

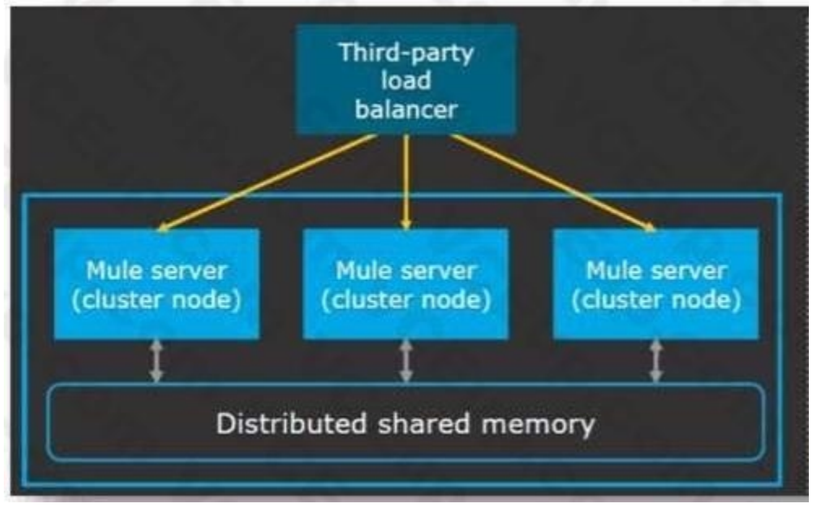

QUESTION 30
Refer to the exhibit.
An organization is designing a Mule application to receive data from one external business partner.
The two companies currently have no shared IT infrastructure and do not want to establish one.
Instead, all communication should be over the public internet (with no VPN).
What Anypoint Connector can be used in the organization's Mule application to securely receive data from this external business partner?

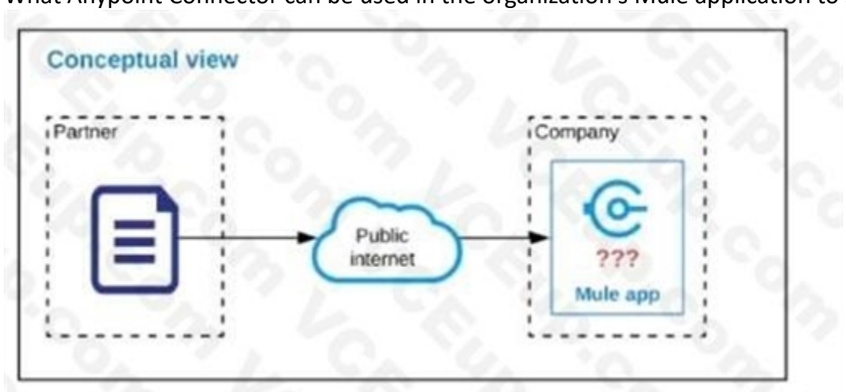

A. File connector

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

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

QUESTION 31
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

Correct Answer: D
Section:
Explanation:
Security Assertion Markup Language (SAML) is an open standard that allows identity providers (IdP) to pass authorization credentials to service providers (SP).
SAML transactions use Extensible Markup Language (XML) for standardized communications between the identity provider and service providers.
SAML is the link between the authentication of a user's identity and the authorization to use a service.
WS-Security is the key extension that supports many authentication models including: basic username/password credentials, SAML, OAuth and more.
A common way that SOAP API's are authenticated is via SAML Single Sign On (SSO). SAML works by facilitating the exchange of authentication and authorization credentials across applications.
However, there is no specification that describes how to add SAML to REST web services.
Reference : https://www.oasis-open.org/committees/download.php/16768/wss-v1.1-spec-os-SAMLTokenProfile.pdf

QUESTION 32
Refer to the exhibit.

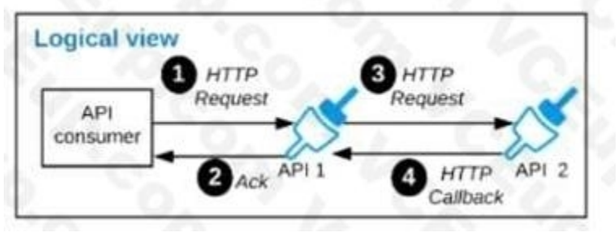

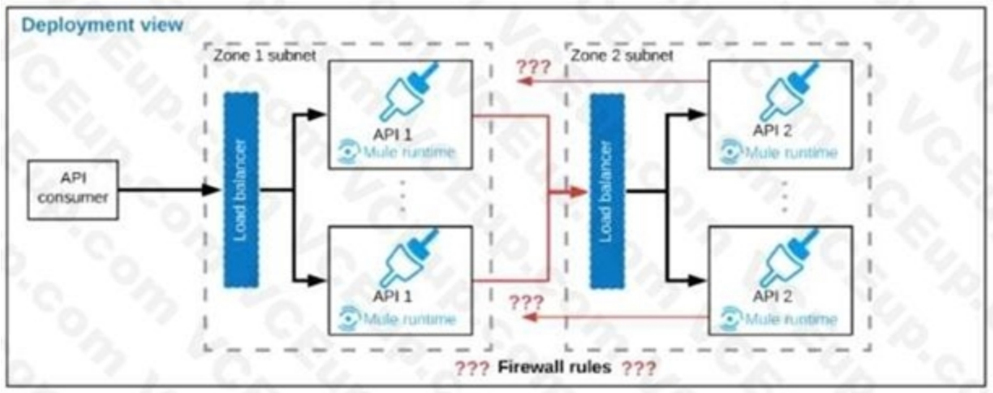

A business process involves two APIs that interact with each other asynchronously over HTTP. EachAPI is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API l to 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

Correct Answer: B
Section:
Explanation:
* If your API implementation involves putting a load balancer in front of your APIkit application, configure the load balancer to redirect URLs that reference the baseUri of the application directly. If the load balancer does not redirect URLs, any calls that reach the load balancer looking for the application do not reach their destination.
* When you receive incoming traffic through the load balancer, the responses will go out the same way. However, traffic that is originating from your instance will not pass through the load balancer.
Instead, it is sent directly from the public IP address of your instance out to the Internet. The ELB is not involved in that scenario.
* The question says "each API is deployed to multiple redundant Mule runtimes", that seems to be a hint for self hosted Mule runtime cluster. Set Inbound allowed for the LB, outbound allowed for runtime to request out.
* Hence correct way is to enable communication from each API's Mule Runtimes and Network zone to the load balancer of the other API. Because communication is asynchronous one
Reference: https://docs.mulesoft.com/apikit/4.x/configure-load-balancer-task

QUESTION 33
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

Correct Answer: C
Section:
Explanation:
Correct answer is Application A accesses the persistent object store via the Object Store REST API Application B uses the Object Store connector to access a persistent object * Object Store v2 lets CloudHub applications store

data and states across batch processes, Mule components and applications, from within an application or by using the Object Store REST API. * On-premise Mule applications cannot use Object Store v2. * You can select Object Store v2 as the implementation for Mule 3 and Mule 4 in CloudHub by checking the Object Store V2 checkbox in Runtime Manager at deployment time. * CloudHub Mule applications can use Object Store connector to write to the object store * The only way on-premises Mule applications can access Object Store v2 is via the Object Store REST API.
* You can configure a Mule app to use the Object Store REST API to store and retrieve values from an object store in another Mule app.

QUESTION 34
What limits if a particular Anypoint Platform user can discover an asset in Anypoint Exchange?

A. Design Center and RAML were both used to create the asset
B. The existence of a public Anypoint Exchange portal to which the asset has been published
C. The type of the asset in Anypoint Exchange
D. The business groups to which the user belongs

Correct Answer: D
Section:
Explanation:
* "The existence of a public Anypoint Exchange portal to which the asset has been published" -question does not mention anything about the public portal. Beside the public portal is open to the internet, to anyone.
* If you cannot find an asset in the current business group scopes, search in other scopes. In the left navigation bar click All assets (assets provided by MuleSoft and your own master organization), Provided by MuleSoft, or a business group scope. User belonging to one Business Group can see assets related to his group only Reference:
https://docs.mulesoft.com/exchange/to-find-info https://docs.mulesoft.com/exchange/asset-detailsCorrect answer is The business groups to which the user belongs

QUESTION 35
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

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

QUESTION 36
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

Correct Answer: D
Section:
Explanation:
Correct answer is 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 : Using Anypoint MQ, you can create two types of queues: Standard queue These queues don't guarantee a specific message order. Standard queues are the best fit for applications in which messages must be delivered quickly. FIFO (first in, first out) queue These queues ensure that your messages arrive in order. FIFO queues are the best fit for applications requiring strict message ordering and exactly-once delivery, but in which message delivery speed is of less importance Use of FIFO queue is no where in the option and also it decreased throughput. Similarly persistent object store is not the preferred solution approach when you maximizing message throughput. This rules out one of the options. Scatter Gather does not support ObjectStore. This rules out one of the options. Standard Anypoint MQ queues don't guarantee a specific message order hence using another for each block to collect response wont work as requirement here is to ensure the order. Hence considering all the above factors the feasible approach is 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

QUESTION 37
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?

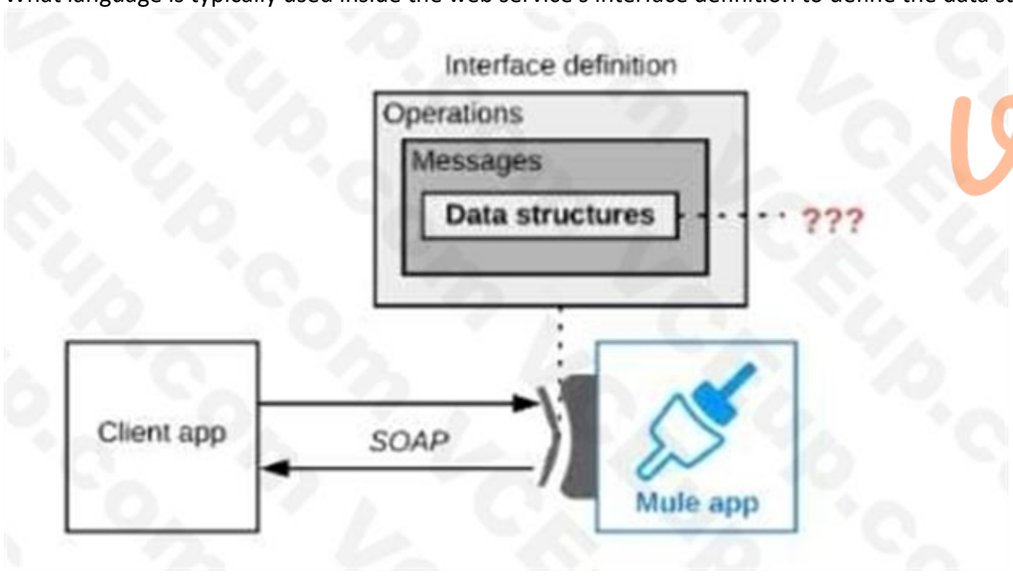

A. WSDL
B. XSD
C. JSON Schema
D. RAMI

Correct Answer: X, S, D, I, N, T, H, I, S, A, P, P, R, O, A, C, H, T, O, D, E, V, E, L, O, P, I, N, G, A, W, E, B, S, E, R, V, I, C, E, Y, O, U, B, E, G, I, N, W, I, T, H
Section:
Explanation:
Answer: XSD In this approach to developing a web service, you begin with
Explanation:an XML schema (XSD file) that defines XML data structures to be used as parameters and return types in the web service operations.

