MuleSoft

Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1

NEW QUESTION 1
An Integration Mule application is being designed to synchronize customer data between two systems. One system is an IBM Mainframe and the other system is a Salesforce Marketing Cloud (CRM) instance. Both systems have been deployed in their typical configurations, and are to be invoked using the native protocols provided by Salesforce and IBM.
What interface technologies are the most straightforward and appropriate to use in this Mute application to interact with these systems, assuming that Anypoint Connectors exist that implement these interface technologies?

A. IBM: DB access CRM: gRPC
B. IBM: REST CRM:REST
C. IBM: Active MQ CRM: REST
D. IBM: CICS CRM: SOAP

Answer: D

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

NEW QUESTION 2
An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to custom hosted Mule runtimes (on-premises) while others execute in the MuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems.
How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access the backend systems?

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

Answer: D

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

NEW QUESTION 3
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.

Answer: C

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

NEW QUESTION 4
What operation can be performed through a JMX agent enabled in a Mule application?

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

Answer: C

Explanation:
JMX Management Java Management Extensions (JMX) is a simple and standard way to manage applications, devices, services, and other resources. JMX is dynamic, so you can use it to monitor and manage resources as they are created, installed, and implemented. You can also use JMX to monitor and manage the Java Virtual Machine (JVM). Each resource is instrumented by one or more Managed Beans, or MBeans. All MBeans are registered in an MBean Server. The JMX server agent consists of an MBean Server and a set of services for handling Mbeans. There are several agents provided with Mule for JMX support. The easiest way to configure JMX is to use the default JMX support agent. Log4J Agent The log4j agent exposes the configuration of the Log4J instance used by Mule for JMX management. You enable the Log4J agent using the <jmx-log4j> element. It does not take any additional properties MuleSoft Reference: https://docs.mulesoft.com/mule-runtime/3.9/jmx-management

NEW QUESTION 5
An Organization has previously provisioned its own AWS VPC hosting various servers. The organization now needs to use Cloudhub to host a Mule application that will implement a REST API once deployed to Cloudhub, this Mule application must be able to communicate securely with the customer-provisioned AWS VPC resources within the same region, without being interceptable on the public internet.
What Anypoint Platform features should be used to meet these network communication requirements between Cloudhub and the existing customer-provisioned AWS VPC?

A. Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC
B. Configure an external identity provider (IDP) in Anypoint Platform with certificates from the customer provisioned AWS VPC
C. Add a default API Whitelisting policy to API Manager to automatically whitelist the customer provisioned AWS VPC IP ranges needed by the Mule applicaton
D. Use VM queues in the Mule application to allow any non-mule assets within the customer provisioned AWS VPC to subscribed to and receive messages

Answer: A

Explanation:
Correct answer is: Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC
* Connecting to your Anypoint VPC extends your corporate network and allows CloudHub workers to access resources behind your corporate firewall.
* You can connect on-premises data centers through a secured VPN tunnel, or a private AWS VPC through VPC peering, or by using AWS Direct Connect.
MuleSoft Doc Reference : https://docs.mulesoft.com/runtime-manager/virtual-private-cloud

NEW QUESTION 6
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

Answer: D

NEW QUESTION 7
An organization is designing the following two Mule applications that must share data via a common persistent object store instance:
- Mule application P will be deployed within their on-premises datacenter.
- Mule application C will run on CloudHub in an Anypoint VPC.
The object store implementation used by CloudHub is the Anypoint Object Store v2 (OSv2).
what type of object store(s) should be used, and what design gives both Mule applications access to the same object store instance?

A. Application P uses the Object Store connector to access a persistent object store Application C accesses this persistent object store via the Object Store REST API through an IPsec tunnel
B. Application C and P both use the Object Store connector to access the Anypoint Object Store v2
C. Application C uses the Object Store connector to access a persistent object Application P accesses the persistent object store via the Object Store REST API
D. Application C and P both use the Object Store connector to access a persistent object store
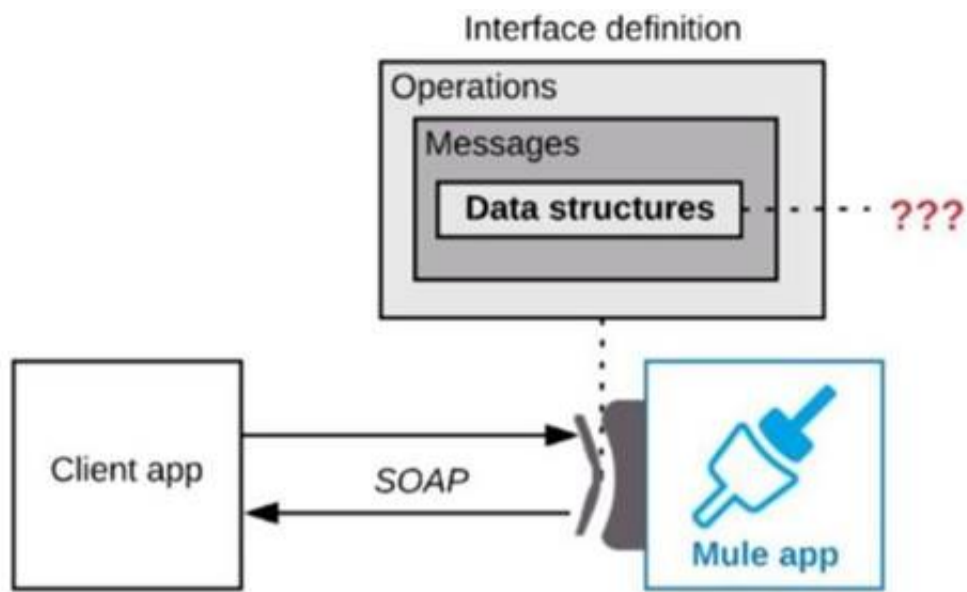
Answer: C

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

NEW QUESTION 8
Refer to the exhibit.
A Mule application is being designed to expose a SOAP web service to its clients.
What language is typically used inside the web service's interface definition to define the data structures that the web service Is expected to exchange with its clients?

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

Answer: B

Explanation:
Correct Answer XSD In this approach to developing a web service, you begin with an XML schema (XSD file) that defines XML data structures to be used as parameters and return types in the web service operations.
-- Reference:
https://www.w3schools.com/xml/schema_intro.asp

NEW QUESTION 9
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

Answer: D

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

NEW QUESTION 10
An organization currently uses a multi-node Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications. The organization is planning to transition to a deployment model based on Docker containers in a Kubernetes cluster. The organization has already created a standard Docker image containing a Mule runtime and all required dependencies (including a JVM), but excluding the Mule application itself.
What is an expected outcome of this transition to container-based Mule application deployments?

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

Answer: A

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

NEW QUESTION 11
A leading bank implementing new mule API.
The purpose of API to fetch the customer account balances from the backend application and display them on the online platform the online banking platform. The

online banking platform will send an array of accounts to Mule API get the account balances.
As a part of the processing the Mule API needs to insert the data into the database for auditing purposes and this process should not have any performance related implications on the account balance retrieval flow
How should this requirement be implemented to achieve better throughput?

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

**Answer:** D

NEW QUESTION 12
A Mule application contains a Batch Job with two Batch Steps (Batch_Step_l and Batch_Step_2). A payload with 1000 records is received by the Batch Job.
How many threads are used by the Batch Job to process records, and how does each Batch Step process records within the Batch Job?

A. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and RECORDS are processed IN PARALLEL within and between the two Batch Steps
B. Each Batch Job uses a SINGLE THREAD for all Batch steps Each Batch step instance receives ONE record at a time as the payload, and RECORDS are processed IN ORDER, first through Batch_Step_l and then through Batch_Step_2
C. Each Batch Job uses a SINGLE THREAD to process a configured block size of record Each Batch Step instance receives A BLOCK OF records as the payload, and BLOCKS of records are processed IN ORDER
D. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and BATCH STEP INSTANCES execute IN PARALLEL to processrecords and Batch Steps in ANY order as fast as possible

Answer: A

Explanation:
* Each Batch Job uses SEVERAL THREADS for the Batch Steps
* Each Batch Step instance receives ONE record at a time as the payload. It's not received in a block, as it does not wait for multiple records to be completed before moving to next batch step. (So Option D is out of choice)
* RECORDS are processed IN PARALLEL within and between the two Batch Steps.
* RECORDS are not processed in order. Let's say if second record completes batch_step_1 before record 1, then it moves to batch_step_2 before record 1. (So option C and D are out of choice)
* A batch job is the scope element in an application in which Mule processes a message payload as a batch of records. The term batch job is inclusive of all three phases of processing: Load and Dispatch, Process, and On Complete.
* A batch job instance is an occurrence in a Mule application whenever a Mule flow executes a batch job. Mule creates the batch job instance in the Load and Dispatch phase. Every batch job instance is identified internally using a unique String known as batch job instance id.
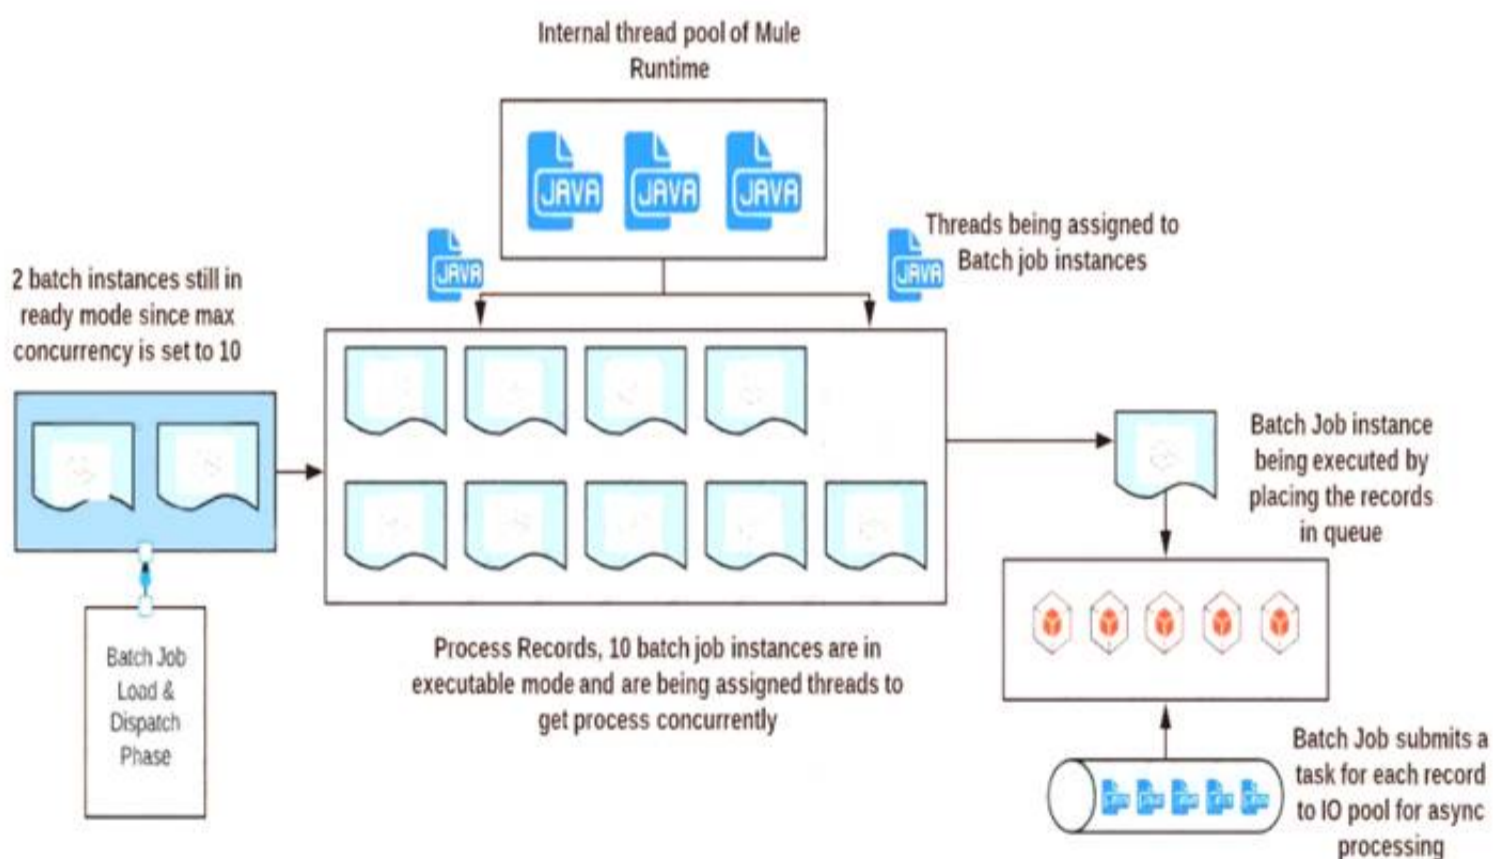
NEW QUESTION 13
A Mule application uses the Database connector.
What condition can the Mule application automatically adjust to or recover from without needing to restart or redeploy the Mule application?

A. One of the stored procedures being called by the Mule application has been renamed
B. The database server was unavailable for four hours due to a major outage but is now fully operational again
C. The credentials for accessing the database have been updated and the previous credentials are no longer valid
D. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade

Answer: B

Explanation:
* Any change in the application will require a restart except when the issue outside the app. For below situations , you would need to redeploy the code after doing necessary changes
-- One of the stored procedures being called by the Mule application has been renamed. In this case, in the Mule application you will have to do changes to

accommodate the new stored procedure name.
-- Required redesign of Mule applications to follow microservice architecture principles. As code is changed, deployment is must
-- If the credentials changed and you need to update the connector or the properties.
-- The credentials for accessing the database have been updated and the previous credentials are no longer valid. In this situation you need to restart or redeploy depending on how credentials are configured in Mule application.
* So Correct answer is The database server was unavailable for four hours due to a major outage but is now fully operational again as this is the only external issue to application.

NEW QUESTION 14

An integration Mute application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mute application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system. Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system).

The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed.

The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages.

What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

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

Answer: D

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

NEW QUESTION 15

49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58.
https://api.acme.com/partnefs/vl. The API implementation is estimated to require deployment to 5 CloudHub workers.
The retailer has obtained a public X.509 certificate for the name apl.acme.com, signed by a reputable CA, to be used as the server certificate.
Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

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

Answer: C

Explanation:
* An X.509 certificate is a vital safeguard against malicious network impersonators. Without x.509 server authentication, man-in-the-middle attacks can be initiated by malicious access points, compromised routers, etc.
* X.509 is most used for SSL/TLS connections to ensure that the client (e.g., a web browser) is not fooled by a malicious impersonator pretending to be a known, trustworthy website.
* Coming to the question , we can not use SLB here as SLB does not allow to define vanity domain names. * Hence we need to use DLB and add certificate in there
----
Hence correct answer is Add the X 509 certificate to the cloudhub Dedicated Load Balancer (DLB), not the Mule application. Create the CNAME for api.acme.com pointing to the DLB's record

NEW QUESTION 16
What aspects of a CI/CD pipeline for Mule applications can be automated using MuleSoft-provided Maven plugins?

A. Compile, package, unit test, validate unit test coverage, deploy
B. Compile, package, unit test, deploy, integration test (Incorrect)
C. Compile, package, unit test, deploy, create associated API instances in API Manager
D. Import from API designer, compile, package, unit test, deploy, publish to Anypoint Exchange

Answer: A

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

Explanation

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

Diagram Description automatically generated
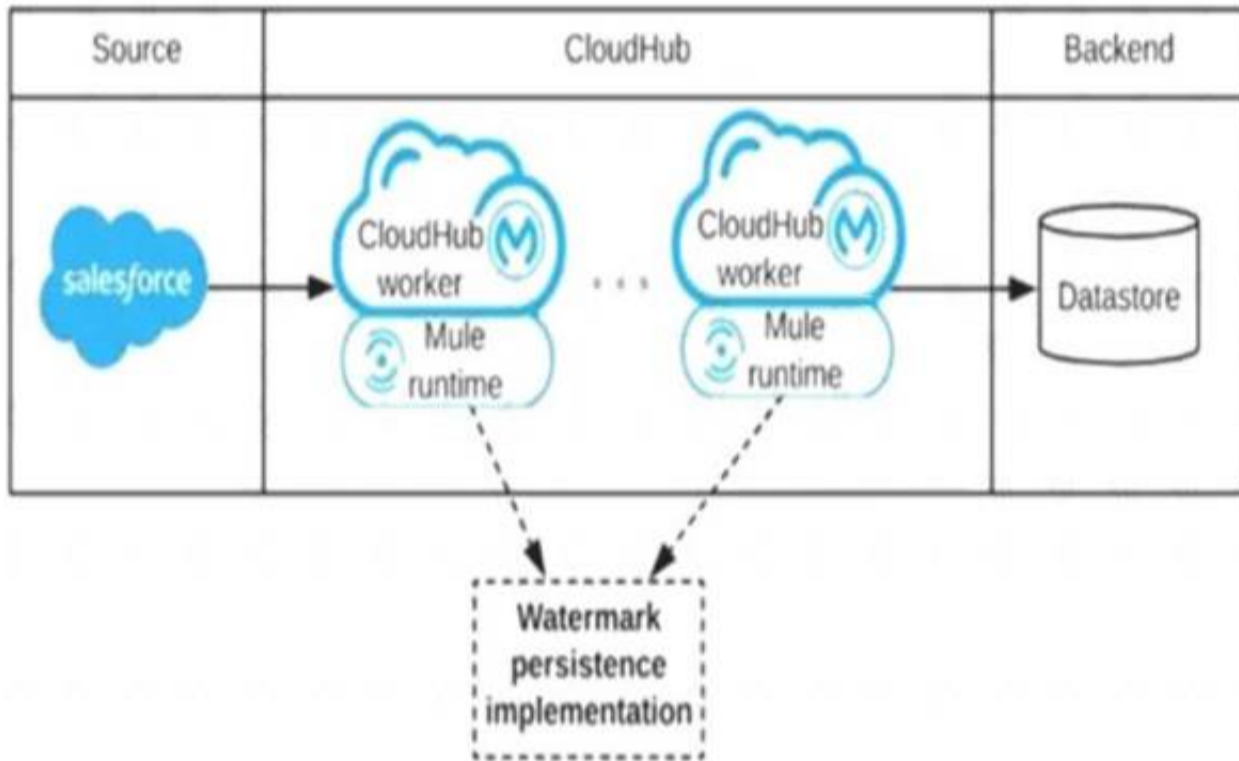
NEW QUESTION 17

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

Answer: C

NEW QUESTION 18

A company is building an application network and has deployed four Mule APIs: one experience API, one process API, and two system APIs. The logs from all the APIs are aggregated in an external log aggregation tool. The company wants to trace messages that are exchanged between multiple API implementations. What is the most idiomatic (based on its intended use) identifier that should be used to implement Mule event tracing across the multiple API implementations?

A. Mule event ID
B. Mule correlation ID
C. Client's IP address
D. DataWeave UUID

Answer: B

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.

NEW QUESTION 19

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

Answer: A

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.

NEW QUESTION 20
What Mule application can have API policies applied by Anypoint Platform to the endpoint exposed by that Mule application?

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

Answer: A

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

NEW QUESTION 21
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

MCIA-Level-1 Practice Exam Features:

* MCIA-Level-1 Questions and Answers Updated Frequently

* MCIA-Level-1 Practice Questions Verified by Expert Senior Certified Staff

* MCIA-Level-1 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* MCIA-Level-1 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The MCIA-Level-1 Practice Test Here