Number: MCIA-LEVEL-1-MAINTENANCE
Passing Score: 800.0
Time Limit: 120.0
File Version: 1.7

Exam Code: MCIA-LEVEL-1-MAINTENANCE
Exam Name: MuleSoft Certified Integration Architect - Level 1 MAINTENANCE
Team-Support: https://VCEplus.io/

Exam A

QUESTION 1
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 2
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 3

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 4
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 5
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 6
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 7
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 8
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 9
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 10
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

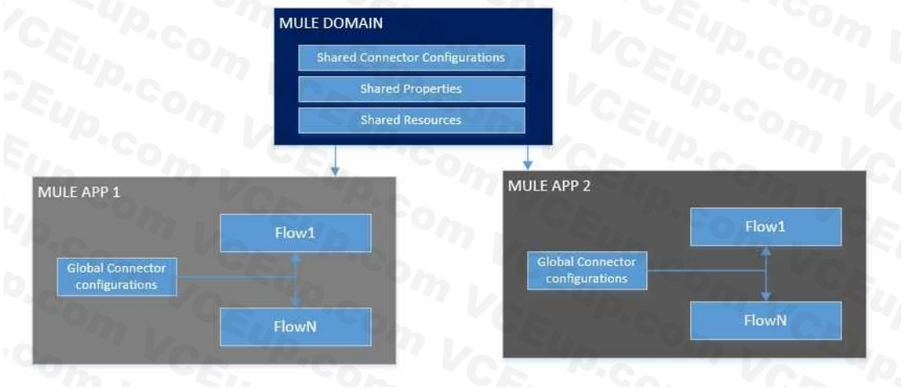

QUESTION 11
An API implementation is being designed that must invoke an Order API which is known to repeatedly experience downtime. For this reason a fallback API is to be called when the Order API is unavailable. What approach to designing invocation of the fallback API provides the best resilience?

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 — even after a certain number of retries — 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. Thesecodes 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

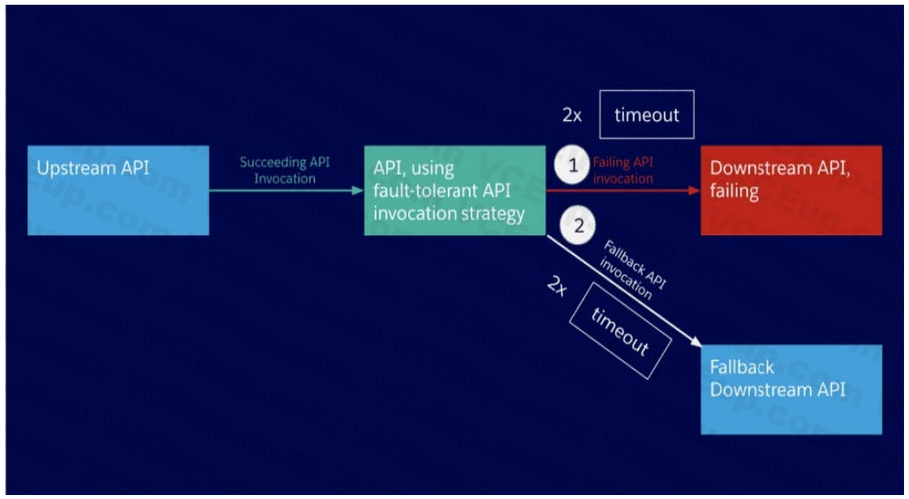

Hence correct answer is Redirect client requests through an HTTP 303 temporary redirect status codeto the fallback API whenever the Order API is unavailable

QUESTION 12
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 • An application component • that accesses a service • by invoking an API of that service - by definition of the term API over HTTP API consumer • A business role, which is often assigned to an individual • that develops API clients, i.e., performs the activities necessary for enabling an API client to invoke APIs API implementation • An application component • that implements the functionality

QUESTION 13
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 14

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 15
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 – 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 – 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)

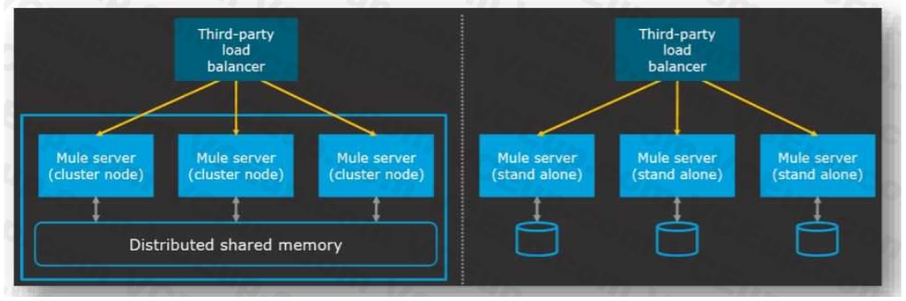

QUESTION 16
An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?

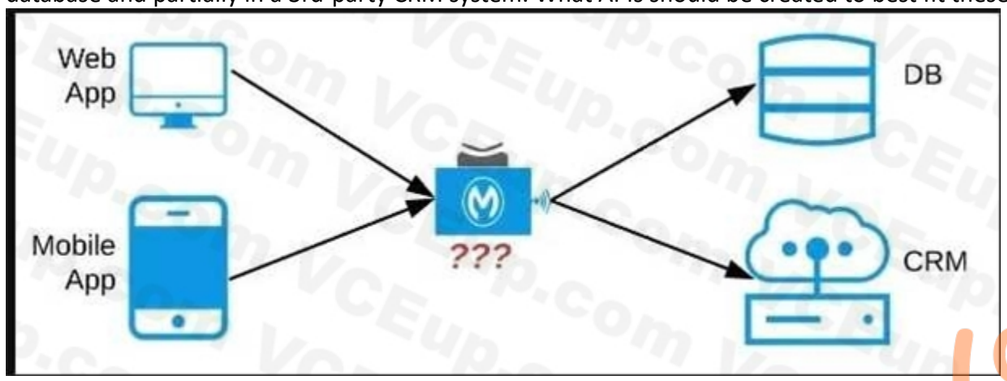

A.
A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.

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

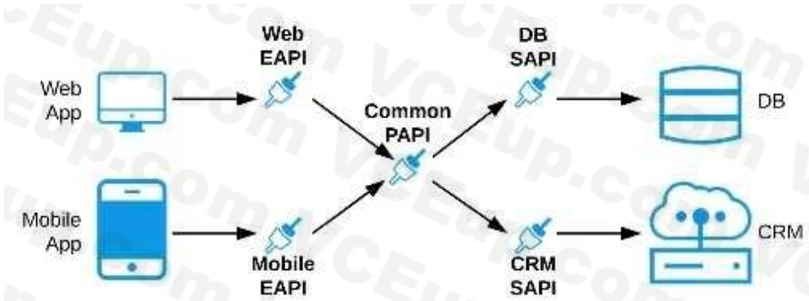

Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements 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