---------- Reference:
https://www.w3schools.com/xml/schema_intro.asp

QUESTION 38
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

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

QUESTION 39
Refer to the exhibit.

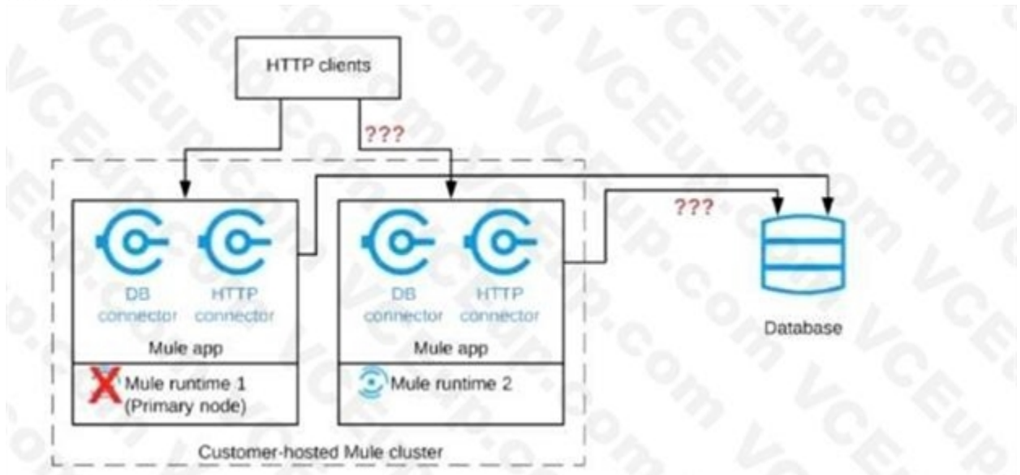

A Mule application is deployed to a cluster of two customer-hosted Mute runtimes. The Mute application has a flow that polls a database and another flow with an HTTP Listener.
HTTP clients send HTTP requests directly to individual cluster nodes.
What happens to database polling and HTTP request handling 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 beaccepted

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 failednode Incur increased latency
D. Database polling stops All HTTP requests are rejected

Correct Answer: A
Section:
Explanation:
Correct answer is Database polling continues Only HTTP requests sent to the remaining nodecontinue to be accepted. : Architecture descripted in the question could be described asfollows.When node 1 is down , DB polling will still continue via node 2 . Also requests which arecoming directly to node 2 will also be accepted and processed in BAU fashion. Only thing that wontwork is when requests are sent to Node 1 HTTP connector. The flaw with this architecture is HTTPclients are sending HTTP requests directly to individual cluster nodes. By default, clustering Muleruntime engines ensures high system availability. If a Mule runtime engine node becomesunavailable due to failure or planned downtime, another node in the cluster can assume theworkload and continue to process existing events and messages

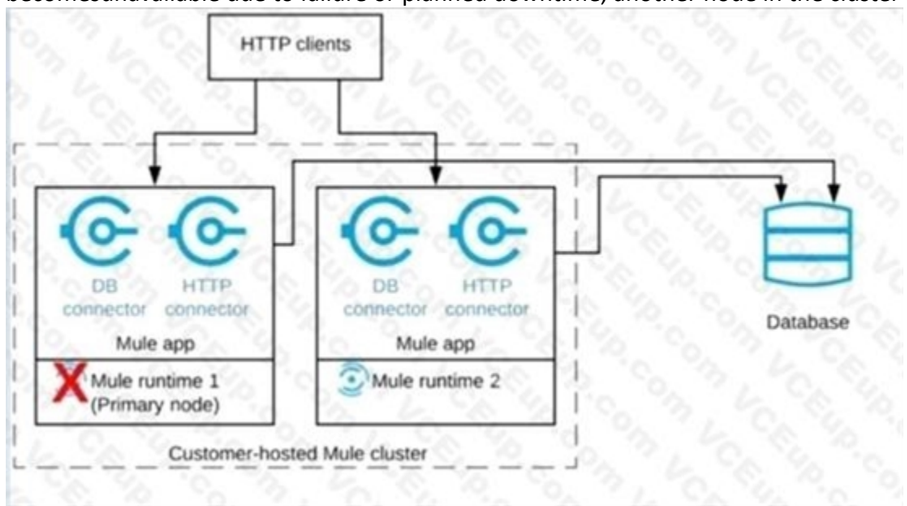

QUESTION 40
A global organization operates datacenters in many countries. There are private network links between these datacenters because all business data (but NOT metadata) must be exchanged over these private network connections.
The organization does not currently use AWS in any way.
The strategic 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

Correct Answer: D
Section:
Explanation:
Correct answer is MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in each datacenter There are two things to note about the question which can help us figure out correct answer..
* Business data must be exchanged over these private network connections which means we can not use MuleSoft provided Cloudhub option. So we are left with either customer hosted runtime in external cloud provider or customer hosted runtime in their own premises. As customer does not use AWS at the moment. Hence that don't have the immediate option of using Customer-hosted runtime plane in multiple AWS regions. hence the most suitable option for runtime plane is Customer-hosted runtime plane in each datacenter
* Metadata has no limitation to reside in organization premises. Hence for control plane MuleSoft hosted Anypoint platform can be used as a strategic solution.
Hybrid is the best choice to start. Mule hosted Control plane and Customer hosted Runtime to start with.

Once they mature in cloud migration, everything can be in Mule hosted.

QUESTION 41
Refer to the exhibit.

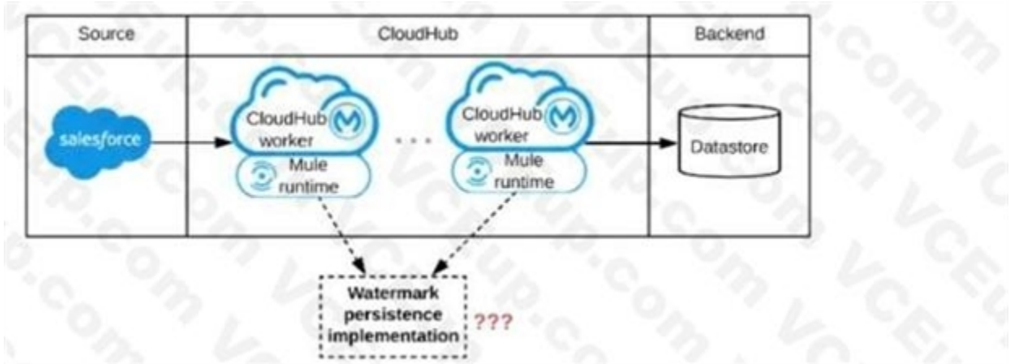

A Mule application is being designed to be deployed to several CloudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.
A watermark will be used to only retrieve those Satesforce Accounts that have been modified since 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?

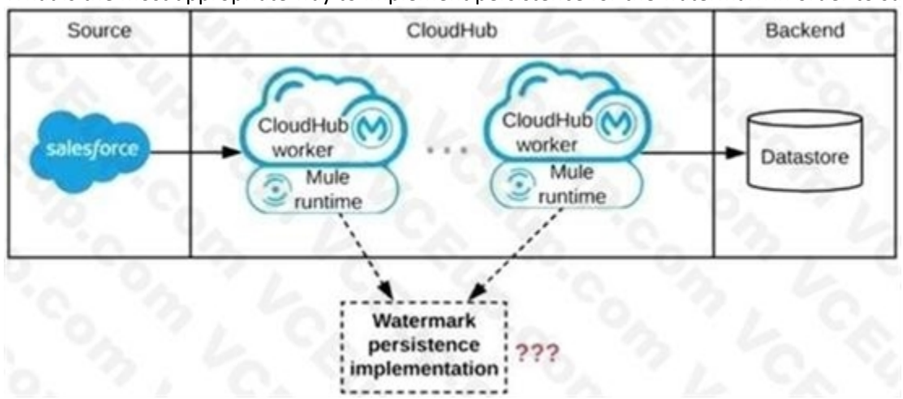

A. Persistent Anypoint MQ Queue
B. Persistent Object Store
C. Persistent Cache Scope
D. Persistent VM Queue

Correct Answer: B
Section:
Explanation:
* An object store is a facility for storing objects in or across Mule applications. Mule uses object stores to persist data for eventual retrieval.
* Mule provides two types of object stores:
1) In-memory store ñ stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime.
2) Persistent store ñ Mule persists data when an object store is explicitly configured to be persistent.
In a standalone Mule runtime, Mule creates a default persistent store in the file system. If you do not specify an object store, the default persistent object store is used.
MuleSoft Reference: https://docs.mulesoft.com/mule-runtime/3.9/mule-object-stores

QUESTION 42
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

Correct Answer: D
Section:
Explanation:
Correct answer is Implement transformation logic in the new Mule application using DataWeave, invoking existing transformation services when possible. * The key here minimal testing effort, "Extend existing transformation logic" is not a feasible option because additional functionality is highly specific to the new Mule application so it should not be a part of commonly used functionality.
So this option is ruled out.
* "Implement transformation logic in the new Mule application using DataWeave, replicating the transformation logic of existing transformation services" Replicating the transformation logic of existing transformation services will cause duplicity of code. So this option is ruled out.
* "Implement and expose all transformation logic as microservices using DataWeave, so it can be reused by any application component that needs it, including the new Mule application" as question specifies that the transformation is app specific and wont be used outside

QUESTION 43
An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub.
For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub's default zero-downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers.This is expected to result in several Mule application deployments per hour.

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

Correct Answer: A
Section:
Explanation:
* When you create an Anypoint VPC, the range of IP addresses for the network must be specified in the form of a Classless Inter-Domain Routing (CIDR) block, using CIDR notation.
* This address space is reserved for Mule workers, so it cannot overlap with any address space used in your data center if you want to peer it with your VPC.
* To calculate the proper sizing for your Anypoint VPC, you first need to understand that the number of dedicated IP addresses is not the same as the number of workers you have deployed.
* For each worker deployed to CloudHub, the following IP assignation takes place: For better fault tolerance, the VPC subnet may be divided into up to four Availability Zones.
* A few IP addresses are reserved for infrastructure. At least two IP addresses per worker to perform at zero-downtime.
* Hence in this scenario 2048 IP's are required to support the requirement.

QUESTION 44
Refer to the exhibit.

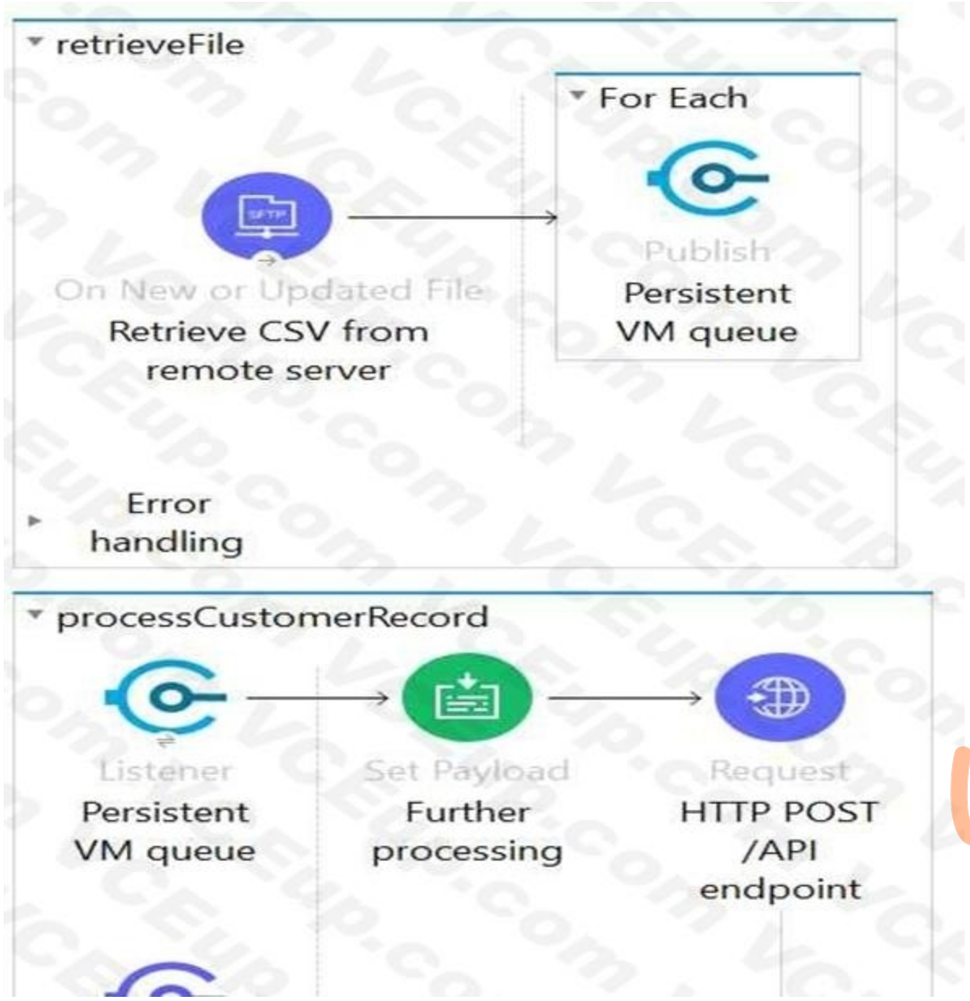

This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow's VM Listner receives messages from 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.

Correct Answer: B
Section:

QUESTION 45
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

Correct Answer: B
Section:
Explanation:
Scatter-Gather sends a request message to multiple targets concurrently. It collects the responses from all routes, and aggregates them into a single message.
Reference: https://docs.mulesoft.com/mule-runtime/4.3/scatter-gather-concept

QUESTION 46
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 RuntimeManager
Define 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 Manager
Define the splunk appender in EACH Mule application's log4j2.xml file
C. Disable Cloudhub logging in Runtime Manager
Define 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

Correct Answer: B
Section:
Explanation:
By default, CloudHub replaces a Mule application's log4j2.xml file with a CloudHub log4j2.xml file. In CloudHub, you can disable the CloudHub provided Mule application log4j2 file. This allows integrating Mule application logs with custom or third-party log management systems

QUESTION 47
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

Correct Answer: D
Section:
Explanation:
* XA transaction add tremendous latency so "Reduce Latency" is incorrect option XA transactions define "All or No" commit protocol.
* Each local XA resource manager supports the A.C.I.D properties (Atomicity, Consistency, Isolation, and Durability).
---------- So correct choice is "Ensure consistency"
Reference: https://docs.mulesoft.com/mule-runtime/4.3/xa-transactions

QUESTION 48
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 worker New 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 worker New 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 worker New 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 worker New web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.

Correct Answer: C
Section:
Explanation:
CloudHub supports updating your applications at runtime so end users of your HTTP APIs experiencezero downtime. While your application update is deploying, CloudHub keeps the old version of yourapplication running. Your domain points to the old version of your application until the newlyuploaded version is fully started. This allows you to keep servicing requests from your old applicationwhile the new version of your application is starting.

QUESTION 49
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 flow
On Error Continue error handler in the child flow
B. For Each scope in the parent flow
On Error Continue error handler in the child flow
C. Parallel For Each scope in the parent flow
On Error Propagate error handler in the child flow
D. Until Successful router in the parent flow
On Error Propagate error handler in the child flow

Correct Answer: B
Section:
Explanation:
Correct answer is For Each scope in the parent flow On Error Continue error handler in the child flow.
You can extract below set of requirements from the question a) Records should be sent to downstream system in the same order that it was received in the array b) Any validation or communication failures should not prevent further processing of the remaining records First requirement can be met using For Each scope in the parent flow and second requirement can be met using On Error Continue scope in child flow so that error will be suppressed.

QUESTION 50
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.comAs part of the migration strategy, the organization plans to migrate all of its Mule applications andload balancer to cloudhubWhat is the most straight-forward and cost effective approach to the Mule applications deploymentand load balancing that preserves the public URLs?

A. Deploy the Mule applications to Cloudhub Update 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 Cloudhub Create CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancer Apply mapping rules in the SLB to map URLs to their corresponding Mule applications

D. Deploy the Mule applications to Cloudhub Update 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.

Correct Answer: A
Section:
Explanation:
https://help.mulesoft.com/s/feed/0D52T000055pzgsSAA.

QUESTION 51
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

Correct Answer: A
Section:
Explanation:
Correct answer is VM VM support Transactional Type. When an exception occur, The transaction rolls back to its original state for reprocessing. This feature is not supported by other connectors.
Here is additional information about Transaction management:

	Shared Load Balancer	Dedicated Load Balancer
VPC	Shared VPC (Mulesoft)	VPC (Customer)
Default Load Balancer	Cloudhub provides Deault Shared Load Balancer available in All Environment	Need to Purchase
Organization Use	Multiple Oragnization	Specific to Organization
Certificate	Mulesoft Certificate	Organization Certificate
TLS Support	Yes	Yes
URL Mapping	Fixed URL Mapping	Customer URL Mapping
Timeout	30 Sec Session Timeout	Custom Timeout
Ports	Public Port {80 : 8081, 443 : 8082}	Private Port {80 : 8091, 443 : 8092}
Fashion	Round Robin	Round Robin
Supports HTTPS Protocol	Yes	Yes
Worker Assignment	No	Yes
IP Blacklisting/ Whitelisting	No	Yes
	https://docs.mulesoft.com/runtime-manager/lb-whitelists	
Configure Custom Domain	No	Yes
Custom Certificate	No	Yes
Rate Limit	Lower Rate Limit and applied According to Region	Higher Rate Limit Threshold
VPC	Anypoint VPC optional	Can't Use DLB without Anypoint VPC

QUESTION 52
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 theexternal 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

Correct Answer: B
Section:
Explanation:
3xx HTTP status codes indicate a redirection that the user agent (a web browser or a crawler) needsto take further action when trying to access a particular resource.
Reference: https://www.w3.org/Protocols/rfc2616/rfc2616-sec10.html

QUESTION 53
An organization is migrating all its Mule applications to Runtime Fabric (RTF). None of the Mule applications use Mule domain projects.
Currently, all the Mule applications have been manually deployed to a server group among several customer hosted Mule runtimes.
Port conflicts between these Mule application deployments are currently managed by the DevOps team who carefully manage Mule application properties files.
When the Mule applications are migrated from the current customer-hosted server group to Runtime Fabric (RTF), fo the Mule applications need to be rewritten and what DevOps port configuration responsibilities change or stay the same?

A. Yes, the Mule applications Must be rewritten
DevOps No Longer needs to manage port conflicts between the Mule applications
B. Yes, the Mule applications Must be rewritten
DevOps Must Still Manage port conflicts.
C. NO, The Mule applications do NOT need to be rewritten
DevOps MUST STILL manage port conflicts
D. NO, the Mule applications do NO need to be rewritten
DevOps NO LONGER needs to manage port conflicts between the Mule applications.

Correct Answer: C
Section:
Explanation:
* Anypoint Runtime Fabric is a container service that automates the deployment and orchestration of your Mule applications and gateways.
* Runtime Fabric runs on customer-managed infrastructure on AWS, Azure, virtual machines (VMs) or bare-metal servers.
* As none of the Mule applications use Mule domain projects. applications are not required to be rewritten. Also when applications are deployed on RTF, by default ingress is allowed only on 8081.
* Hence port conflicts are not required to be managed by DevOps team

QUESTION 54
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, customerprovided, 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)

Correct Answer: A
Section:
Explanation:
https://docs.mulesoft.com/runtime-manager/dedicated-load-balancer-tutorial

QUESTION 55
A Mule application is being designed To receive nightly a CSV file containing millions of records from an external vendor over SFTP, The records from the file need to be validated, transformed. And then written to a database. Records can be inserted into the database in any order.
In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

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

Correct Answer: C
Section:
Explanation:
Correct answer is Use a Batch job scope to bulk insert records into the database * Batch Job is most efficient way to manage millions of records.
A few points to note here are as follows :
Reliability: If you want reliabilty while processing the records, i.e should the processing survive a runtime crash or other unhappy scenarios, and when restarted process all the remaining records, if yes then go for batch as it uses persistent queues.
Error Handling: In Parallel for each an error in a particular route will stop processing the remaining records in that route and in such case you'd need to handle it using on error continue, batch process does not stop during such error instead you can have a step for failures and have a dedicated handling in it.
Memory footprint: Since question said that there are millions of records to process, parallel for each will aggregate all the processed records at the end and can possibly cause Out Of Memory.
Batch job instead provides a BatchResult in the on complete phase where you can get the count of failures and success. For huge file processing if order is not a concern definitely go ahead with Batch Job

QUESTION 56
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

Correct Answer: D
Section:
Explanation:
Anypoint Platform provides a scripting and command-line tool for both Anypoint Platform and Anypoint Platform Private Cloud Edition (Anypoint Platform PCE). The command-line interface (CLI) supports both the interactive shell and standard CLI modes and works with: Anypoint Exchange Access management Anypoint Runtime Manager

QUESTION 57
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

Correct Answer: D
Section:
Explanation:
As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO).
Configure identity management using one of the following single sign-on standards:
1) OpenID Connect: End user identity verification by an authorization server including SSO 2) SAML 2.0: Web-based authorization including cross-domain SSO

QUESTION 58
An API implementation is being developed to expose data from a production database via HTTPrequests. The API implementation executes a database SELECT statement that is dynamically createdbased upon data received from each incoming HTTP request. The developers are planning to usevarious types of testing to make sure the Mule application works as expected, can handle specificworkloads, and behaves correctly from an API consumer perspective. What type of testing wouldtypically mock the results from each SELECT statement rather than actually execute it in theproduction database?

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

Correct Answer: A
Section:
Explanation:
In Unit testing instead of using actual backends, stubs are used for the backend services. This ensures that developers are not blocked and have no dependency on other systems.
In Unit testing instead of using actual backends, stubs are used for the backend services. This ensures that developers are not blocked and have no dependency on other systems.
Below are the typical characteristics of unit testing.
-- Unit tests do not require deployment into any special environment, such as a staging environment
-- Unit tests san be run from within an embedded Mule runtime
-- Unit tests can/should be implemented using MUnit
-- For read-only interactions to any dependencies (such as other APIs): allowed to invoke production endpoints
-- For write interactions: developers must implement mocks using MUnit
-- Require knowledge of the implementation details of the API implementation under test

QUESTION 59
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

Correct Answer: A
Section:
Explanation:

SOAP API specifications are provided as WSDL. Design center doesn't provide the functionality to create WSDL file. Hence WSDL needs to be created using XML editor

QUESTION 60
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

Correct Answer: C
Section:
Explanation:
* Mule Domain Project is ONLY available for customer-hosted Mule runtimes, but not for Anypoint Runtime Fabric
* Mule domain project is available for Hybrid and Private Cloud (PCE). Rest all provide application isolation and can't support domain project.
What is Mule Domain Project?
* A Mule Domain Project is implemented to configure the resources that are shared among different projects. These resources can be used by all the projects associated with this domain. Mule applications can be associated with only one domain, but a domain can be associated with multiple projects. Shared resources allow multiple development teams to work in parallel using the same set of reusable connectors. Defining these connectors as shared resources at the domain level allows the team to: - Expose multiple services within the domain through the same port. - Share the connection to persistent storage. - Share services between apps through a well-defined interface. - Ensure consistency between apps upon any changes because the configuration is only set in one place.
* Use domains Project to share the same host and port among multiple projects. You can declare the http connector within a domain project and associate the domain project with other projects. Doing this also allows to control thread settings, keystore configurations, time outs for all the requests made within multiple applications. You may think that one can also achieve this by duplicating the http connector configuration across all the applications. But, doing this may pose a nightmare if you have to make a change and redeploy all the applications.
* If you use connector configuration in the domain and let all the applications use the new domain instead of a default domain, you will maintain only one copy of the http connector configuration. Any changes will require only the domain to the redeployed instead of all the applications.
You can start using domains in only three steps:
1) Create a Mule Domain project
2) Create the global connector configurations which needs to be shared across the applications inside the Mule Domain project
3) Modify the value of domain in mule-deploy.properties file of the applications

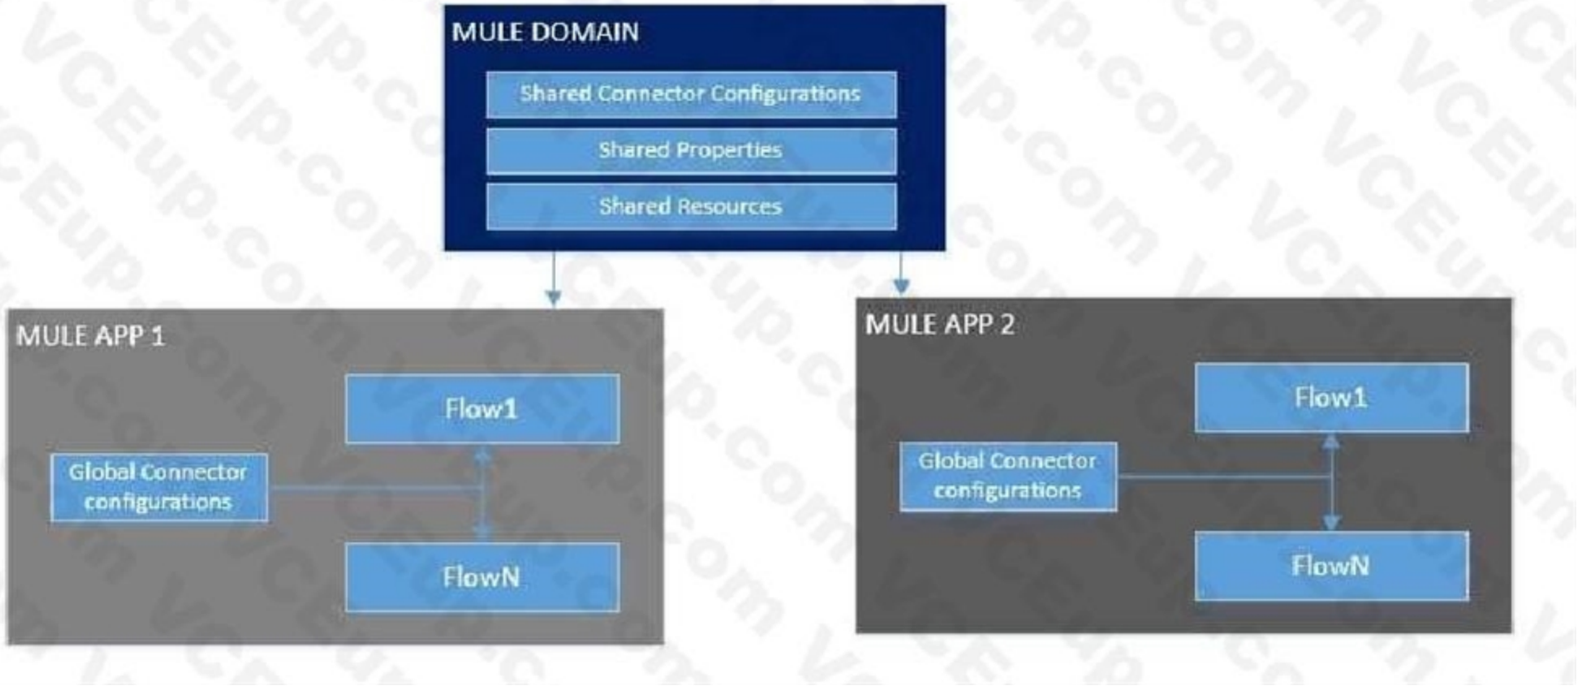

QUESTION 61
An API implementation is being designed that must invoke an Order API which is known to repeatedly experience downtime. For this reason a fallback API is to be called when the Order API is unavailable. What approach to designing invocation of the fallback API provides the best resilience?

A. Redirect client requests through an HTTP 303 temporary redirect status code to the fallback APIwhenever the Order API is unavailable
B. Set an option in the HTTP Requester component that invokes the order API to instead invoke afallback API whenever an HTTP 4XX or 5XX response status code is received from Order API
C. Create a separate entry for the order API in API manager and then invoke this API as a fallback API if the primary Order API is unavailable
D. Search Anypoint Exchange for a suitable existing fallback API and them implement invocations to their fallback API in addition to the Order API

Correct Answer: A
Section:
Explanation:
* Resilience testing is a type of software testing that observes how applications act under stress. It's meant to ensure the product's ability to perform in chaotic conditions without a loss of core functions or data; it ensures a quick recovery after unforeseen, uncontrollable events.
* In case an API invocation fails ó even after a certain number of retries ó it might be adequate to invoke a different API as a fallback. A fallback API, by definition, will never be ideal for the purpose of the API client, otherwise it would be the primary API.
* Here are some examples for fallback APIs:
- An old, deprecated version of the same API.
- An alternative endpoint of the same API and version (e.g. API in another CloudHub region).
- An API doing more than required, and therefore not as performant as the primary API.
- An API doing less than required and therefore forcing the API Client to offer a degraded service, which is still better than no service at all.
* API clients implemented as Mule applications offer the 'Until Successful Scope and Exception' strategies at their disposal, which together allow configuring fallback actions such as a fallback API invocation.
* All HTTP response status codes within the 3xx category are considered redirection messages. These codes indicate to the user agent (i.e. your web browser) that an additional action is required in orderto complete the request and access the desired resource

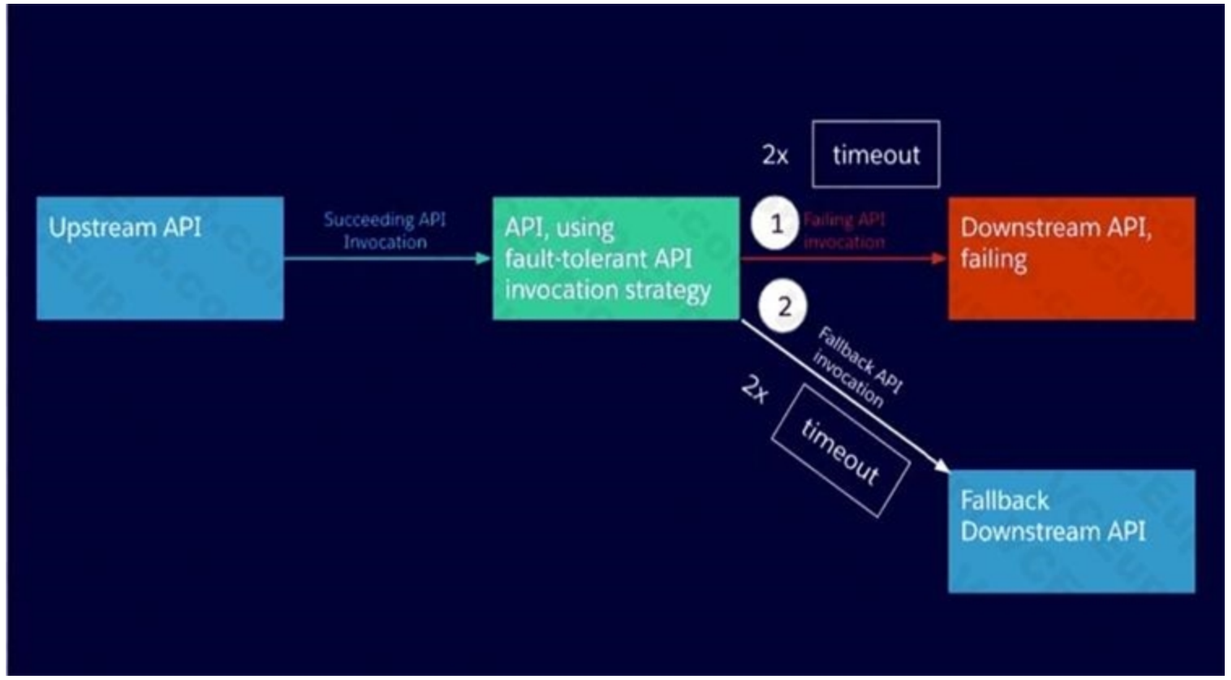

Hence correct answer is Redirect client requests through an HTTP 303 temporary redirect status codeto the fallback API whenever the Order API is unavailable

QUESTION 62

How are the API implementation , API client, and API consumer combined to invoke and process an API ?

A. The API consumer creates an API implementation , which receives API invocations from an API such that they are processed for an API client
B. The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation
C. An API client creates an API consumer, which receives API invocation from an API such that they are processed for an API implementation
D. The API client creates an API consumer which sends API invocations to an API such that they are processed by API implementation

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

QUESTION 63

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

Correct Answer: A
Section:
Explanation:
* HTTP/1.1 keeps all requests and responses in plain text format.
* HTTP/2 uses the binary framing layer to encapsulate all messages in binary format, while stillmaintaining HTTP semantics, such as verbs, methods, and headers. It came into use in 2015, andoffers several methods to decrease latency, especially when dealing with mobile platforms andserver-intensive graphics and videos
* Currently, Mule application can have API policies only for Mule application that accepts requestsover HTTP/1x

QUESTION 64
The implementation of a Process API must change. What is a valid approach that minimizes the impact of this change on API clients?

A. Implement required changes to the Process API implementation so that whenever possible, the Process API's RAML definition remains unchanged
B. Update the RAML definition of the current Process API and notify API client developers by sending them links to the updated RAML definition
C. Postpone changes until API consumers acknowledge they are ready to migrate to a new Process API or API version
D. Implement the Process API changes in a new API implementation, and have the old API implementation return an HTTP status code 301 - Moved Permanently to inform API clients they should be calling the new API implementation

Correct Answer: A
Section:
Explanation:
* Option B shouldn't be used unless extremely needed, if RAML is changed, client needs to accommodate changes. Question is about minimizing impact on Client. So this is not a valid choice.
* Option C isn't valid as Business can't stop for consumers acknowledgment.
* Option D again needs Client to accommodate changes and isn't viable option.
* Best choice is A where RAML definition isn't changed and underlined functionality is changed without any dependency on client and without impacting client.

QUESTION 65
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

Correct Answer: B
Section:
Explanation:
High availability is about up-time of your application
A) High availability can be achieved only in CloudHub isn't correct statement. It can be achieved in customer hosted runtime planes as well
B) An object store is a facility for storing objects in or across Mule applications. Mule runtime engine
(Mule) uses object stores to persist data for eventual retrieval. It can be used for disaster recovery but not for High Availability. Using object store can't guarantee that all instances won't go down at once. So not an appropriate choice.
Reference: https://docs.mulesoft.com/mule-runtime/4.3/mule-object-stores
C) High availability can be achieved by below two models for on-premise MuleSoft implementations.
1) Mule Clustering ñ Where multiple Mule servers are available within the same cluster environment and the routing of requests will be done by the load balancer. A cluster is a set of up to eight servers that act as a single deployment target and high-availability processing unit. Application instances in a cluster are aware of each other, share common information, and synchronize statuses.
If one server fails, another server takes over processing applications. A cluster can run multiple applications. (refer left half of the diagram) In given scenario, it's mentioned that 'data cannot be shared among of different instances'. So this is not

a correct choice.
Reference: https://docs.mulesoft.com/runtime-manager/cluster-about
2) Load balanced standalone Mule instances ñ The high availability can be achieved even without cluster, with the usage of third party load balancer pointing requests to different Mule servers. This approach does not share or synchronize data between Mule runtimes. Also high availability achieved as load balanced algorithms can be implemented using external load balancer. (refer right half of the diagram)

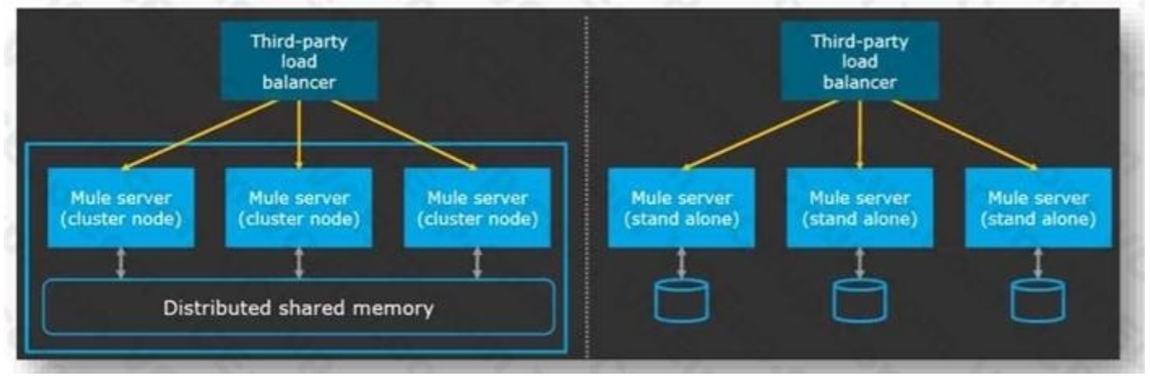

QUESTION 66
An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?

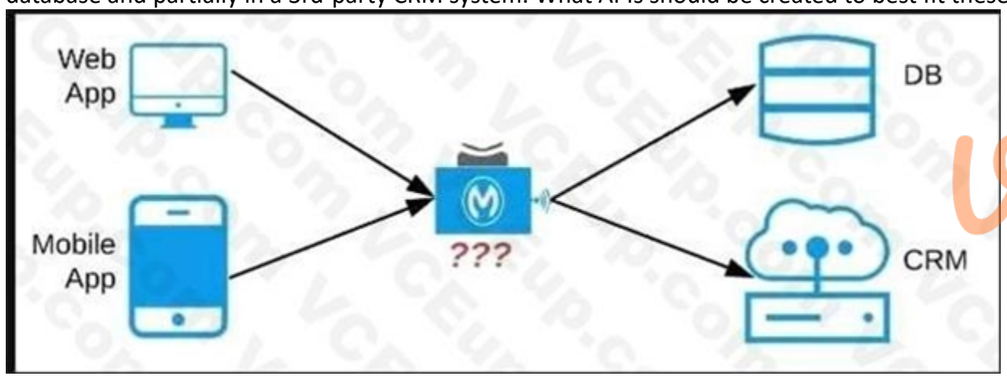

A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
B. One set of APIs (Experience API, Process 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.

Correct Answer: C
Section:
Explanation:
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.
Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.
Hence the correct answer is: 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

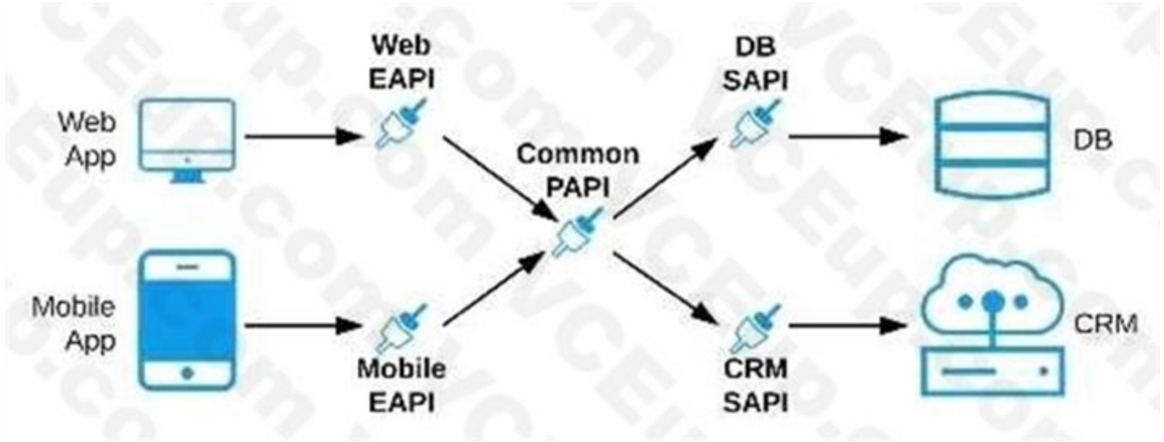

Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.
Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.
Hence the correct answer is: 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

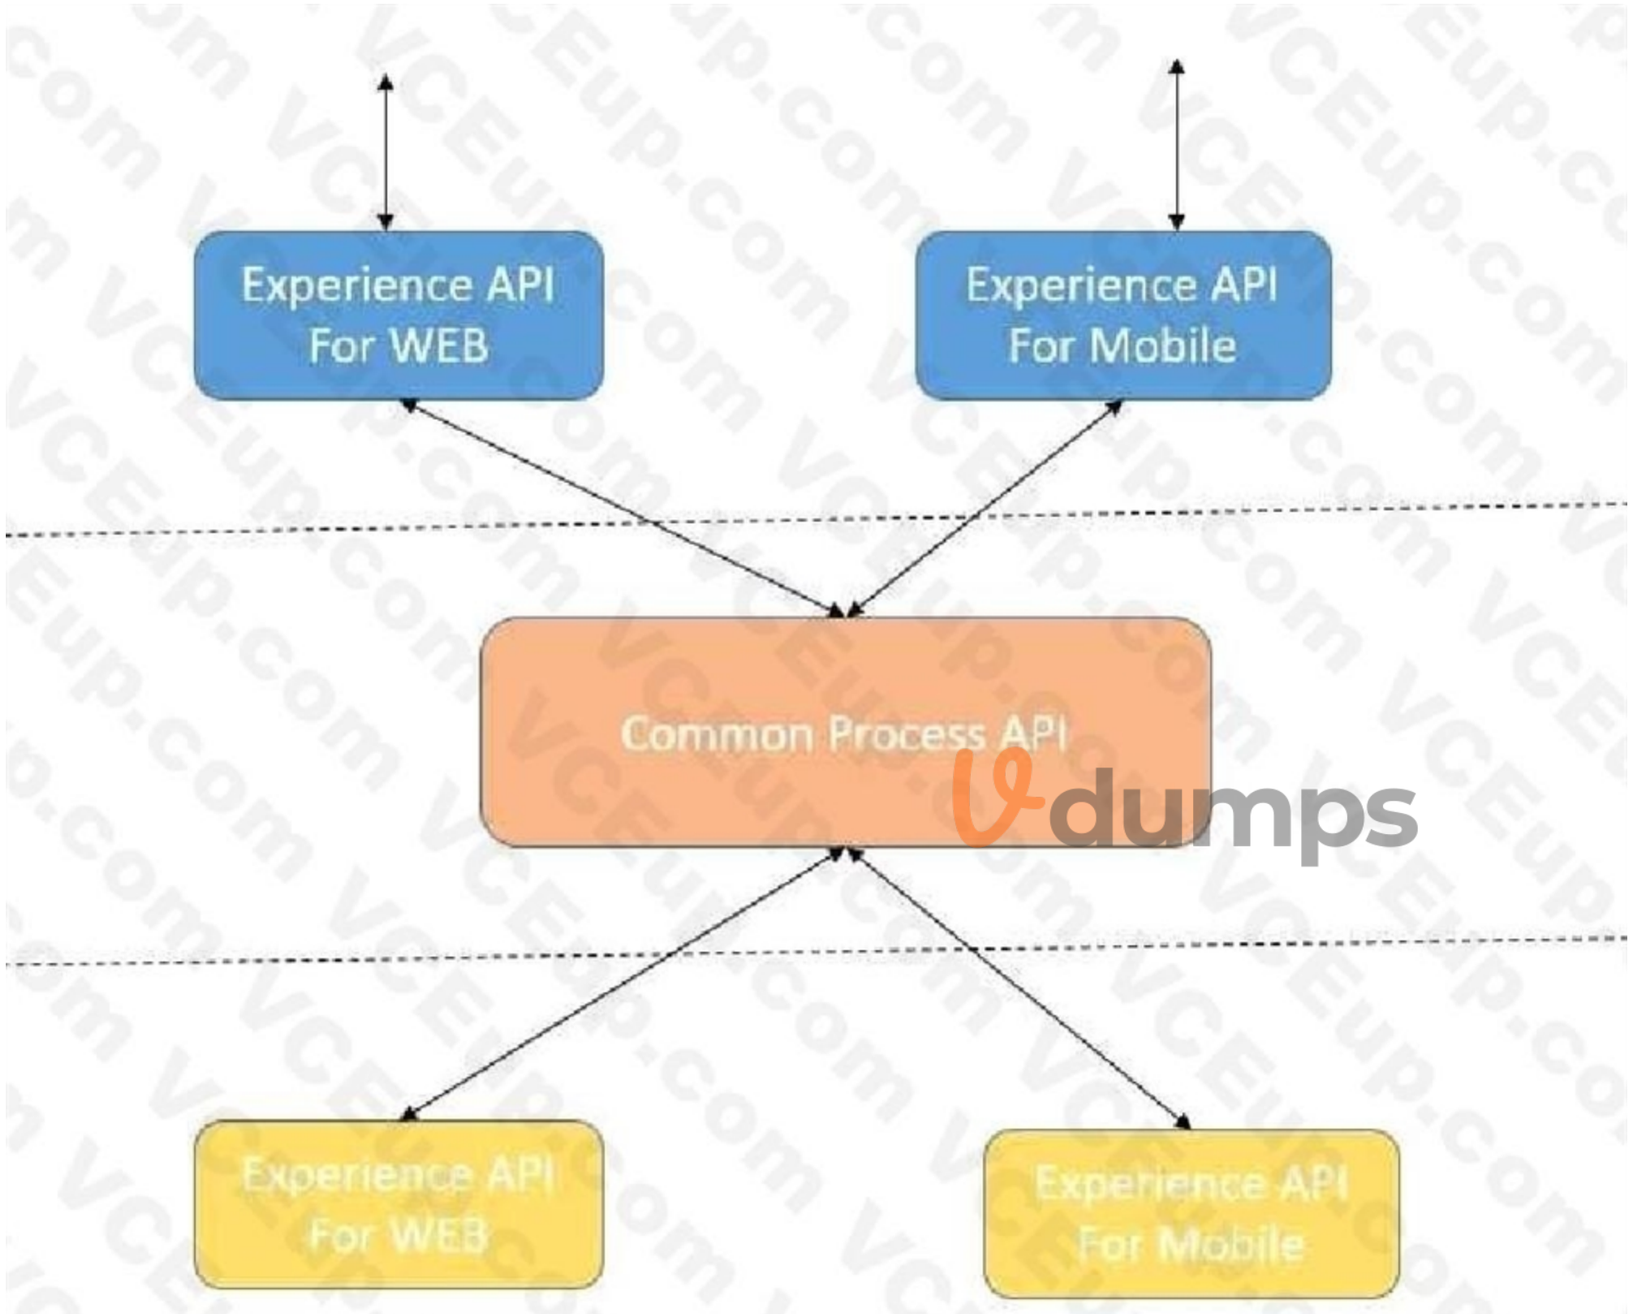

QUESTION 67
What is true about automating interactions with Anypoint Platform using tools such as Anypoint Platform REST API's, Anypoint CLI or the Mule Maven plugin?

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

Correct Answer: A
Section:
Explanation:
Correct answer is By default, the Anypoint CLI and Mule Maven plugin are not included in the Mule runtime Maven is not part of runtime though it is part of studio. You do not need it to deploy in order to deploy your app. Same is the case with CLI.

QUESTION 68
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

Correct Answer: C
Section:
Explanation:
Correct answer is Workers are randomly distributed across available AZs within that region. This ensure high availability for deployed mule applications Mulesoft documentation reference : https://docs.mulesoft.com/runtime-manager/cloudhub-hadr

QUESTION 69
What best describes the Fully Qualified Domain Names (FQDNs), also known as DNS entries, created when a Mule application is deployed to the CloudHub Shared Worker Cloud?

A. A fixed number of FQDNs are created, IRRESPECTIVE of the environment and VPC design
B. The FQDNs are determined by the application name chosen, IRRESPECTIVE of the region
C. The FQDNs are determined by the application name, but can be modified by an administrator after deployment
D. The FQDNs are determined by both the application name and the region

Correct Answer: D
Section:
Explanation:
Every Mule application deployed to CloudHub receives a DNS entry pointing to the CloudHub. The DNS entry is a CNAME for the CloudHub Shared Load Balancer in the region to which the Mule application is deployed. When we deploy the application on CloudHub, we get a generic url to access the endpoints. Generic URL looks as below:
<application-name>.<region>.cloudhub.io <application-name> is the deployed application name which is unique across all the MuleSoft clients. <region> is the region name in which an application is deployed.
The public CloudHub (shared) load balancer already redirects these requests, where myApp is the name of the Mule application deployment to CloudHub: HTTP requests to Error! Hyperlink reference not valid. redirects to Error! Hyperlink reference not valid.
HTTPS traffic to Error! Hyperlink reference not valid. redirects to
Error! Hyperlink reference not valid.

QUESTION 70
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

Correct Answer: D
Section:
Explanation:
Key to this question lies in the fact that Process API are not meant to be accessed directly by clients.
Lets analyze options one by one. Client ID enforcement : This is applied at process API level generally to ensure that identity of API clients is always known and available for API-based analytics Rate Limiting : This policy is applied on Process Level API to secure API's against degradation of service that can happen in case load received is more than it can handle Custom circuit breaker : This is also quite useful feature on process level API's as it saves the API client the wasted time and effort of invoking a failing API. JSON threat protection : This policy is not required at Process API and rather implemented as Experience API's. This policy is used to safeguard application from malicious attacks by injecting malicious code in JSON object. As ideally Process API's are never called from external world , this policy is never used on Process API's Hence correct answer is JSON threat protection MuleSoft Documentation Reference : https:// docs.mulesoft.com/api-manager/2.x/policy-mule3- json-threat

QUESTION 71
What is a key difference between synchronous and asynchronous logging from Mule applications?

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

Correct Answer: B
Section:
Explanation:
Types of logging:
A) Synchronous: The execution of thread that is processing messages is interrupted to wait for the log message to be fully handled before it can continue.
ï The execution of the thread that is processing your message is interrupted to wait for the log message to be fully output before it can continue
ï Performance degrades because of synchronous logging
ï Used when the log is used as an audit trail or when logging ERROR/CRITICAL messages
ï If the logger fails to write to disk, the exception would raise on the same thread that's currently processing the Mule event. If logging is critical for you, then you can rollback the transaction.

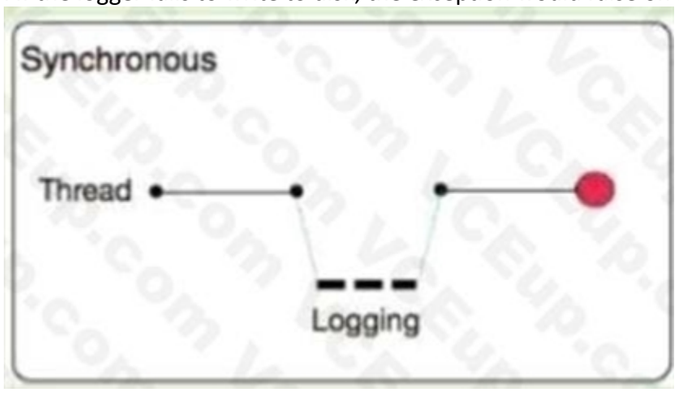

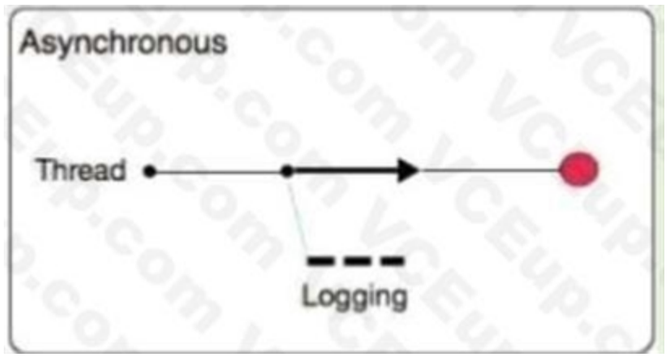

B) Asynchronous:
ï The logging operation occurs in a separate thread, so the actual processing of your message won't be delayed to wait for the logging to complete
ï Substantial improvement in throughput and latency of message processing
ï Mule runtime engine (Mule) 4 uses Log4j 2 asynchronous logging by default
ï The disadvantage of asynchronous logging is error handling.
ï If the logger fails to write to disk, the thread doing the processing won't be aware of any issues writing to the disk, so you won't be able to rollback anything. Because the actual writing of the log gets differed, there's a chance that log messages might never make it to disk and get lost, if Mule were to crash before the buffers are flushed.
---------- So Correct answer is: Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event

QUESTION 72
A global, high-volume shopping Mule application is being built and will be deployed to CloudHub. To improve performance, the Mule application uses a Cache scope that maintains cache state in a CloudHub object store. Web clients will access the Mule application over HTTP from all around the world, with peak volume coinciding with business hours in the web client's geographic location. To achieve optimal performance, what Anypoint Platform region should be chosen for the CloudHub object store?

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

Correct Answer: A
Section:
Explanation:
CloudHub object store should be in same region where the Mule application is deployed. This will give optimal performance.
Before learning about Cache scope and object store in Mule 4 we understand what is in general Caching is and other related things.
WHAT DOES "CACHING" MEAN?
Caching is the process of storing frequently used data in memory, file system or database which saves processing time and load if it would have to be accessed from original source location every time.
In computing, a cache is a high-speed data storage layer which stores a subset of data, so that future requests for that data are served up faster than is possible by accessing the data's primary storage location. Caching allows you to efficiently reuse previously retrieved or computed data.
How does Caching work?
The data in a cache is generally stored in fast access hardware such as RAM (Random-access memory) and may also be used in correlation with a software component. A cache's primary purpose is to increase data retrieval performance by reducing the need to access the underlying slower storage layer.
Caching in MULE 4
In Mule 4 caching can be achieved in mule using cache scope and/or object-store. Cache scope internally uses Object Store to store the data.
What is Object Store Object Store lets applications store data and states across batch processes, Mule components, and applications, from within an application. If used on cloud hub, the object store is shared between applications deployed on Cluster.
Cache Scope is used in below-mentioned cases:
ï Need to store the whole response from the outbound processor
ï Data returned from the outbound processor does not change very frequently
ï As Cache scope internally handle the cache hit and cache miss scenarios it is more readable Object Store is used in below-mentioned cases:

ï Need to store custom/intermediary data
ï To store watermarks
ï Sharing the data/stage across applications, schedulers, batch.
If CloudHub object store is in same region where the Mule application is deployed it will aid in fast access of data and give optimal performance.

QUESTION 73
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, customerprovided, 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

Correct Answer: B
Section:
Explanation:
Correct answer is Only MuleSoft provided certificates can be used for server side certificate * The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
* You would need to use dedicated load balancer which can enable you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
* To use a dedicated load balancer in your environment, you must first create an Anypoint VPC.
Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.
Additional Info on SLB Vs DLB:

	Shared Load Balancer	Dedicated Load Balancer
VPC	Shared VPC (Mulesoft)	VPC (Customer)
Default Load Balancer	Cloudhub provides Deault Shared Load Balancer available in All Environment	Need to Purchase
Organization Use	Multiple Oragnization	Specific to Organization
Certificate	Mulesoft Certificate	Organization Certificate
TLS Support	Yes	Yes
URL Mapping	Fixed URL Mapping	Customer URL Mapping
Timeout	30 Sec Session Timeout	Custom Timeout
Ports	Public Port {80 : 8081, 443 : 8082}	Private Port {80 : 8091, 443 : 8092}
Fashion	Round Robin	Round Robin
Supports HTTPS Protocol	Yes	Yes
Worker Assignment	No	Yes
IP Blacklisting/ Whitelisting	No	Yes
	https://docs.mulesoft.com/runtime-manager/lb-whitelists	
Configure Custom Domain	No	Yes
Custom Certificate	No	Yes
Rate Limit	Lower Rate Limit and applied According to Region	Higher Rate Limit Threshold
VPC	Anypoint VPC optional	Can't Use DLB without Anypoint VPC

QUESTION 74
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.

Correct Answer: C
Section:
Explanation:

Object Store Connector is a Mule component that allows for simple key-value storage. Although it can serve a wide variety of use cases, it is mainly design for: - Storing synchronization information, such as watermarks. - Storing temporal information such as access tokens. - Storing user information.
Additionally, Mule Runtime uses Object Stores to support some of its own components, for example:
- The Cache module uses an Object Store to maintain all of the cached data. - The OAuth module (and every OAuth enabled connector) uses Object Stores to store the access and refresh tokens. Object Store data is in the same region as the worker where the app is initially deployed. For example, if you deploy to the Singapore region, the object store persists in the Singapore region. MuleSoft Reference : https://docs.mulesoft.com/object-store-connector/1.1/ Data can be shared between different instances of the Mule application. This is not recommended for Inter Mule app communication.
Coming to the question, object store cannot be used to share cached data if it is deployed as separate Mule applications or deployed under separate Business Groups. Hence correct answer is When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.

QUESTION 75
An organization has several APIs that accept JSON data over HTTP POST. The APIs are all publiclyavailable and are associated with several mobile applications and web applications. The organizationdoes 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 theapplications or servers running the API implementations.
What out-of-the-box Anypoint Platformpolicy 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

Correct Answer: D
Section:
Explanation:
We need to note few things about the scenario which will help us in reaching the correct solution.
Point 1 : The APIs are all publicly available and are associated with several mobile applications and web applications. This means Apply an IP blacklist policy is not viable option. as blacklisting IPs is limited to partial web traffic. It can't be useful for traffic from mobile application Point 2 : The organization does NOT want to use any authentication or compliance policies for these APIs. This means we can not apply HTTPS mutual authentication scheme.
Header injection or removal will not help the purpose.
By its nature, JSON is vulnerable to JavaScript injection. When you parse the JSON object, the malicious code inflicts its damages. An inordinate increase in the size and depth of the JSON payload can indicate injection. Applying the
JSON threat protection policy can limit the size of your JSON payload and thwart recursive additions to the JSON hierarchy.
Hence correct answer is Apply a JSON threat protection policy to all APIs to detect potential threat vectors

QUESTION 76
A new upstream API Is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity. The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms. If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API's desired SLA?

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

Correct Answer: D
Section:
Explanation:
Before we answer this question , we need to understand what median (50th percentile) and 80th percentile means. If the 50th percentile (median) of a response time is 500ms that means that 50% of my transactions are either as fast or faster than 500ms.
If the 90th percentile of the same transaction is at 1000ms it means that 90% are as fast or faster and only 10% are slower. Now as per upstream SLA , 99th percentile is 800 ms which means 99% of the incoming requests should have response time less than or equal to 800 ms. But as per one of the backend API , their 95th percentile is 1000 ms which means that backend API will take 1000 ms or less than that for 95% of. requests. As there are

three API invocation from upstream API , we can not conclude a timeout that can be set to meet the desired SLA as backend SLA's do not support it.
Let see why other answers are not correct.
1) Do not set a timeout --> This can potentially violate SLA's of upstream API
2) Set a timeout of 100 ms; ---> This will not work as backend API has 100 ms as median meaning only 50% requests will be answered in this time and we will get timeout for 50% of the requests.
Important thing to note here is, All APIs need to be executed sequentially, so if you get timeout in first API, there is no use of going to second and third API. As a service provider you wouldn't want to keep 50% of your consumers dissatisfied. So not the best option to go with.
*To quote an example: Let's assume you have built an API to update customer contact details.
- First API is fetching customer number based on login credentials
- Second API is fetching Info in 1 table and returning unique key
- Third API, using unique key provided in second API as primary key, updating remaining details * Now consider, if API times out in first API and can't fetch customer number, in this case, it's useless to call API 2 and 3 and that is why question mentions specifically that all APIs need to be executed sequentially.
3) Set a timeout of 50 ms --> Again not possible due to the same reason as above Hence correct answer is No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

QUESTION 77
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.

Correct Answer: D
Section:
Explanation:
* Semantic Versioning is a 3-component number in the format of X.Y.Z, where :
X stands for a major version.
Y stands for a minor version:
Z stands for a patch.
So, SemVer is of the form Major.Minor.Patch Coming to our question , minor version of the API has been changed which is backward compatible. Hence there is no change required on API client end. If they want to make use of new featured that have been added as a part of minor version change they may need to change code at their end. Hence correct answer is The API client code ONLY needs to be changed if it needs to take advantage of new features.

QUESTION 78
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.

Correct Answer: B
Section:
Explanation:
An SLA for the upstream API CANNOT be provided.

QUESTION 79
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

Correct Answer: A
Section:
Explanation:
Correct answer is "Compile, package, unit test, validate unit test coverage, deploy" : Anypoint Platform supports continuous integration and continuous delivery using industry standard tools Mule Maven Plugin The Mule Maven plugin can automate building, packaging and deployment of Mule applications from source projects Using the Mule Maven plugin, you can automate your Mule application deployment to CloudHub, to Anypoint Runtime Fabric, or on-premises, using any of the following deployment strategies ï CloudHub deployment ï Runtime Fabric deployment ï Runtime Manager REST API deployment ï Runtime Manager agent deployment MUnit Maven Plugin The MUnit Maven plugin can automate test execution, and ties in with the Mule Maven plugin. It provides a full suite of integration and unit test capabilities, and is fully integrated with Maven and Surefire for integration with your continuous deployment environment. Since MUnit 2.x, the coverage report goal is integrated with the maven reporting section. Coverage Reports are generated during Maven's site lifecycle, during the coverage-report goal. One of the features of MUnit Coverage is to fail the build if a certain coverage level is not reached. MUnit is not used for integration testing Also publishing to Anypoint Exchange or to create associated

API instances in API Manager is not a part of CICD pipeline which can ne achieved using mulesoft provided maven plugin Architecture mentioned in the question can be diagrammatically put as below. Persistent Object Store is the correct answer .
* Mule Object Stores: An object store is a facility for storing objects in or across Mule applications.
Mule uses object stores to persist data for eventual retrieval.
Mule provides two types of object stores:
1) In-memory store ñ stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime. So we cant use in memory store in our scenario as we want to share watermark within all cloudhub workers
2) Persistent store ñ Mule persists data when an object store is explicitly configured to be persistent.
Hence this watermark will be available even any of the worker goes down

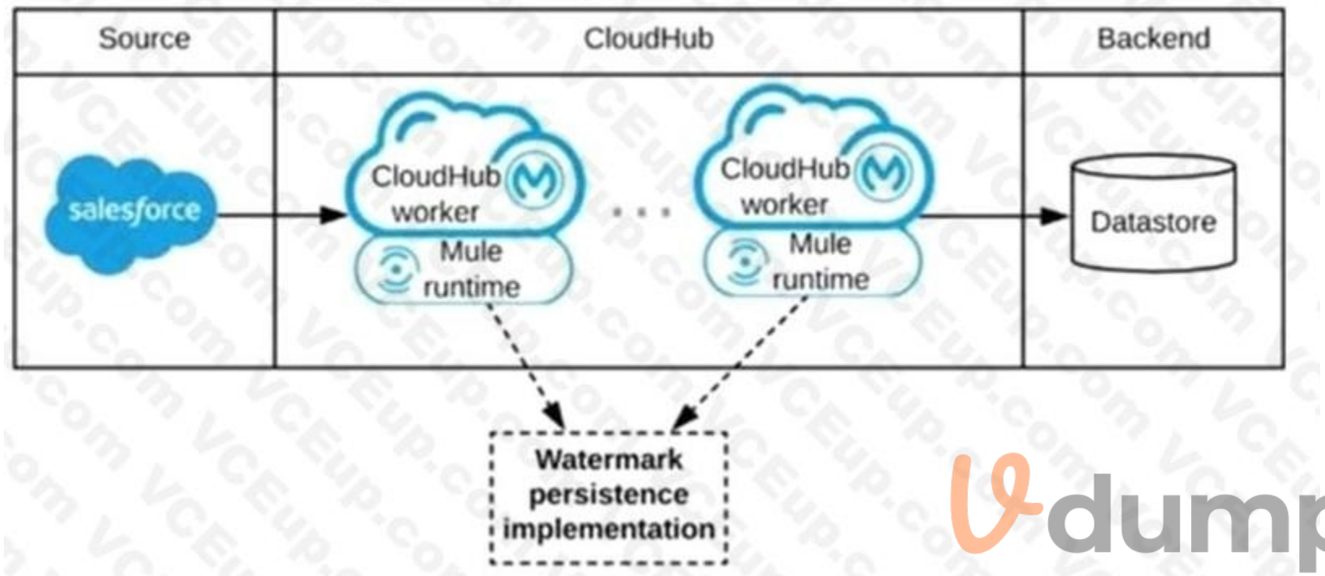

QUESTION 80
What condition requires using a CloudHub Dedicated Load Balancer?

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

Correct Answer: D
Section:
Explanation:
Correct answer is When server-side load-balanced TLS mutual authentication is required between API implementations and API clients CloudHub dedicated load balancers (DLBs) are an optional component of Anypoint Platform that enable you to route external HTTP and HTTPS traffic to multiple Mule applications deployed to CloudHub workers in a Virtual Private Cloud (VPC). Dedicated load balancers enable you to: * Handle load balancing among the different CloudHub workers that run your application. * Define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication. * Configure proxy rules that map your applications to custom domains. This enables you to host your applications under a single domain

QUESTION 81
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

Correct Answer: B
Section:
Explanation:
Correct answer is Mule correlation ID By design, Correlation Ids cannot be changed within a flow in Mule 4 applications and can be set only at source. This ID is part of the Event Context and is generated as soon as the message is received by the application. When a HTTP Request is received, the request is inspected for "X-Correlation-Id" header. If "X-Correlation-Id" header is present, HTTP connector uses this as the Correlation Id. If "X-Correlation-Id" header is NOT present, a Correlation Id is randomly generated. For Incoming HTTP Requests: In order to set a custom Correlation Id, the client invoking the HTTP request must set "X-Correlation-Id" header. This will ensure that the Mule Flow uses this Correlation Id. For Outgoing HTTP Requests: You can also propagate the existing Correlation Id to downstream APIs. By default, all outgoing HTTP Requests send "X-Correlation-Id" header. However, you can choose to set a different value to "X-Correlation-Id" header or set "Send Correlation Id" to NEVER.

QUESTION 82
Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improved throughput of the system. But it was observed over the period of time that few of the important exception log messages which were used to rollback transactions are not working as expected causing huge loss to the Organization. Organization wants to avoid these losses. Application also has constraints due to which they cant compromise on throughput much.
What is the possible option in this case?

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

Correct Answer: D
Section:
Explanation:
Correct approach is to use Mixed configuration of asynchronous or synchronous loggers shoud be used to log exceptions via synchronous way Asynchronous logging poses a performance-reliability trade-off. You may lose some messages if Mule crashes before the logging buffers flush to the disk. In this case, consider that you can have a mixed configuration of asynchronous or synchronous loggers in your app. Best practice is to use asynchronous logging over synchronous with a minimum logging level of WARN for a production application. In some cases, enable INFO logging level when you need to confirm events such as successful policy installation or to perform troubleshooting. Configure your logging strategy by editing your application's src/main/resources/log4j2.xml file

QUESTION 83
As a part of business requirement , old CRM system needs to be integrated using Mule application.
CRM system is capable of exchanging data only via SOAP/HTTP protocol. As an integration architectwho follows API led approach , what is the the below step you will perform so that you can sharedocument with CRM team?

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

Correct Answer: C
Section:
Explanation:
Correct answer is Create WSDL specification using text editor SOAP services are specified using WSDL. A client program connecting to a web service can read the WSDL to determine what functions are available on the server. We can not create WSDL specification in Design Center. We need to use external text editor to create WSDL.

QUESTION 84
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 deployment. External log management system is required.
B. When deploying an application to CloudHub , logs retention period should be selected as 2 years
C. When deploying an application to CloudHub, worker size should be sufficient to store 2 years data
D. Logging strategy should be configured accordingly in log4j file deployed with the application.

Correct Answer: A
Section:
Explanation:
Correct answer is It is not possible to store logs for 2 years in CloudHub deployment. External log management system is required. CloudHub has a specific log retention policy, as described in the documentation: the platform stores logs of up to 100 MB per app & per worker or for up to 30 days, whichever limit is hit first. Once this limit has been reached, the oldest log information is deleted in chunks and is irretrievably lost. The recommended approach is to persist your logs to a external logging system of your choice (such as Splunk, for instance) using a log appender. Please note that this solution results in the logs no longer being stored on our platform, so any support cases you lodge will require for you to provide the appropriate logs for review and case resolution

QUESTION 85
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

Correct Answer: C
Section:

QUESTION 86
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 partner PGP 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 partner TLS 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

Correct Answer: D
Section:

QUESTION 87
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 tomeet 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

Correct Answer: A
Section:

QUESTION 88
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

Correct Answer: C
Section:

QUESTION 89
A company is implementing a new Mule application that supports a set of critical functions driven by a rest API enabled, claims payment rules engine hosted on oracle ERP. As designed the mule application requires many data transformation operations as it performs its batch processing logic.
The company wants to leverage and reuse as many of its existing java-based capabilities (classes, objects, data model etc.) as possible What approach should be considered when implementing required data mappings and transformations between Mule application and Oracle ERP in the new Mule application?

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

Correct Answer: C
Section:

QUESTION 90
An insurance company has an existing API which is currently used by customers. API is deployed to customer hosted Mule runtime cluster. The load balancer that is used to access any APIs on the mule cluster is only configured to point to applications hosted on the server at port 443.
Mule application team of a company attempted to deploy a second API using port 443 but the application will not start and checking logs shows an error indicating the address is already in use.
Which steps must the organization take to resolve this error and allow customers to access both the API's?

A. Change the base path of the HTTP listener configuration in the second API to a different one fromthe 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 projectusing port 443
D. Set the HTTP listener of the second API to use different port than the one used in the first API

Correct Answer: C

Section:

QUESTION 91
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

Correct Answer: C, D
Section:

QUESTION 92
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 processing
There 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 processing
There 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

Correct Answer: C
Section:

QUESTION 93
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

Correct Answer: B
Section:

QUESTION 94
Refer to the exhibit.

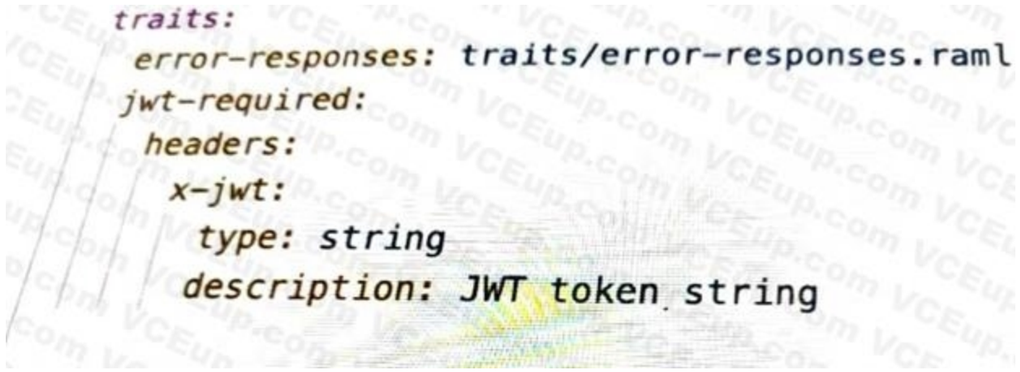

What is the type data format shown in the exhibit?

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

Correct Answer: A
Section:

QUESTION 95
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 API! 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

Correct Answer: C
Section:

QUESTION 96
Which type of communication is managed by a service mesh in a microservices architecture?

A. Communication between microservices runtime administrators
B. Communication between microservices developers
C. Communication between microservices
D. Communication between trading partner services

Correct Answer: C
Section:

QUESTION 97
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

Correct Answer: D
Section:

QUESTION 98
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

Correct Answer: C
Section:

QUESTION 99
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

Correct Answer: B
Section:

QUESTION 100
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

Correct Answer: B
Section:

QUESTION 101
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

Correct Answer: B
Section:

QUESTION 102
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

Correct Answer: A
Section:

QUESTION 103
According to MuleSoft, which major benefit does a Center for Enablement (C4E) provide for an enterprise and its lines of business?

A. Enabling Edge security between the lines of business and public devices
B. Centralizing project management across the lines of business
C. Centrally managing return on investment (ROI) reporting from lines of business to leadership
D. Accelerating self-service by the lines of business

Correct Answer: B
Section:

QUESTION 104
According to MuleSoft's recommended REST conventions, which HTTP method should an API use tospecify how AP\ clients can request data from a specified resource?

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

Correct Answer: D
Section:

QUESTION 105
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

Correct Answer: C
Section:
Explanation:
Anypoint Platform offers universal API management and Integration-Platform-as-a-Service (iPaaS) capabilities in a unified platform, meaning that it provides a single control plane to manage both fulllifecycle API management and integration. This allows organizations to easily manage their APIs and integrations, as well as deploy APIs and integrations quickly and efficiently. According to the MuleSoft Certified Integration Architect - Level 1 Course Book, "Anypoint Platform provides a unified platform for managing, deploying, and monitoring both API and integration solutions, allowing organizations to quickly and easily build and manage their APIs and integrations."

QUESTION 106
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

Correct Answer: D
Section:

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

Correct Answer: D
Section:

QUESTION 108
According to the Internet Engineering Task Force (IETF), which supporting protocol does File Transfer Protocol (FTP) use for reliable communication?

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

Correct Answer: B
Section:

QUESTION 109
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

Correct Answer: C, E
Section:

QUESTION 110
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

Correct Answer: B
Section:

QUESTION 111
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

Correct Answer: B
Section:

QUESTION 112
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

Correct Answer: C
Section:

QUESTION 113
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

Correct Answer: A
Section:

QUESTION 114
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

Correct Answer: B
Section:

QUESTION 115
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

Correct Answer: D
Section:

QUESTION 116
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.

Correct Answer: C
Section:

QUESTION 117
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)

Correct Answer: A
Section:

QUESTION 118
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

Correct Answer: A
Section:

QUESTION 119
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

Correct Answer: A
Section:

QUESTION 120
An API client makes an HTTP request to an API gateway with an Accept header containing the value''application''.
What is a valid HTTP response payload for this request in the client requested data format?

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

Correct Answer: B
Section:

QUESTION 121
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

Correct Answer: C
Section:

QUESTION 122
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

Correct Answer: D
Section:

QUESTION 123
According to MuleSoft, what is a major distinguishing characteristic of an application network in relation to the integration of systems, data, and devices?

A. It uses a well-organized monolithic approach with standards
B. It is built for change and self-service
C. It leverages well-accepted internet standards like HTTP and JSON
D. It uses CI/CD automation for real-time project delivery

Correct Answer: B
Section:

QUESTION 124
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-toend 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

Correct Answer: A, B
Section:

QUESTION 125
An organization's IT team follows an API-led connectivity approach and must use Anypoint Platform to implement a System AP\ that securely accesses customer dat a. The organization uses Salesforce as the system of record for all customer data, and its most important objective is to reduce the overall development time to release the System API.
The team's integration architect has identified four different approaches to access the customer data from within the implementation of the System API by using different Anypoint Connectors that all meet the technical requirements of the project.

A. Use the Anypoint Connector for Database to connect to a MySQL database to access a copy of the customer data
B. Use the Anypoint Connector for HTTP to connect to the Salesforce APIs to directly access thecustomer data
C. Use the Anypoint Connector for Salesforce to connect to the Salesforce APIs to directly access the customer data
D. Use the Anypoint Connector tor FTP to download a file containing a recent near-real time extract of the customer data

Correct Answer: C
Section:

QUESTION 126
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

Correct Answer: A
Section:

QUESTION 127
According to MuleSoft, a synchronous invocation of a RESTful API using HTTP to get an individualcustomer record from a single system is an example of which system integration interaction pattern?

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

Correct Answer: A
Section:

QUESTION 128
Cloud Hub is an example of which cloud computing service model?

A. Platform as a Service (PaaS)
B. Software as a Service (SaaS)
C. Monitoring as a Service (MaaS)
D. Infrastructure as a Service (laaS)

Correct Answer: A
Section:

QUESTION 129
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

Correct Answer: B
Section:

QUESTION 130
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

Correct Answer: D
Section:

QUESTION 131
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

Correct Answer: C
Section:

QUESTION 132
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

Correct Answer: D
Section:

QUESTION 133
An application load balancer routes requests to a RESTful web API secured by Anypoint Flex Gateway.
Which protocol is involved in the communication between the load balancer and the Gateway?

A. SFTP

B. HTTPS
C. LDAP
D. SMTP

Correct Answer: B
Section:

QUESTION 134
A key Cl/CD capability of any enterprise solution is a testing framework to write and run repeatable tests.
Which component of Anypoint Platform provides the te6t automation capabilities for customers to use in their pipelines?

A. Anypoint CLl
B. Mule Maven Plugin
C. Exchange Mocking Service
D. MUnit

Correct Answer: D
Section:

QUESTION 135
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

Correct Answer: B
Section:

QUESTION 136
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.

Correct Answer: A
Section:
Explanation:
The advantage of using OAuth 2.0 client credentials and access tokens over only API keys for API authentication is that if the access token is compromised, the client credentials do not have to be reissued.
OAuth 2.0 is a secure protocol for authenticating clients and authorizing them to access protected resources. It works by having the client authenticate with the authorization server and receive an access token, which is then used to authenticate requests to the API. If the access token is compromised, it can be revoked and replaced without needing to reissue the client credentials.
Reference: MuleSoft Certified Integration Architect - Level 1 Official Text Book and Resources:
Chapter 7: Security

Section 7.2: OAuth 2.0

QUESTION 137
An organization is choosing between API-led connectivity and other integration approaches.
According to MuleSoft, which business benefits is associated with an API-led connectivity approach using Anypoint Platform?

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

Correct Answer: B
Section:

QUESTION 138
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

Correct Answer: A
Section:

QUESTION 139
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 bespecified to indicate when client requests are successfully received, understood, and accepted by theweb service?

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

Correct Answer: B
Section:

QUESTION 140
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

Correct Answer: A

Section:

QUESTION 141
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

Correct Answer: A
Section: