

Mulesoft.Premium.MCIA-LEVEL-1-MAINTENANCE.30q - DEMO

Number: MCIA-LEVEL-1-M
Passing Score: 800
Time Limit: 120 min



Exam Code: MCIA-LEVEL-1-MAINTENANCE

Exam Name: MuleSoft Certified Integration Architect - Level 1 MAINTENANCE

Website: <https://VCEup.com/>

Team-Support: <https://VCEplus.io/>



QUESTION 1

An organization is designing a mule application to support an all or nothing transaction between several 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: (none)

Explanation

Explanation/Reference:

Explanation:

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

QUESTION 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 the external API?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

3xx HTTP status codes indicate a redirection that the user agent (a web browser or a crawler) needs to 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), for 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: (none)

Explanation

Explanation/Reference:

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 deployed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates.

What type of restrictions exist on the types of certificates that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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 reliability 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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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 can 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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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:

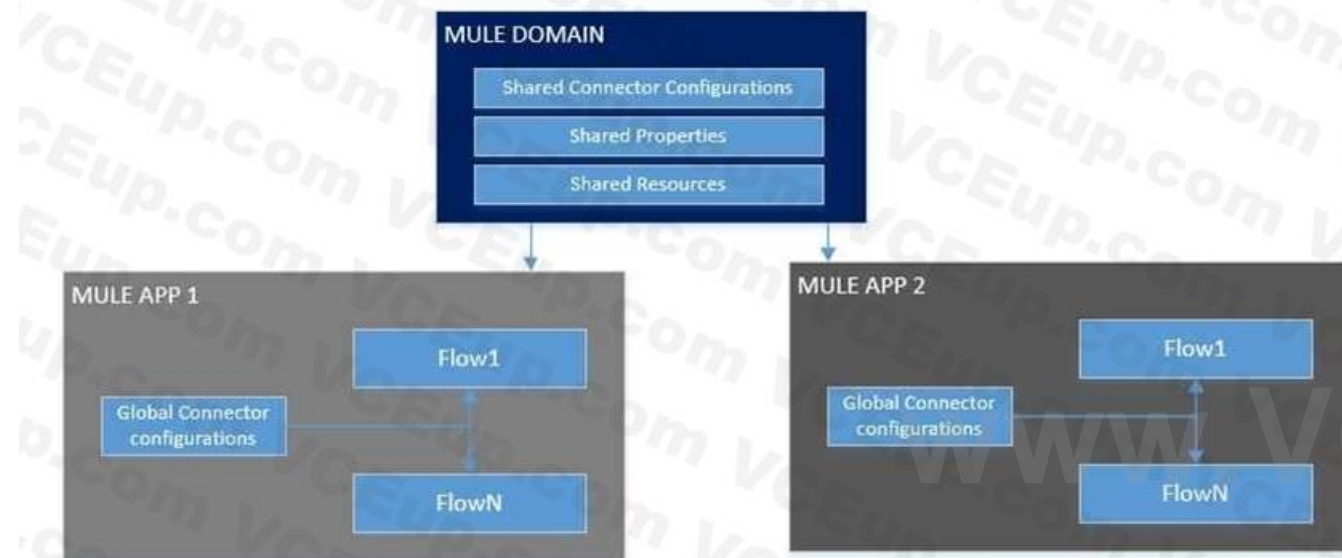
- Expose multiple services within the domain through the same port.
- Share the connection to persistent storage.
- Share services between apps through a well-defined interface.
- Ensure consistency between apps upon any changes because the configuration is only set in one place.

* 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 be 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?

- Redirect client requests through an HTTP 303 temporary redirect status code to the fallback API whenever the Order API is unavailable
- Set an option in the HTTP Requester component that invokes the order API to instead invoke a fallback API whenever an HTTP 4XX or 5XX response status code is received from Order API
- 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
- Search Anypoint Exchange for a suitable existing fallback API and then implement invocations to their fallback API in addition to the Order API

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

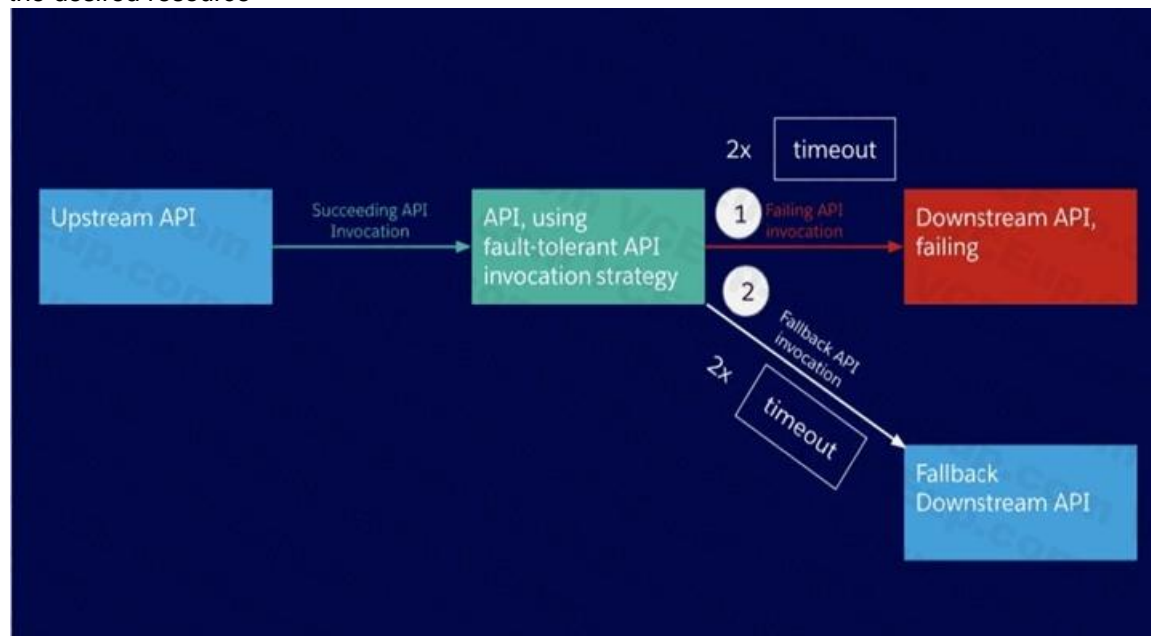
* Resilience testing is a type of software testing that observes how applications act under stress. It's meant to ensure the product's ability to perform in chaotic conditions without a loss of core functions or data; it ensures a quick recovery after unforeseen, uncontrollable events.

* 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. These codes indicate to the user agent (i.e. your web browser) that an additional action is required in order to complete the request and access the desired resource



Hence correct answer is Redirect client requests through an HTTP 303 temporary redirect status code to 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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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 still maintaining HTTP semantics, such as verbs, methods, and headers. It came into use in 2015, and offers several methods to decrease latency, especially when dealing with mobile platforms and server-intensive graphics and videos* Currently, Mule application can have API policies only for Mule application that accepts requests over 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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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)

www.VCEup.com

**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: (none)

Explanation


Explanation/Reference:

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.

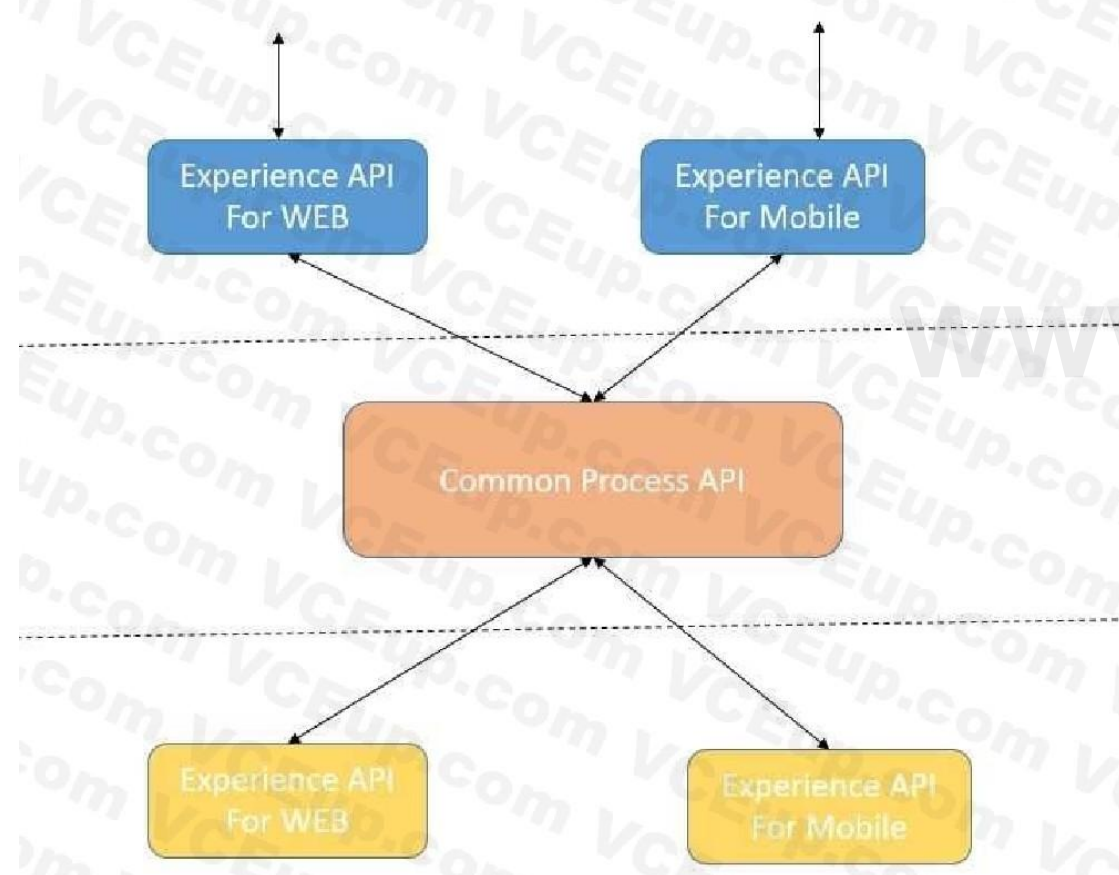
Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system



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 through 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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation


Explanation/Reference:

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 deferred, 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
- D. Choose a region that is the traffic-weighted geographic center of all web clients

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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

QUESTION 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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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 downstream 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: (none)

Explanation

Explanation/Reference:

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: (none)

Explanation

Explanation/Reference:

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 CI/CD pipeline which can be 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.

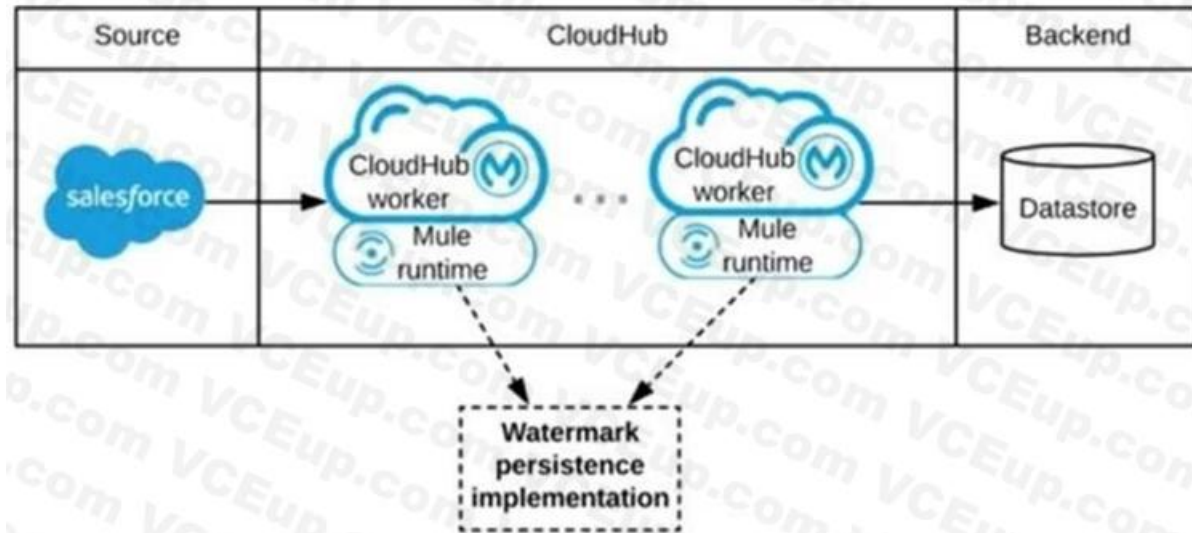
Mule uses object stores to persist data for eventual retrieval.

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 can't 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: (none)

Explanation

Explanation/Reference:

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