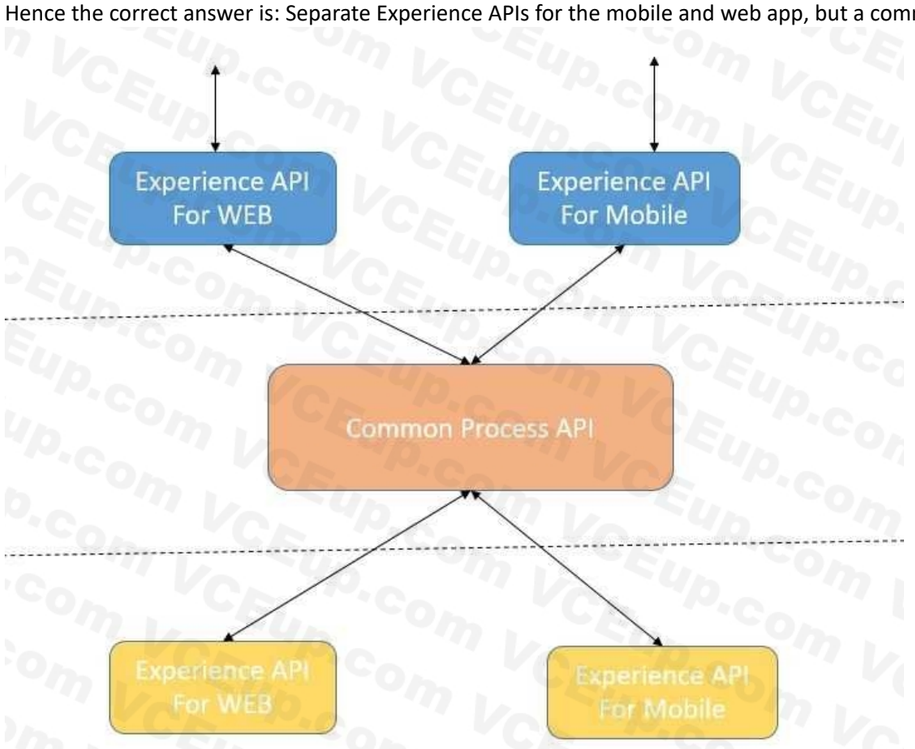

QUESTION 17
What is true about automating interactions with Anypoint Platform using tools such as Anypoint Platform REST API's, Anypoint CLI or the Mule Maven plugin?

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 18
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 19
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 http://myApp.<region>.cloudhub.io redirects to http://mule-worker-myApp.<region>.cloudhub.io:8081
HTTPS traffic to https://myApp.<region>.cloudhub.io redirects to
https://mule-worker-myApp.<region>.cloudhub.io:8082

QUESTION 20
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 21
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.
- 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
- Performance degrades because of synchronous logging
- Used when the log is used as an audit trail or when logging ERROR/CRITICAL messages
- 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.

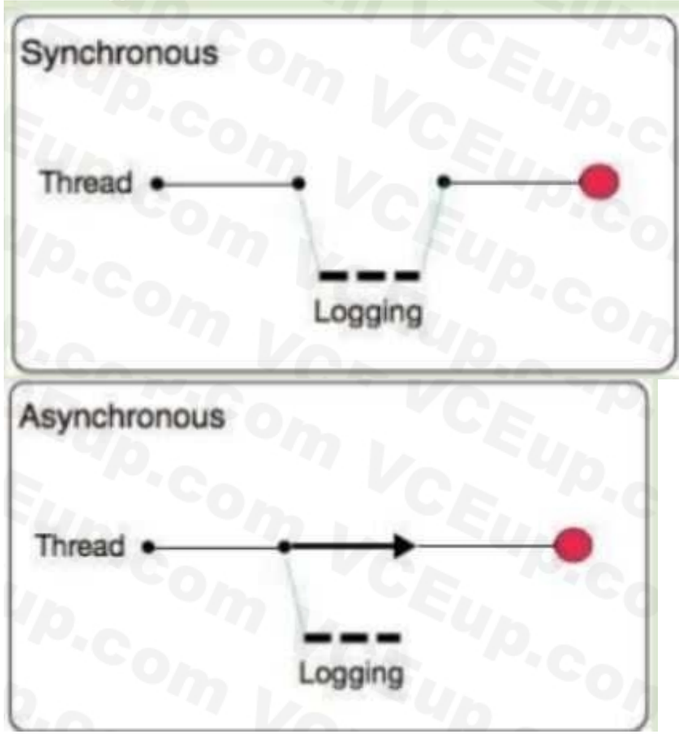

B) Asynchronous:

- 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
- Substantial improvement in throughput and latency of message processing
- Mule runtime engine (Mule) 4 uses Log4j 2 asynchronous logging by default
- The disadvantage of asynchronous logging is error handling.
- 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 22

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

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:

- Need to store the whole response from the outbound processor
- Data returned from the outbound processor does not change very frequently ? As Cache scope internally handle the cache hit and cache miss scenarios it is more readable Object Store is used in below-mentioned cases:
- Need to store custom/intermediary data
- To store watermarks
- 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 23

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 24
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 25
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 26
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 27
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.

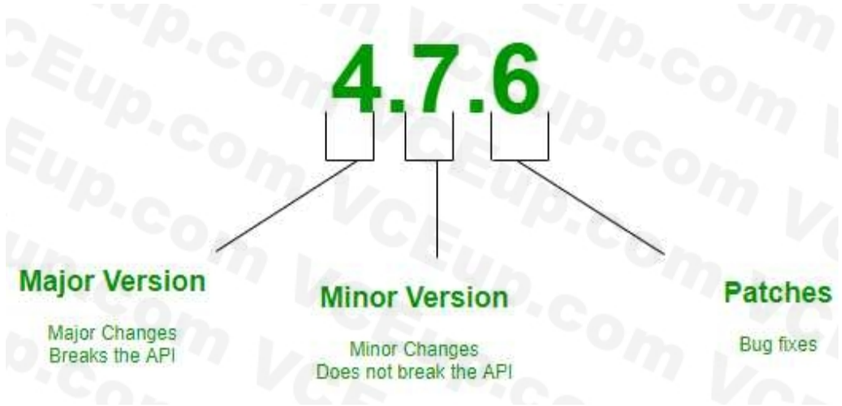

QUESTION 28
When designing an upstream API and its implementation, the development team has been advised to not set timeouts when invoking downstream API. Because the downstream API has no SLA that can be relied upon. This is the only donwstream API dependency of that upstream API. Assume the downstream 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 29
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 • CloudHub deployment • Runtime Fabric deployment • Runtime Manager REST API deployment • 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 – 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 – 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

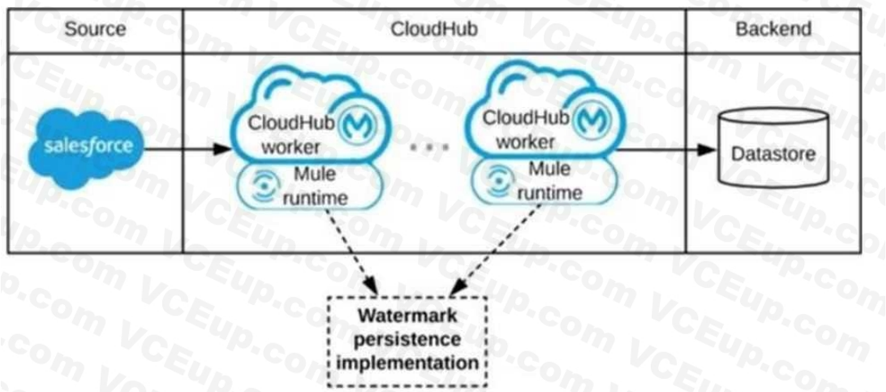

QUESTION 30
What condition requires using a CloudHub Dedicated Load Balancer?

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 31
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 32
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 33
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 34
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 35
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 36
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 37
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 38
A Mule application is synchronizing customer data between two different database systems.
What is the main benefit of using eXtended Architecture (XA) transactions over local transactions to synchronize these two different database systems?

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

Correct Answer: B
Section:
Explanation:
Reference: https://docs.oracle.com/middleware/1213/wls/PERFM/llrtune.htm#PERFM997

QUESTION 39
An organization has implemented a continuous integration (CI) lifecycle that promotes Mule applications through code, build, and test stages. To standardize the organization's CI journey, a new dependency control approach is being designed to store artifacts that include information such as dependencies, versioning, and build promotions.
To implement these process improvements, the organization will now require developers to maintain all dependencies related to Mule application code in a shared location.
What is the most idiomatic (used for its intended purpose) type of system the organization should use in a shared location to standardize all dependencies related to Mule application code?

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

Correct Answer: C
Section:

QUESTION 40
An organization has deployed both Mule and non-Mule API implementations to integrate its customer and order management systems. All the APIs are available to REST clients on the public internet.

The organization wants to monitor these APIs by running health checks: for example, to determine if an API can properly accept and process requests. The organization does not have subscriptions to any external monitoring tools and also does not want to extend its IT footprint.
What Anypoint Platform feature provides the most idiomatic (used for its intended purpose) way to monitor the availability of both the Mule and the non-Mule API implementations?

A. API Functional Monitoring
B. Runtime Manager
C. API Manager
D. Anypoint Visualizer

Correct Answer: D
Section:
Explanation:
Reference: https://docs.mulesoft.com/visualizer/

QUESTION 41
The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure.
Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).
Given the company's current setup and requirements, what is the most idiomatic (used for its intended purpose) way to send Mule application logs to the external LMT?

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

Correct Answer: A
Section:
Explanation:
Reference: https://help.mulesoft.com/s/article/Enable-external-log-forwarding-for-Muleapplications-deployed-in-RTF

QUESTION 42
An organization is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to rejections from the back-end system will need to be processed manually (outside the back-end system).
The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization’s firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.
What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

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

An Until Successful scope containing Object Store components for long retries A dead-letter object store configured in the Mule application

Correct Answer: A
Section:

QUESTION 43
A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should be published in a way that is guaranteed in normal situations and also minimizes duplicate delivery in less frequent failure scenarios.
The organizational firewall is configured to only allow outbound traffic on ports 80 and 443. Some external event consumers are within the organizational network, while others are located outside the firewall.
What Anypoint Platform service is most idiomatic (used for its intended purpose) for publishing these Mule events to all external consumers while addressing the desired reliability goals?

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

Correct Answer: B
Section:
Explanation:
Set the Anypoint MQ connector operation to publish or consume messages, or to accept (ACK) or not accept (NACK) a message.
Reference: https://docs.mulesoft.com/mq/

QUESTION 44
A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.
The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.
What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

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

Correct Answer: D
Section:
Explanation:
Reference: https://docs.spring.io/spring-framework/docs/4.2.x/spring-frameworkreference/html/integration-testing.html

QUESTION 45
In Anypoint Platform, a company wants to configure multiple identity providers (IdPs) for multiple lines of business (LOBs). Multiple business groups, teams, and environments have been defined for these LOBs.
What Anypoint Platform feature can use multiple IdPs across the company's business groups, teams, and environments?

A. MuleSoft-hosted (CloudHub) dedicated load balancers
B. Client (application) management
C. Virtual private clouds
D. Permissions

Correct Answer: A

Section:
Explanation:
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.
Reference: https://docs.mulesoft.com/runtime-manager/cloudhub-dedicated-load-balancer

QUESTION 46
An external web UI application currently accepts occasional HTTP requests from client web browsersto change (insert, update, or delete) inventory pricing information in an inventory system's database.
Each inventory pricing change must be transformed and then synchronized with multiple customer experience systems in near real-time (in under 10 seconds). New customer experience systems are expected to be added in the future.
The database is used heavily and limits the number of SELECT queries that can be made to the database to 10 requests per hour per user.
What is the most scalable, idiomatic (used for its intended purpose), decoupled. reusable, and maintainable integration mechanism available to synchronize each inventory pricing change with the various customer experience systems in near real-time?

A. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the watermark attribute set to an appropriate database column In the same now, use a Scatter-Gather to call each customer experience system's REST API with transformed inventory-pricing records
B. Add a trigger to the inventory-pricing database table so that for each change to the inventory pricing database, a stored procedure is called that makes a REST call to a Mule application Write the Mule application to publish each Mule event as a message to an Anypoint MQ exchange Write other Mule applications to subscribe to the Anypoint MQ exchange, transform each received message, and then update the Mule application's corresponding customer experience system(s)
C. Replace the external web UI application with a Mule application to accept HTTP requests fromclient web browsersIn the same Mule application, use a Batch Job scope to test if the database request will succeed,aggregate pricing changes within a short time window, and then update both the inventory pricingdatabase and each customer experience system using a Parallel For Each scope
D. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the ID attribute set to an appropriate database column In the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queue Write other Mule applications to subscribe to the Anypoint MQ queue, transform each received message, and then update the Mule application's corresponding customer experience system(s)

Correct Answer: B
Section:

QUESTION 47
An ABC Farms project team is planning to build a new API that is required to work with data from different domains across the organization.
The organization has a policy that all project teams should leverage existing investments by reusing existing APIs and related resources and documentation that other project teams have already developed and deployed.
To support reuse, where on Anypoint Platform should the project team go to discover and read existing APIs, discover related resources and documentation, and interact with mocked versions of those APIs?

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

Correct Answer: D
Section:
Explanation:
The mocking service is a feature of Anypoint Platform and runs continuously. You can run the mocking service from the text editor, the visual editor, and from Anypoint Exchange. You can simulate calls to the API in API Designer before publishing the API specification to Exchange or in Exchange after publishing the API specification.
Reference: https://docs.mulesoft.com/design-center/design-mocking-service

QUESTION 48
A Mule application is being designed for deployment to a single CloudHub worker. The Mule application will have a flow that connects to a SaaS system to perform some operations each time the flow is invoked.
The SaaS system connector has operations that can be configured to request a short-lived token (fifteen minutes) that can be reused for subsequent connections within the fifteen minute time window. After the token expires, a new token must be requested and stored.

What is the most performant and idiomatic (used for its intended purpose) Anypoint Platform component or service to use to support persisting and reusing tokens in the Mule application to help speed up reconnecting the Mule application to the SaaS application?

A. Nonpersistent object store
B. Persistent object store
C. Variable
D. Database

Correct Answer: D
Section:
Explanation:
Reference: https://docs.mulesoft.com/mule-runtime/4.4/reconnection-strategy-about

QUESTION 49
An organization has an HTTPS-enabled Mule application named Orders API that receives requestsfrom another Mule application named Process Orders.
The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS).
At a minimum, what must be stored in each truststore and keystore of these two Mule applications to properly support two-way TLS between the two Mule applications while properly protecting each Mule application's keys?

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

Correct Answer: C
Section:
Explanation:
Reference: https://www.caeliusconsulting.com/blogs/one-way-and-two-way-tls-and-theirimplementation-in-mulesoft/

QUESTION 50
A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction.
What is the behavior of the Scatter-Gather when running within this local transaction?

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

Correct Answer: A
Section:

Explanation:
Reference: https://docs.mulesoft.com/mule-runtime/4.4/transaction-management

QUESTION 51
An organization is creating a Mule application that will be deployed to CloudHub. The Mule application has a property named dbPassword that stores a database user's password.
The organization's security standards indicate that the dbPassword property must be hidden from every Anypoint Platform user after the value is set in the Runtime Manager Properties tab.
What configuration in the Mule application helps hide the dbPassword property value in Runtime Manager?

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

Correct Answer: B
Section:
Explanation:
Reference: https://docs.mulesoft.com/runtime-manager/secure-application-properties

QUESTION 52
An organization is designing a Mule application to periodically poll an SFTP location for new files containing sales order records and then process those sales orders. Each sales order must be processed exactly once.
To support this requirement, the Mule application must identify and filter duplicate sales orders on the basis of a unique ID contained in each sales order record and then only send the new sales orders to the downstream system.
What is the most idiomatic (used for its intended purpose) Anypoint connector, validator, or scope that can be configured in the Mule application to filter duplicate sales orders on the basis of the unique ID field contained in each sales order record?

A. Configure a Cache scope to filter and store each record from the received file by the order ID
B. Configure a Database connector to filter and store each record by the order ID
C. Configure an Idempotent Message Validator component to filter each record by the order ID
D. Configure a watermark In an On New or Updated File event source to filter unique records by the order ID

Correct Answer: C
Section:
Explanation:
Reference: https://docs.mulesoft.com/mule-runtime/3.9/idempotent-filter

QUESTION 53
An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.
What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

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

Correct Answer: B
Section:
Explanation:
Reference: https://blogs.mulesoft.com/dev-guides/api-security-ways-to-authenticate-and-authorize/

QUESTION 54
An organization has defined a common object model in Java to mediate the communication between different Mule applications in a consistent way. A Mule application is being built to use this common object model to process responses from a SOAP API and a REST API and then write the processed results to an order management system.
The developers want Anypoint Studio to utilize these common objects to assist in creating mappings for various transformation steps in the Mule application.
What is the most idiomatic (used for its intended purpose) and performant way to utilize these common objects to map between the inbound and outbound systems in the Mule application?

A. Use JAXB (XML) and Jackson (JSON) data bindings
B. Use the WSS module
C. Use the Java module
D. Use the Transform Message component

Correct Answer: A
Section:
Explanation:
Reference: https://docs.mulesoft.com/mule-runtime/3.9/understanding-mule-configuration

QUESTION 55
A marketing organization is designing a Mule application to process campaign data. The Mule application will periodically check for a file in a SFTP location and process the records in the file. The size of the file can vary from 10MB to 5GB. Due to the limited availabiltty of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.
The application must transform and send different formats of this file to three different downstream SFTP locations.
What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

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

Correct Answer: A
Section:
Explanation:
Reference: https://docs.mulesoft.com/mule-runtime/4.4/streaming-about

QUESTION 56
A manufacturing company is planning to deploy Mule applications to its own Azure Kubernetes Service infrastructure.
The organization wants to make the Mule applications more available and robust by deploying each Mule application to an isolated Mule runtime in a Docker container while managing all the Mule applications from the MuleSoft-hosted control plane.
What is the most idiomatic (used for its intended purpose) choice of runtime plane to meet these organizational requirements?

A. Anypoint Platform Private Cloud Edition
B. Anypoint Runtime Fabric
C. CloudHub
D. Anypoint Service Mesh

Correct Answer: B
Section:
Explanation:
Reference: https://blogs.mulesoft.com/dev-guides/how-to-tutorials/anypoint-runtime-fabric/

QUESTION 57
The AnyAirline organization's passenger reservations center is designing an integration solution that combines invocations of three different System APIs (bookFlight, bookHotel, and bookCar) in a business transaction. Each System API makes calls to a single database.
The entire business transaction must be rolled back when at least one of the APIs fails.
What is the most idiomatic (used for its intended purpose) way to integrate these APIs in near realtime that provides the best balance of consistency, performance, and reliability?

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

Correct Answer: C
Section:
Explanation:
Reference: https://aws.amazon.com/blogs/compute/building-a-serverless-distributed-applicationusing-a-saga-orchestration-pattern/

QUESTION 58
As a part of project requirement, client will send a stream of data to mule application. Payload size can vary between 10mb to 5GB. Mule application is required to transform the data and send across multiple sftp servers. Due to the cost cuttings in the organization, mule application can only be allocated one worker with size of 0.2 vCore.
As an integration architect , which streaming strategy you would suggest to handle this scenario?

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

Correct Answer: D
Section:
Explanation:
As the question says that data needs to be sent across multiple sftp serves , we cannot use nonrepeatable streams. The non-repeatable strategy disables repeatable streams, which enables you to read an input stream only once.
You cant use in memory storage because with 0.2 vcore you will get only 1 GB of heap memory.
Hence application will error out for file more than 1 GB.
Hence the correct option is file base repeatable stream

QUESTION 59
Mule application muleA deployed in cloudhub uses Object Store v2 to share data across instances. As a part of new requirement , application muleB which is deployed in same region wants to access this Object Store.
Which of the following option you would suggest which will have minimum latency in this scenario?

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

Correct Answer: A
Section:
Explanation:
V2 Rest API is recommended for on premise applications to access Object Store. It also comes with overhead of encryption and security of using rest api. With Object Store v2, the API call is localized to the same data center as the Runtime Manager app.
But in this case requirement is to access the OS of other mule application and not the same mule application.
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.
However, Object Store v2 is not designed for app-to-app communication.

QUESTION 60
As a part of project , existing java implementation is being migrated to Mulesoft. Business is very tight on the budget and wish to complete the project in most economical way possible.
Canonical object model using java is already a part of existing implementation. Same object model is required by mule application for a business use case. What is the best way to achieve this?

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

Correct Answer: A
Section:
Explanation:
Mule 4 is built to:
• Minimize the need for custom code.
• Avoid the need for you to know or understand Java.
However, some advanced uses cases require integration with custom Java code, such as:
• Reuse of a library, such as a tax calculation library.
• Reuse of a canonical object model that is standard in the organization.
• Execution of custom logic using Java.
Mule ref doc : https://docs.mulesoft.com/java-module/1.2/

QUESTION 61
In one of the critical payment related mule application, transaction is being used . As an enhancement to implementation , scatter gather route is introduced which is also the part of transaction group. Scatter gather route has 4 routes.
What will be the behavior of the Mule application in case of error occurs in 4th route of the scattergather
router and transaction needs to be rolled back?

A. Only errored route will be rolled back
B. All routes will be rolled back
C. Scatter Gather router cannot be part of transaction

Correct Answer: B
Section:
Explanation:
• Scatter Gather: When running within a transaction, Scatter Gather does not execute in parallel. This means that the second route is executed after the first one is processed, the third after the second one, etc. In case of error, all routes will be rolled back

QUESTION 62
Which Mulesoft feature helps users to delegate their access without sharing sensitive credentials or giving full control of accounts to 3rd parties?

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

Correct Answer: C
Section:
Explanation:
Connected Apps
The Connected Apps feature provides a framework that enables an external application to integrate with Anypoint Platform using APIs through OAuth 2.0 and OpenID Connect. Connected apps help users delegate their access without sharing sensitive credentials or giving full control of their accounts to third parties. Actions taken by connected apps are audited, and users can also revoke access at any time. Note that some products do not currently include client IDs in this release of the Connected Apps feature. The Connected Apps feature enables you to use secure authentication protocols and control an app's access to user data. Additionally, end users can authorize the app to access their Anypoint Platform data.
Mule Ref Doc : https://docs.mulesoft.com/access-management/connected-apps-overview

QUESTION 63
What is maximum vCores can be allocated to application deployed to CloudHub?

A. 1 vCores
B. 2 vCores
C. 4 vCores
D. 16 vCores

Correct Answer: D
Section:

QUESTION 64
An organization has just developed a Mule application that implements a REST API. The mule application will be deployed to a cluster of customer hosted Mule runtimes.
What additional infrastructure component must the customer provide in order to distribute inbound API requests across the Mule runtimes of the cluster?

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

Correct Answer: B
Section:
Explanation:
Correct answer is An HTTP Load Balancer.
Key thing to note here is that we are deploying application to customer hosted Mule runtime. This means we will need load balancer to route the requests to different instances of the cluster.
Rest all options are distractors and their requirement depends on project use case.

QUESTION 65
An insurance company is implementing a MuleSoft API to get inventory details from the two vendors. Due to network issues, the invocations to vendor applications are getting timed-out intermittently. But the transactions are successful upon reprocessing What is the most performant way of implementing this requirement?

A. Implement a scatter-gather scope to invoke the two vendor applications on two different route Use the Until-Successful scope to implement the retry mechanism for timeout errors on each route
B. Implement a Choice scope to invoke the two vendor applications on two different route Use the try-catch scope to implement the retry mechanism for timeout errors on each route

C. Implement a For-Each scope to invoke the two vendor applications
Use until successful scope to implement the retry mechanism for the timeout errors
D. Implement Round-Robin scope to invoke the two vendor applications on two different routes Use the Try-Catch scope to implement retry mechanism for timeout errors on each route

Correct Answer: A
Section:

QUESTION 66
An airline is architecting an API connectivity project to integrate its flight data into an online aggregation website. The interface must allow for secure communication high-performance and asynchronous message exchange.
What are suitable interface technologies for this integration assuming that Mulesoft fully supports these technologies and that Anypoint connectors exist for these interfaces?

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

Correct Answer: A
Section:

QUESTION 67
An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger of flow for processing a high-volume set of records from the source system and synchronize with the SaaS system using the Batch job scope After processing 1000 records in a periodic synchronization of 1 lakh records, the replicas in which batch job instance was started went down due to unexpected failure in the runtime fabric environment What is the consequence of losing the replicas that run the Batch job instance?

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

Correct Answer: B
Section:

QUESTION 68
One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make.
Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments.
The production environment however does not have any rate limiting.
What is the cost-effective approach to conduct performance test of the API implementation in the non-production staging environment?

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

Correct Answer: C
Section:

QUESTION 69
A system API EmployeeSAPI is used to fetch employee's data from an underlying SQL database.
The architect must design a caching strategy to query the database only when there is an update to the employees stable or else return a cached response in order to minimize the number of redundant transactions being handled by the database.
What must the architect do to achieve the caching objective?

A. Use an On Table Row on employees table and call invalidate cache
Use an object store caching strategy and expiration interval to empty
B. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and expiration interval to empty
C. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and set expiration interval to 1-hour
D. Use an on table rule on employees table call invalidate cache and said new employees data to cache Use an object store caching strategy and set expiration interval to 1-hour

Correct Answer: A
Section:

QUESTION 70
A leading bank implementing new mule API.
The purpose of API to fetch the customer account balances from the backend application and display them on the online platform the online banking platform. The online banking platform will send an array of accounts to Mule API get the account balances.
As a part of the processing the Mule API needs to insert the data into the database for auditing purposes and this process should not have any performance related implications on the account balance retrieval flow How should this requirement be implemented to achieve better throughput?

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

Correct Answer: D
Section:

QUESTION 71
A Mule application is built to support a local transaction for a series of operations on a single database. The mule application has a Scatter-Gather scope that participates in the local transaction.
What is the behavior of the Scatter-Gather when running within this local transaction?

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

Correct Answer: A
Section:

QUESTION 72
How does timeout attribute help inform design decisions while using JMS connector listening for incoming messages in an extended architecture (XA) transaction?

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

Correct Answer: D
Section:

QUESTION 73
An auto mobile company want to share inventory updates with dealers Dl and D2 asynchronously and concurrently via queues Q1 and Q2. Dealer Dl must consume the message from the queue Q1 and dealer D2 to must consume a message from the queue Q2.
Dealer D1 has implemented a retry mechanism to reprocess the transaction in case of any errors while processing the inventers updates. Dealer D2 has not implemented any retry mechanism.
How should the dealers acknowledge the message to avoid message loss and minimize impact on the current implementation?

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

Correct Answer: D
Section:

QUESTION 74
A company is using Mulesoft to develop API's and deploy them to Cloudhub and on premises targets.
Recently it has decided to enable Runtime Fabric deployment option as well and infrastructure is set up for this option.
What can be used to deploy Runtime Fabric?

A. AnypointCLI
B. Anypoint platform REST API's
C. Directly uploading ajar file from the Runtime manager
D. Mule maven plug-in

Correct Answer: D
Section:

QUESTION 75
As an enterprise architect, what are the two reasons for which you would use a canonical data model in the new integration project using Mulesoft Anypoint platform (choose two answers)

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

Correct Answer: A, B
Section:

QUESTION 76
A company is planning to migrate its deployment environment from on-premises cluster to a Runtime Fabric (RTF) cluster. It also has a requirement to enable Mule applications deployed to a Mule runtime instance to store and share data across application replicas and restarts.
How can these requirements be met?

A. Anypoint object store V2 to share data between replicas in the RTF cluster
B. Install the object store pod on one of the cluster nodes
C. Configure Persistence Gateway in any of the servers using Mule Object Store
D. Configure Persistent Gateway at the RTF

Correct Answer: D
Section:

QUESTION 77
An organization designing a hybrid, load balanced, single cluster production environment. Due to performance service level agreement goals, it is looking into running the Mule applications in an active-active multi node cluster configuration.
What should be considered when running its Mule applications in this type of environment?

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

Correct Answer: B
Section:

QUESTION 78
An organization has decided on a cloud migration strategy to minimize the organization's own IT resources. Currently the organization has all of its new applications running on its own premises and uses an on-premises load balancer that exposes all APIs under the base URL (https://api.rutujar.com).
As part of migration strategy, the organization is planning to migrate all of its new applications and load balancer CloudHub.
What is the most straightforward and cost-effective approach to Mule application deployment and load balancing that preserves the public URL's?

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

Correct Answer: B
Section:

QUESTION 79

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

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

Correct Answer: B
Section:

QUESTION 80
Refer to the exhibit.

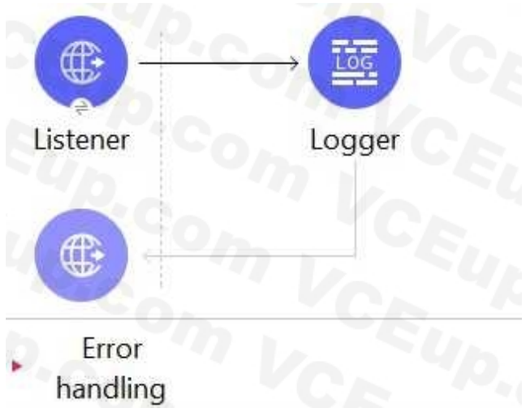

The HTTP Listener and the Logger are being handled from which thread pools respectively?

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

Correct Answer: C
Section:

QUESTION 81
An organization has strict unit test requirement that mandate every mule application must have an MUnit test suit with a test case defined for each flow and a minimum test coverage of 80%.
A developer is building Munit test suit for a newly developed mule application that sends API request to an external rest API.
What is the effective approach for successfully executing the Munit tests of this new application while still achieving the required test coverage for the Munit tests?

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

Correct Answer: C
Section:

QUESTION 82
A company is planning to extend its Mule APIs to the Europe region. Currently all new applications are deployed to Cloudhub in the US region following this naming convention
{API name}-{environment}. for example, Orders-SAPI-dev, Orders-SAPI-prod etc.
Considering there is no network restriction to block communications between API's, what strategy should be implemented in order to apply the same new API's running in the EU region of CloudHub as well to minimize latency between API's and target users and systems in Europe?

A. Set region property to Europe (eu-de) in API manager for all the mule application No need to change the naming convention
B. Set region property to Europe (eu-de) in API manager for all the mule application Change the naming convention to {API name}-{environment}-{region} and communicate this change to the consuming applications and users
C. Set region property to Europe (eu-de) in runtime manager for all the mule application No need to change the naming convention
D. Set region property to Europe (eu-de) in runtime manager for all the mule application Change the naming convention to {API name}-{environment}-{region} and communicate this change to the consuming applications and users

Correct Answer: D
Section:

QUESTION 83
As a part of project requirement, Java Invoke static connector in a mule 4 application needs to invoke a static method in a dependency jar file. What are two ways to add the dependency to be visible by the connectors class loader?
(Choose two answers)

A. In the Java Invoke static connector configuration, configure a path and name of the dependency jar file
B. Add the dependency jar file to the java classpath by setting the JVM parameters
C. Use Maven command to include the dependency jar file when packaging the application
D. Configure the dependency as a shared library in the project POM
E. Update mule-artefact.json to export the Java package

Correct Answer: B, D
Section:

QUESTION 84
When the mule application using VM is deployed to a customer-hosted cluster or multiple cloudhub workers, how are messages consumed by the Mule engine?

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

Correct Answer: C
Section:

QUESTION 85
An insurance provider is implementing Anypoint platform to manage its application infrastructure and is using the customer hosted runtime for its business due to certain financial requirements it must meet. It has built a number of synchronous API's and is currently hosting these on a mule runtime on one server These applications make use of a number of components including heavy use of object stores and VM queues.
Business has grown rapidly in the last year and the insurance provider is starting to receive reports of reliability issues from its applications.
The DevOps team indicates that the API's are currently handling too many requests and this is over loading the server. The team has also mentioned that there is a significant downtime when the server is down for maintenance.
As an integration architect, which option would you suggest to mitigate these issues?

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

Correct Answer: B
Section:

QUESTION 86
A company is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTP5 POST and must be acknowledged immediately.
Once acknowledged the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to the rejections from the back-end system will need to be processed manually (outside the banking system).
The mule application will be deployed to a customer hosted runtime and will be able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization's firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.
Which combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

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

Correct Answer: B
Section:

QUESTION 87
An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events What are the two considerations that must be kept in mind while using test recorder
(Choose two answers)

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

Correct Answer: A, D
Section:

QUESTION 88
A company is modernizing its legal systems lo accelerate access lo applications and data while supporting the adoption of new technologies. The key to achieving this business goal is unlocking the companies' key systems and dala including microservices miming under Docker and kubernetes containers using apis.
Considering the current aggressive backlog and project delivery requirements the company wants to take a strategic approach in the first phase of its transformation projects by quickly deploying API's in mule runtime that are able lo scale, connect to on premises systems and migrate as needed.
Which runtime deployment option supports company's goals?

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

Correct Answer: C
Section:

QUESTION 89
A corporation has deployed multiple mule applications implementing various public and private API's to different cloudhub workers. These API's arc Critical applications that must be highly available and in line with the reliability SLA as defined by stakeholders.
How can API availability (liveliness or readiness) be monitored so that Ops team receives outage notifications?

A. Enable monitoring of individual applications from Anypoint monitoring
B. Configure alerts with failure conditions in runtime manager
C. Configure alerts failure conditions in API manager
D. Use any point functional monitoring test API's functional behavior

Correct Answer: A
Section:

QUESTION 90
An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/'CD pipeline. Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?

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

Correct Answer: D
Section:

QUESTION 91
A mule application must periodically process a large dataset which varies from 6 GB lo 8 GB from a back-end database and write transform data lo an FTPS server using a properly configured bad job scope.
The performance requirements of an application are approved to run in the cloud hub 0.2 vCore with 8 GB storage capacity and currency requirements are met.
How can the high rate of records be effectively managed in this application?

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

Correct Answer: A
Section:

QUESTION 92

An organization has deployed runtime fabric on an eight note cluster with performance profile. An API uses and non persistent object store for maintaining some of its state data. What will be the impact to the stale data if server crashes?

A. State data is preserved
B. State data is rolled back to a previously saved version
C. State data is lost
D. State data is preserved as long as more than one more is unaffected by the crash

Correct Answer: D
Section:

QUESTION 93
A rale limiting policy has been applied to a soap VI.2 API published in Clondhub. The API implementation catches errors in a global error handler on error propagate in the main flow for HTTP: RETRY_EXHAUSTED with HTTP status set to 429 and any with the HTTP status set to 500.
What is the expected H1TP status when the client exceeds the quota of the API calls?

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

Correct Answer: A
Section:

QUESTION 94
An organization is designing multiple new applications to run on CloudHub in a single Anypoint VPC and that must share data using a common persistent Anypoint object store V2 (OSv2).
Which design gives these mule applications access to the same object store instance?

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

Correct Answer: D
Section:

QUESTION 95
A leading e-commerce giant will use Mulesoft API's on runtime fabric (RTF) to process customer orders. Some customer's sensitive information such as credit card information is also there as a part of a API payload.
What approach minimizes the risk of matching sensitive data to the original and can convert back to the original value whenever and wherever required?

A. Apply masking to hide the sensitive information and then use API
B. manager to detokenize the masking format to return the original value
C. create a tokenization format and apply a tokenization policy to the API Gateway
D. Used both masking and tokenization
E. Apply a field level encryption policy in the API Gateway

Correct Answer: A

Section:

QUESTION 96
Customer has deployed mule applications to different customer hosted mule run times. Mule applications are managed from Anypoint platform.
What needs to be configured to monitor these Mule applications from Anypoint monitoring and what sends monitoring data to Anypoint monitoring?

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

Correct Answer: C
Section:

QUESTION 97
A trading company handles millions of requests a day. Due to nature of its business, it requires excellent performance and reliability within its application.
For this purpose, company uses a number of event-based API's hosted on various mule clusters that communicate across a shared message queue sitting within its network.
Which method should be used to meet the company's requirement for its system?

A. XA transactions and XA connected components
B. JMS transactions
C. JMS manual acknowledgements with a reliability pattern
D. VM queues with reliability pattern

Correct Answer: C
Section:

QUESTION 98
An organization has chosen Mulesoft for their integration and API platform.
According to the Mulesoft catalyst framework, what would an integration architect do to create achievement goals as part of their business outcomes?

A. Measure the impact of the centre for enablement
B. build and publish foundational assets
C. agree upon KPI's and help develop and overall success plan
D. evangelize API's

Correct Answer: C
Section:

QUESTION 99
An organization has implemented the cluster with two customer hosted Mule runtimes is hosting an application.
This application has a flow with a JMS listener configured to consume messages from a queue destination. As an integration architect can you advise which JMS listener configuration must be used to receive messages in all the nodes of the cluster?

A. Use the parameter primaryNodeOnly= "false" on the JMS listener

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

Correct Answer: A
Section:

QUESTION 100
A customer wants to use the mapped diagnostic context (MDC) and logging variables to enrich its logging and improve tracking by providing more context in the logs.
The customer also wants to improve the throughput and lower the latency of message processing.
As an Mulesoft integration architect can you advise, what should the customer implement to meet these requirements?

A. Use synchronous logging and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
B. Useasync logger at the level greater than INFO and use pattern layout with [%MDC] in the log4j2,xml configuration file and then configure the logging variables
C. Useasync logger at the level equal to DEBUG orTRACEand use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
D. Use synchronous logging at the INFO DEBUG or Trace level and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables

Correct Answer: B
Section:

QUESTION 101
A project team uses RAML specifications to document API functional requirements and deliver API definitions. As per the current legal requirement, all designed API definitions to be augmented with an additional non-functional requirement to protect the services from a high rate of requests according to define service level agreements.
Assuming that the project is following Mulesoft API governance and policies, how should the project team convey the necessary non-functional requirement to stakeholders?

A. Create proxies in API manager for the non functional requirement and publish to exchange
B. Add all non functional requirements as comments to RAML specification and publish to exchange
C. Create various SLA's in API manager for the non functional requirement and publish to exchange
D. Update API definitions with the fragment for the appropriate policy and publish to exchange

Correct Answer: D
Section:

QUESTION 102
An organization is using Mulesoft cloudhub and develops API's in the latest version. As a part of requirements for one of the API's, third party API needs to be called. The security team has made it clear that calling any external API needs to have include listing As an integration architect please suggest the best way to accomplish the design plan to support these requirements?

A. Implement includelist IP on the cloudhub VPC firewall to allow the traffic
B. Implement the validation of includelisted IP operation
C. Implement the Any point filter processor to implement the include list IP
D. Implement a proxy for the third party API and enforce the IPinclude list policy and call this proxy from the flow of the API

Correct Answer: D
Section:

QUESTION 103
A project uses Jenkins to implement CI/CD process. It was observed that each Mule package contains some of the Jenkins files and folders for configurations of CI/CD jobs.

As these files and folders are not part of the actual package, expectation is that these should not be part of deployed archive.
Which file can be used to exclude these files and folders from the deployed archive?

A. muleignore
B. _unTrackMule
C. muleInclude
D. _muleExclude

Correct Answer: D
Section:

QUESTION 104
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 105
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 106
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 107
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 108
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.

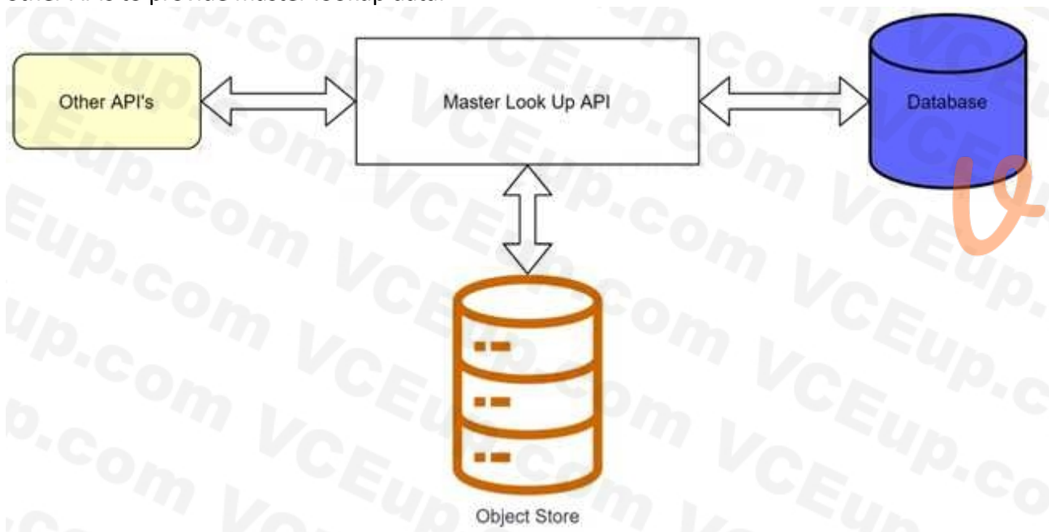

Master lookup API is deployed on two cloudhub workers of 0.1 vCore each because there is a lot of master data to be cached. Master lookup data is stored as a key value pair. The cache gets refreshed if they key is not found in the cache.
Doing performance testing it was observed that the Master lookup 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 109

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 110
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 111
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 112
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 113
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 114
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 115
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 116
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:
Explanation: