



Salesforce

Exam Questions MuleSoft-Integration-Architect-I

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

An organization plans to use the Anypoint Platform audit logging service to log Anypoint MQ actions. What consideration must be kept in mind when leveraging Anypoint MQ Audit Logs?

- A. Anypoint MQ Audit Logs include logs for sending, receiving, or browsing messages
- B. Anypoint MQ Audit Logs include logs for failed Anypoint MQ operations
- C. Anypoint MQ Audit Logs include logs for queue create, delete, modify, and purge operations

Answer: C

NEW QUESTION 2

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

Answer: A

NEW QUESTION 3

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.

Answer: D

NEW QUESTION 4

Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improve 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 can't 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

Answer: D

NEW QUESTION 5

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

NEW QUESTION 6

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

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

Answer: B

NEW QUESTION 7

A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should be published in a way that is guaranteed in normal situations and also minimizes duplicate delivery in less frequent failure scenarios.

The organizational firewall is configured to only allow outbound traffic on ports 80 and 443. Some external event consumers are within the organizational network, while others are located outside the firewall.

What Anypoint Platform service is most idiomatic (used for its intended purpose) for publishing these Mule events to all external consumers while addressing the desired reliability goals?

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

Answer: B

NEW QUESTION 8

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

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

Answer: B

NEW QUESTION 9

A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application.

What non-functional requirement can be expected to be negatively affected when implementing this change?

- A. Throughput
- B. Consistency
- C. Response time
- D. Availability

Answer: B

NEW QUESTION 10

A developer needs to discover which API specifications have been created within the organization before starting a new project. Which Anypoint Platform component can the developer use to find and try out the currently released API specifications?

- A. Anypoint Exchange
- B. Runtime Manager
- C. API Manager
- D. Object Store

Answer: A

NEW QUESTION 10

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

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

Answer: CE

NEW QUESTION 15

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 flows
- 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

Answer: C

NEW QUESTION 17

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

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

Answer: D

NEW QUESTION 22

When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

- A. Environment setup
- B. Identity management setup
- C. Role and permission setup
- D. Dedicated Load Balancer setup

Answer: B

NEW QUESTION 27

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

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

Answer: C

NEW QUESTION 31

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

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

Answer: AE

NEW QUESTION 33

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 CloudhubCreate a CNAME record for base URL(<https://api.rutujar.com>) in the Cloudhub shared load balancer that points to the A record of the on-premises load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications
- B. Deploy the Mule application to CloudhubUpdate 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 balancerApply mapping rules in DLB to map URLto their corresponding Mule applications
- C. Deploy the Mule application to CloudhubUpdate 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 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 balancerApply mapping rules in DLB to map each API proxy application who is responding new application

Answer: C

NEW QUESTION 36

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

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

Answer: A

NEW QUESTION 37

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

Answer: C

NEW QUESTION 42

In a Mule Application, a flow contains two (2) JMS consume operations that are used to connect to a JMS broker and consume messages from two(2) JMS destination. The Mule application then joins the two JMS messages together.

The JMS broker does not implement high availability (HA) and periodically experiences scheduled outages of upto 10 mins for routine maintenance.

What is the most idiomatic (used for its intended purpose) way to build the mule flow so it can best recover from the expected outages?

- A. Configure a reconnection strategy for the JMS connector
- B. Enclose the two(2) JMS operation in an Until Successful scope
- C. Consider a transaction for the JMS connector
- D. Enclose the two(2) JMS operations in a Try scope with an Error Continue error handler

Answer: A

NEW QUESTION 46

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

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

Answer: A

NEW QUESTION 50

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

Answer: A

NEW QUESTION 51

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

Answer: A

NEW QUESTION 53

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

Answer: A

NEW QUESTION 58

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

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

Answer: B

NEW QUESTION 59

A platform architect includes both an API gateway and a service mesh in the architect of a distributed application for communication management.

Which type of communication management does a service mesh typically perform in this architecture?

- A. Between application services and the firewall
- B. Between the application and external API clients
- C. Between services within the application
- D. Between the application and external API implementations.

Answer: C

NEW QUESTION 60

An external web UI application currently accepts occasional HTTP requests from client web browsers to 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 flow, 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 from client web browsers. In 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 pricing database 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)

Answer: B

NEW QUESTION 65

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

Answer: C

NEW QUESTION 66

A leading eCommerce giant will use MuleSoft APIs on Runtime Fabric (RTF) to process customer orders. Some customer-sensitive information, such as credit card information, is required in request payloads or is included in response payloads in some of the APIs. Other API requests and responses are not authorized to access some of this customer-sensitive information but have been implemented to validate and transform based on the structure and format of this customer-sensitive information (such as account IDs, phone numbers, and postal codes).

What approach configures an API gateway to hide sensitive data exchanged between API consumers and API implementations, but can convert tokenized fields back to their original value for other API requests or responses, without having to recode the API implementations?

Later, the project team requires all API specifications to be augmented with an additional non-functional requirement (NFR) to protect the backend services from a high rate of requests, according to defined service-level agreements (SLAs). The NFR's SLAs are based on a new tiered subscription level "Gold", "Silver", or "Platinum" that must be tied to a new parameter that is being added to the Accounts object in their enterprise data model.

Following MuleSoft's recommended best practices, how should the project team now convey the necessary non-functional requirement to stakeholders?

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

Answer: C

NEW QUESTION 71

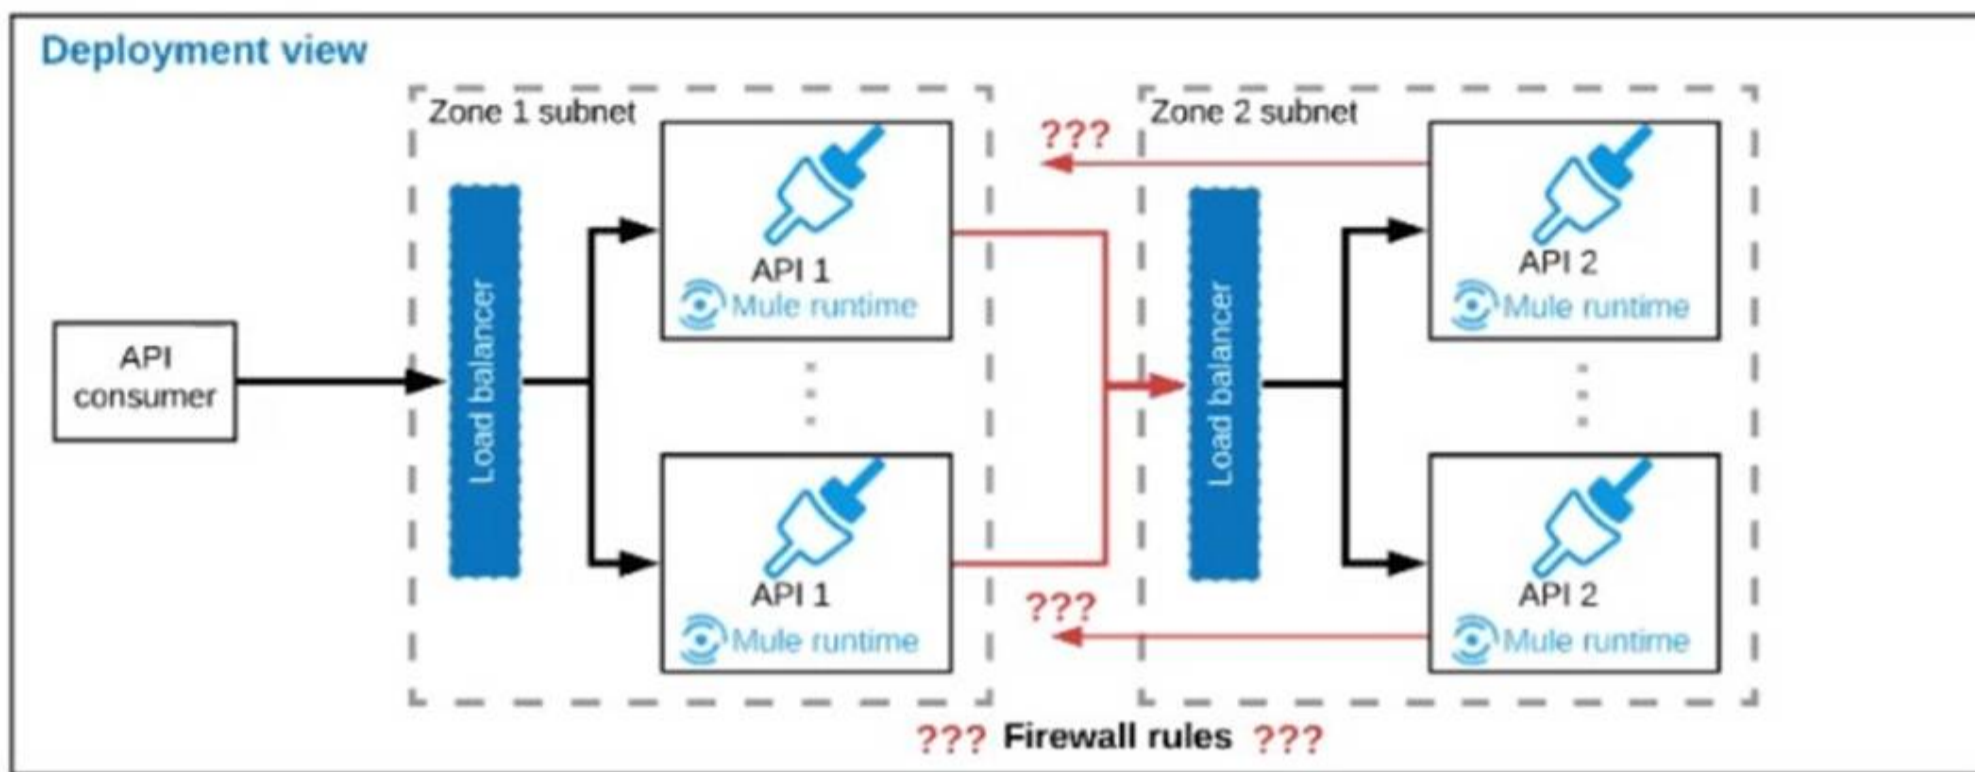
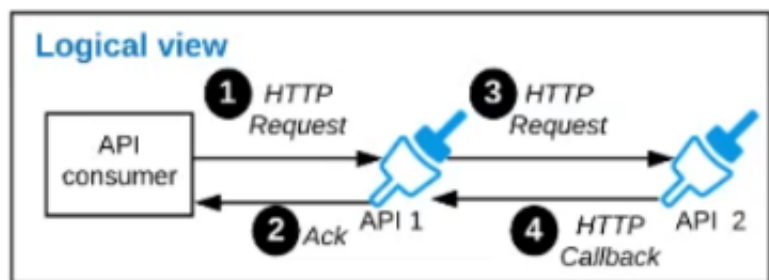
An organization is implementing a Quote of the Day API that caches today's quote. What scenario can use the CloudHub Object Store connector to persist the cache's state?

- 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.

Answer: C

NEW QUESTION 76

Refer to the exhibit.



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

Each API is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.

In the network architecture, how must the firewall rules be configured to enable the above Interaction between API 1 and API 2?

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

Answer: B

NEW QUESTION 77

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

- A. User management
- B. Roles and permissions
- C. Dedicated load balancers
- D. Client Management

Answer: D

NEW QUESTION 78

What is the MuleSoft-recommended best practice to share the connector and configuration information among the APIs?

- A. Build a Mule domain project, add the Database connector and configuration to it, and reference this one domain project from each System API
- B. Build a separate Mule domain project for each API, and configure each of them to use a file on a shared file store to load the configuration information dynamically
- C. Build another System API that connects to the database, and refactor all the other APIs to make requests through the new System API to access the database
- D. Create an API proxy for each System API and share the Database connector configuration with all the API proxies via an automated policy

Answer: A

NEW QUESTION 80

An organization plans to extend its Mule APIs to the EU (Frankfurt) region.

Currently, all Mule applications are deployed to CloudHub 1.0 in the default North American region, from the North America control plane, following this naming convention: {API-name}-{environment} (for example, Orderssapi-dev, Orders-sapi-qa, Orders-sapi-prod, etc.).

There is no network restriction to block communications between APIs.

What strategy should be implemented in order to deploy the same Mule APIs to the CloudHub 1.0 EU region from the North America control plane, as well as to minimize latency between APIs and target users and systems in Europe?

- A. In Runtime Manager, for each Mule application deployment, set the Region property to EU (Frankfurt) and reuse the same Mule application name as in the North American region. Communicate the new urls {API-name}-{environment}.de-ci.cloudhub.io to the consuming API clients in Europe.
- B. In API Manager, set the Region property to EU (Frankfurt) to create an API proxy named {API-name}-proxy-{environment} for each Mule

- application. Communicate the new url {API-name}—proxy—{environment}.de-c1.cloudhub.io to the consuming API clients In Europe.
- C. In Runtime Manager, for each Mule application deployment, leave the Region property blank (default) and change the Mule application name to {API-name}—{environment}.de-cl. Communicate the new urls {API-name}—{environment}.de-ci1.cloudhub.io to the consuming API clients in Europe.
- D. In API Manager, leave the Region property blank (default) to deploy an API proxy named {API-name}~proxy~ (environment).de-cl for each Mule application.
- E. Communicate the new url {API-name}—proxy—{environment}.de-cl.cloudhub.io to the consuming API clients in Europe.

Answer: A

NEW QUESTION 84

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

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

Answer: C

NEW QUESTION 85

A DevOps team has adequate observability of individual system behavior and performance, but it struggles to track the entire lifecycle of each request across different microservices.

Which additional observability approach should this team consider adopting?

- A. Analytics
- B. Metrics
- C. Tracing
- D. Data mining

Answer: C

NEW QUESTION 86

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

Answer: C

NEW QUESTION 90

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

Answer: C

NEW QUESTION 91

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

NEW QUESTION 95

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 fails
- D. An XA transaction writes to both database systems as fast as possible

Answer: B

NEW QUESTION 98

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

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

Answer: D

NEW QUESTION 103

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 non-shared subscription
- D. Use the parameter `primaryNodeOnly= "true"` on the JMS listener

Answer: B

NEW QUESTION 107

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

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

Answer: CE

NEW QUESTION 108

An organization will deploy Mule applications to Cloudhub, Business requirements mandate that all application logs be stored ONLY in an external splunk consolidated logging service and NOT in Cloudhub.

In order to most easily store Mule application logs ONLY in Splunk, how must Mule application logging be configured in Runtime Manager, and where should the log4j2 splunk appender be defined?

- A. Keep the default logging configuration in RuntimeManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support at Mule application deployments.
- B. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in EACH Mule application's log4j2.xml file
- C. Disable Cloudhub logging in Runtime ManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manger to support at Mule application deployments.
- D. Keep the default logging configuration in Runtime Manager Define the Splunk appender in EACH Mule application log4j2.xml file

Answer: B

NEW QUESTION 109

An IT integration delivery team begins a project by gathering all of the requirements, and proceeds to execute the remaining project activities as sequential, non-repeating phases.

Which IT project delivery methodology is this team following?

- A. Kanban
- B. Scrum
- C. Waterfall
- D. Agile

Answer: C

NEW QUESTION 114


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

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

Answer: C

NEW QUESTION 116

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

Answer: C

NEW QUESTION 117

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 4,000 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 4,000 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.

Answer: A

NEW QUESTION 120

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.

Answer: B

NEW QUESTION 122

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.

Answer: C

NEW QUESTION 125

An integration Mule application is deployed to a customer-hosted multi-node Mule 4 runtime duster. The Mule application uses a Listener operation of a JMS connector to receive incoming messages from a JMS queue.

How are the messages consumed by the Mule application?

- A. Depending on the JMS provider's configuration, either all messages are consumed by ONLY the primary cluster node or else ALL messages are consumed by ALL cluster nodes
- B. Regardless of the Listener operation configuration, all messages are consumed by ALL cluster nodes
- C. Depending on the Listener operation configuration, either all messages are consumed by ONLY the primary cluster node or else EACH message is consumed by ANY ONE cluster node
- D. Regardless of the Listener operation configuration, all messages are consumed by ONLY the primary cluster node

Answer: C

NEW QUESTION 127

An organization is evaluating using the CloudHub shared Load Balancer (SLB) vs creating a CloudHub dedicated load balancer (DLB). They are evaluating how this choice affects the various types of certificates used by CloudHub deployed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates. What type of restrictions exist on the types of certificates for the service that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

- A. Underlying Mule applications need to implement own certificates
- B. Only MuleSoft provided certificates can be used for server side certificate
- C. Only self signed certificates can be used
- D. All certificates which can be used in shared load balancer need to get approved by raising support ticket

Answer: B

NEW QUESTION 130

What requirement prevents using Anypoint MQ as the messaging broker for a Mule application?

- A. When the payload sent through the message broker must use XML format
- B. When the payload sent through the message broker must be encrypted
- C. When the messaging broker must support point-to-point messaging
- D. When the messaging broker must be deployed on-premises

Answer: D

NEW QUESTION 131

Which role is primarily responsible for building API implementation as part of a typical MuleSoft integration project?

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

Answer: A

NEW QUESTION 134

A team would like to create a project skeleton that developers can use as a starting point when creating API Implementations with Anypoint Studio. This skeleton should help drive consistent use of best practices within the team.

What type of Anypoint Exchange artifact(s) should be added to Anypoint Exchange to publish the project skeleton?

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

Answer: D

NEW QUESTION 137

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 architect who follows API led approach , what is the the below step you will perform so that you can share document 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

Answer: C

NEW QUESTION 141

An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger of flow for processing a high-volume set of records from the source system and synchronize with the SaaS system using the Batch job scope

After processing 1000 records in a periodic synchronization of 1 lakh records, the replicas in which batch job instance was started went down due to unexpected failure in the runtime fabric environment

What is the consequence of losing the replicas that run the Batch job instance?

- A. The remaining 99000 records will be lost and left and processed
- B. The second replicas will take over processing the remaining 99000 records

- C. A new replacement replica will be available and will be process all 1,00,000 records from scratch leading to duplicate record processing
- D. A new placement replica will be available and will take or processing the remaining 99,000 records

Answer: D

NEW QUESTION 142

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

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

Answer: B

NEW QUESTION 144

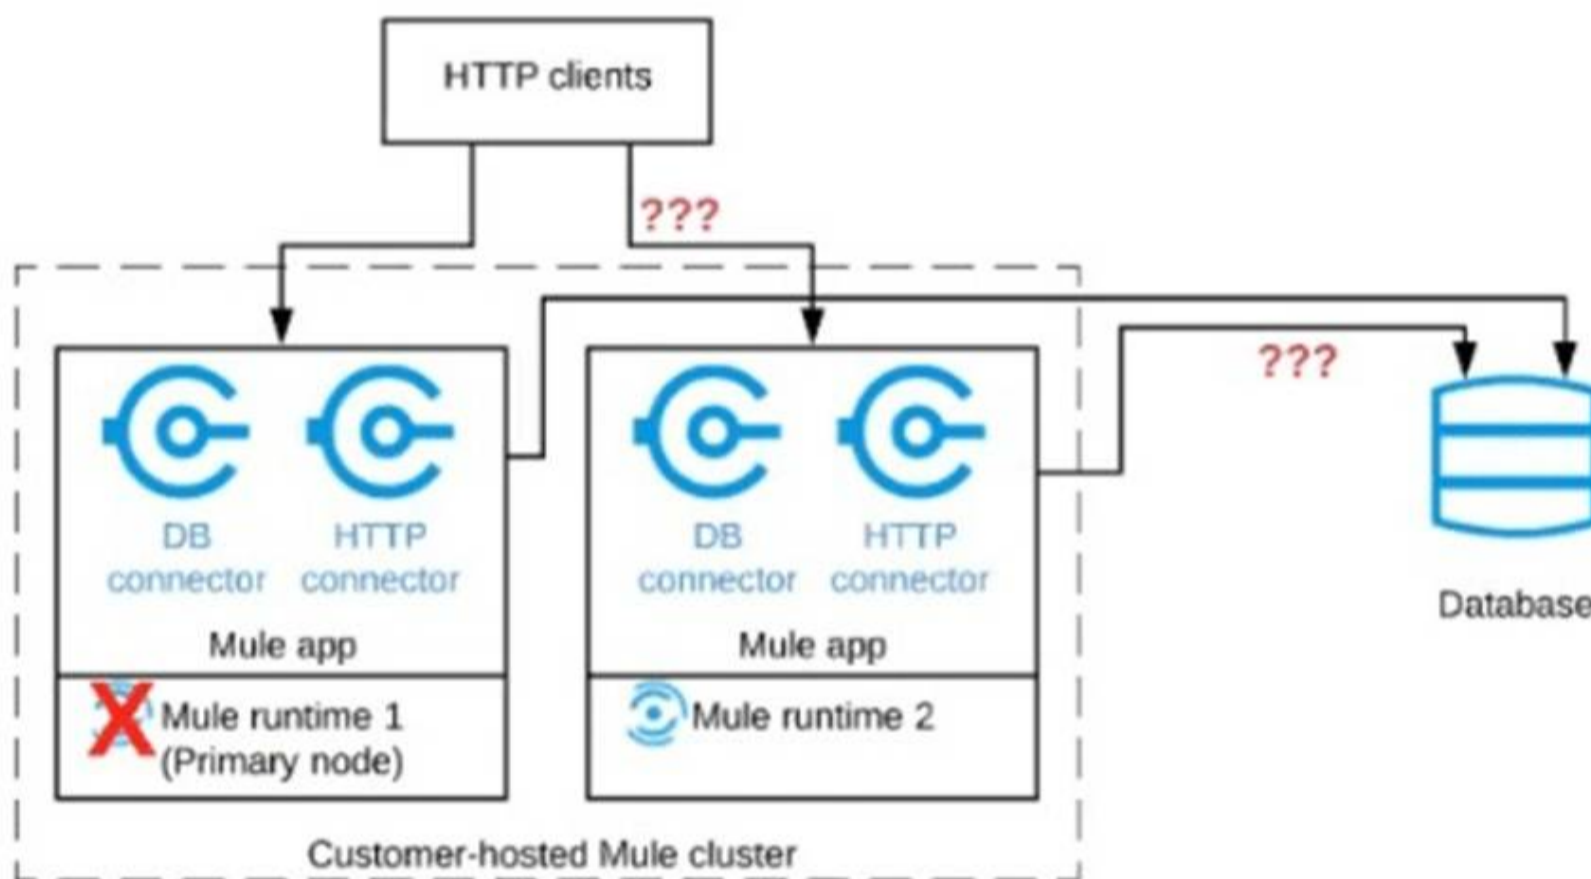
An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders. The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS). At a minimum, what must be stored in each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

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

Answer: C

NEW QUESTION 146

Refer to the exhibit.



A Mule application is deployed to a cluster of two customer-hosted Mute runtimes. The Mute application has a flow that polls a database and another flow with an HTTP Listener.

HTTP clients send HTTP requests directly to individual cluster nodes.

What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has railed, but before that node is restarted?

- A. Database polling continues Only HTTP requests sent to the remaining node continue to be accepted
- B. Database polling stops All HTTP requests continue to be accepted
- C. Database polling continues All HTTP requests continue to be accepted, but requests to the failed node Incur increased latency
- D. Database polling stops All HTTP requests are rejected

Answer: A

NEW QUESTION 147

A manufacturing company is developing a new set of APIs for its retail business. One of the APIs is a Master Look Up API, which is a System API, The API uses a persistent object-store. This API will be used by almost all other APIs to provide master lookup data.

The Master Look Up API is deployed on two CloudHub workers of 0.1 vCore each because there is a lot of master data to be cached. Most of the master

lookup data is stored as a key-value pair. The cache gets refreshed if the key is not found in the cache.

During performance testing, it was determined that the Master Look Up API has a high response time due to the latency of database queries executed to fetch the master lookup data.

What two methods can be used to resolve these performance issues?

Choose 2 answers

- A. Implement the HTTP caching policy for all GET endpoints for the Master Look Up API
- B. Implement an HTTP caching policy for all GET endpoints in the Master Look Up API
- C. Implement locking to synchronize access to the Object Store
- D. Upgrade the vCore size from 0.1 vCore to 0.2 vCore

Answer: AD

NEW QUESTION 150

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 availability of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.

The application must transform and send different formats of this file to three different downstream SFTP locations.

What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

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

Answer: A

NEW QUESTION 155

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

What are suitable interface technologies for this data exchange that are secure, cross-platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

- A. EDJFACT XML over SFTP JSON/REST over HTTPS
- B. SOAP over HTTPS HOP over TLS gRPC over HTTPS
- C. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
- D. CSV over FTP YAML over TLS JSON over HTTPS

Answer: C

NEW QUESTION 159

A set of integration Mule applications, some of which expose APIs, are being created to enable a new business process. Various stakeholders may be impacted by this. These stakeholders are a combination of semi-technical users (who understand basic integration terminology and concepts such as JSON and XML) and technically skilled potential consumers of the Mule applications and APIs.

What is an effective way for the project team responsible for the Mule applications and APIs being built to communicate with these stakeholders using Anypoint Platform and its supplied toolset?

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

Answer: B

NEW QUESTION 162

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

NEW QUESTION 166

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

Answer: A

NEW QUESTION 167

A Mule application is built to support a local transaction for a series of operations on a single database. The mule application has a Scatter-Gather scope that participates in the local transaction.

What is the behavior of the Scatter-Gather when running within this local transaction?

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

Answer: A

NEW QUESTION 172

An organization's IT team must secure all of the internal APIs within an integration solution by using an API proxy to apply required authentication and authorization policies.

Which integration technology, when used for its intended purpose, should the team choose to meet these requirements if all other relevant factors are equal?

- A. API Management (APIM)
- B. Robotic Process Automation (RPA)
- C. Electronic Data Interchange (EDI)
- D. Integration Platform-as-a-service (PaaS)

Answer: A

NEW QUESTION 175

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

Answer: D

NEW QUESTION 178

An organization has chosen Mulesoft for their integration and API platform.

According to the Mulesoft catalyst framework, what would an integration architect do to create achievement goals as part of their business outcomes?

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

Answer: C

NEW QUESTION 179

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

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

Answer: A

NEW QUESTION 181

A corporation has deployed Mule applications to different customer-hosted Mule runtimes. Mule applications deployed to these Mule runtimes are managed by Anypoint Platform.

What needs to be installed or configured (if anything) to monitor these Mule applications from Anypoint Monitoring, and how is monitoring data from each Mule application sent to Anypoint Monitoring?

- A. Enable monitoring of individual Mule applications from the Runtime Manager application settings. Runtime Manager sends monitoring data to Anypoint Monitoring for each deployed Mule application.
- B. Install a Runtime Manager agent on each Mule runtime. Each Runtime Manager agent sends monitoring data from the Mule applications running in its Mule runtime to Runtime Manager, then Runtime Manager sends monitoring data to Anypoint Monitoring.
- C. Leave the out-of-the-box Anypoint Monitoring agent unchanged in its default Mule runtime installation. Each Anypoint Monitoring agent sends monitoring data from the Mule applications running in its Mule runtime to Runtime Manager, then Runtime Manager sends monitoring data to Anypoint Monitoring.
- D. Install an Anypoint Monitoring agent on each Mule runtime. Each Anypoint Monitoring agent sends monitoring data from the Mule applications running in its Mule runtime to Anypoint Monitoring.

Answer: D

NEW QUESTION 184

What approach configures an API gateway to hide sensitive data exchanged between API consumers and API implementations, but can convert tokenized fields back to their original value for other API requests or responses, without having to recode the API implementations?

- A. Create both masking and tokenization formats and use both to apply a tokenization policy in an API gateway to mask sensitive values in message payloads with characters, and apply a corresponding detokenization policy to return the original values to other APIs
- B. Create a masking format and use it to apply a tokenization policy in an API gateway to mask sensitive values in message payloads with characters, and apply a corresponding detokenization policy to return the original values to other APIs
- C. Use a field-level encryption policy in an API gateway to replace sensitive fields in message payload with encrypted values, and apply a corresponding field-level decryption policy to return the original values to other APIs
- D. Create a tokenization format and use it to apply a tokenization policy in an API gateway to replace sensitive fields in message payload with similarly formatted tokenized values, and apply a corresponding detokenization policy to return the original values to other APIs

Answer: D

NEW QUESTION 186

An organization's IT team follows an API-led connectivity approach and must use Anypoint Platform to implement a System API that securely accesses customer data. The organization uses Salesforce as the system of record for all customer data, and its most important objective is to reduce the overall development time to release the System API.

The team's integration architect has identified four different approaches to access the customer data from within the implementation of the System API by using different Anypoint Connectors that all meet the technical requirements of the project.

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

Answer: C

NEW QUESTION 187

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: C

NEW QUESTION 188

An organization is in the process of building automated deployments using a CI/CD process. As a part of automated deployments, it wants to apply policies to API Instances.

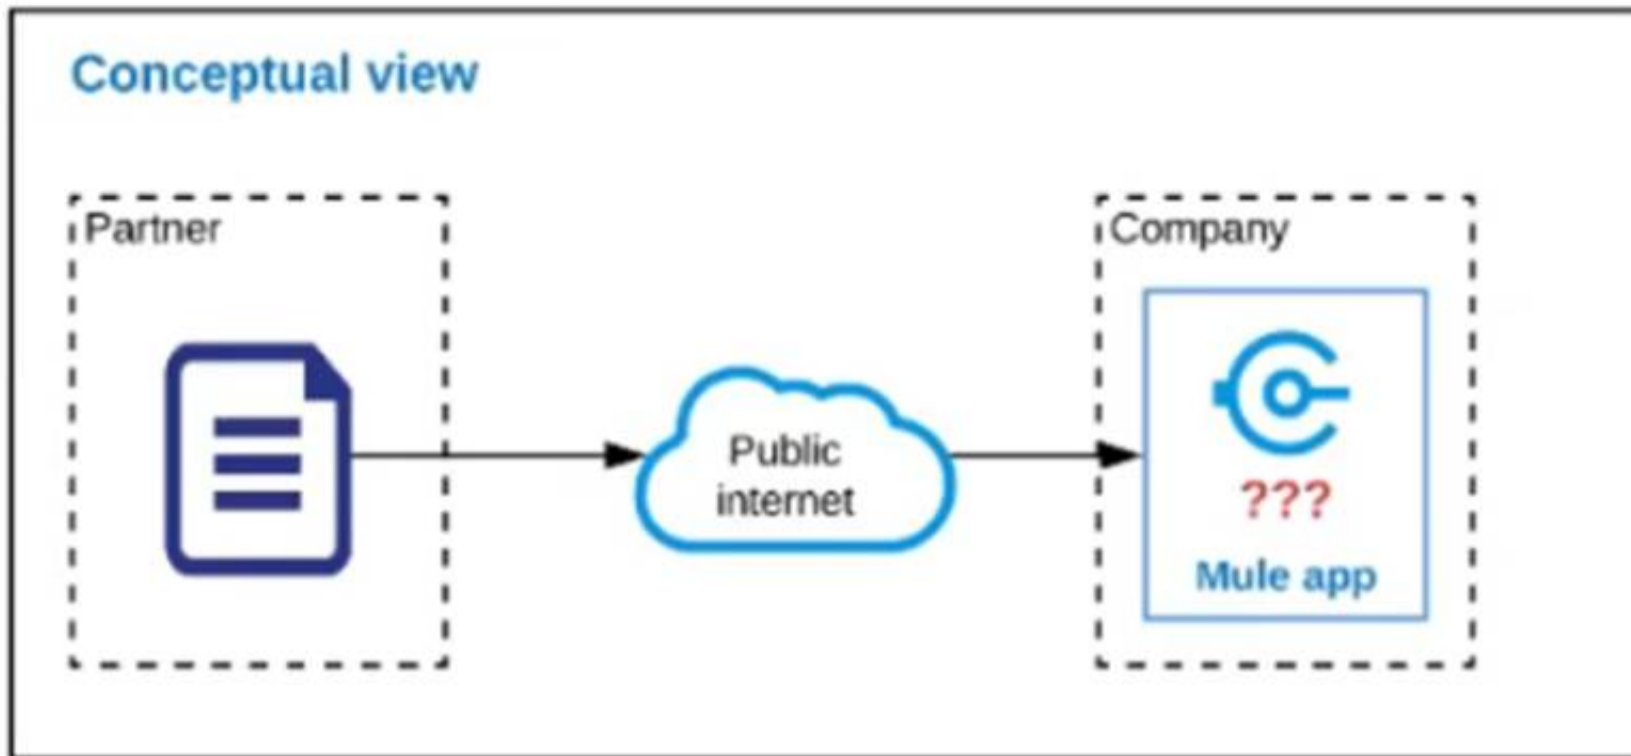
What tool can the organization use to promote and deploy API Manager policies?

- A. Anypoint CLI
- B. MUnit Maven plugin
- C. Mule Maven plugin
- D. Runtime Manager agent

Answer: A

NEW QUESTION 190

Refer to the exhibit.



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

- A. File connector
- B. VM connector
- C. SFTP connector
- D. Object Store connector

Answer: C

NEW QUESTION 193

An organization's governance process requires project teams to get formal approval from all key stakeholders for all new Integration design specifications. An integration Mule application is being designed that interacts with various backend systems. The Mule application will be created using Anypoint Design Center or Anypoint Studio and will then be deployed to a customer-hosted runtime.

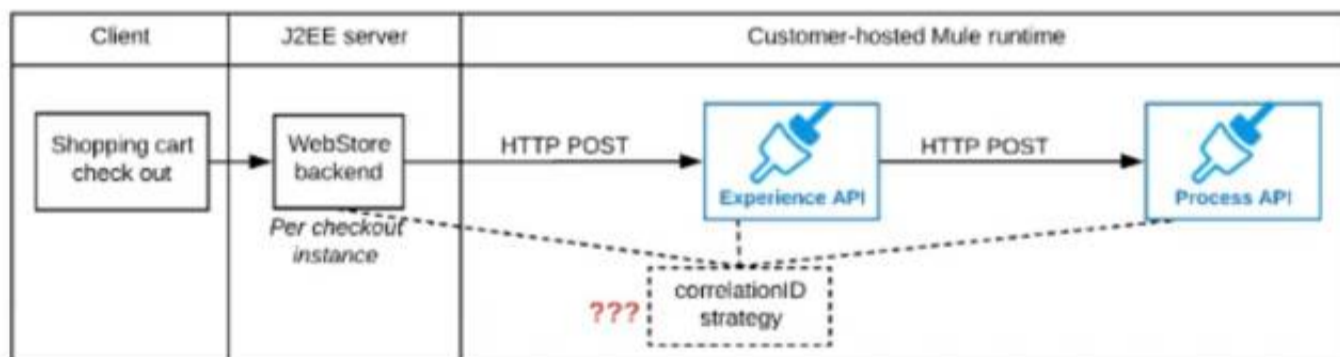
What key elements should be included in the integration design specification when requesting approval for this Mule application?

- A. SLAs and non-functional requirements to access the backend systems
- B. Snapshots of the Mule application's flows, including their error handling
- C. A list of current and future consumers of the Mule application and their contact details
- D. The credentials to access the backend systems and contact details for the administrator of each system

Answer: A

NEW QUESTION 196

Refer to the exhibit.



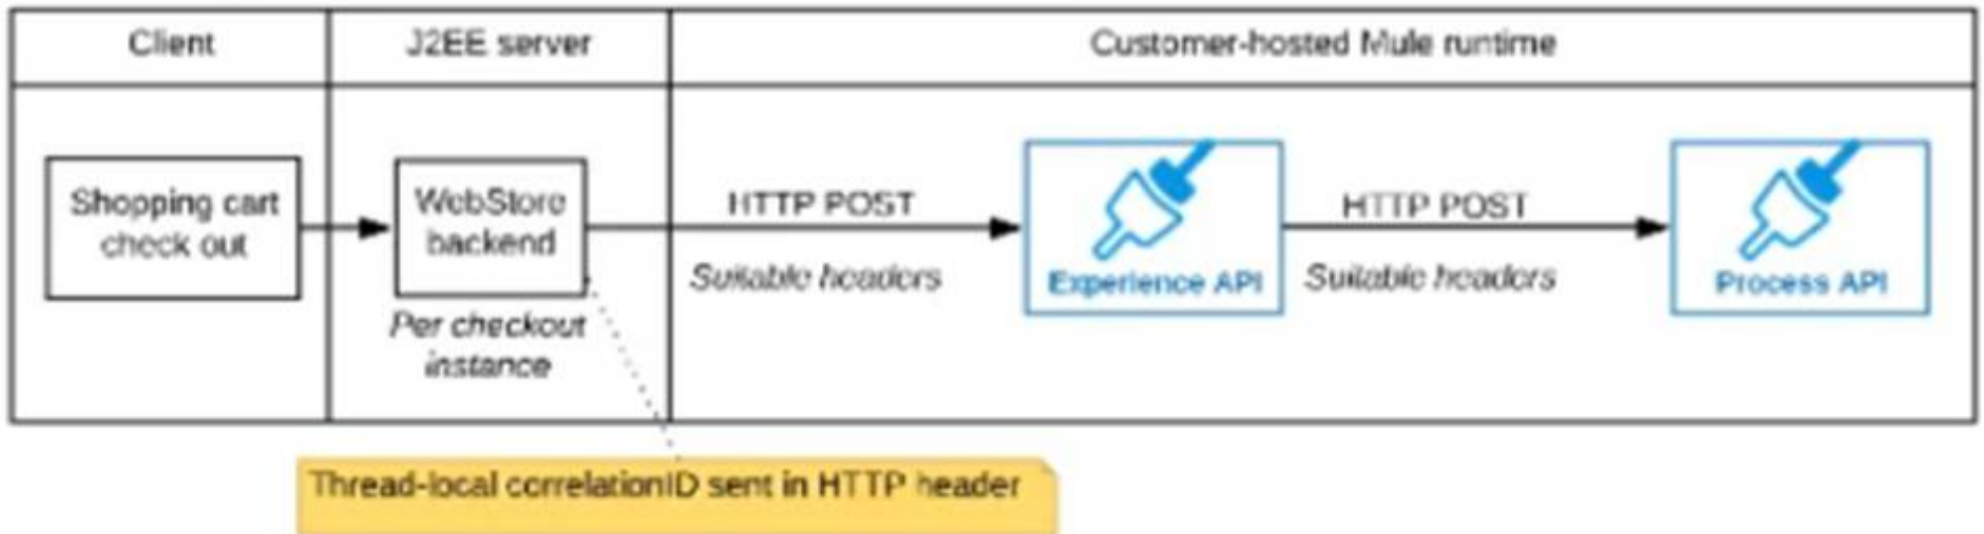
A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations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

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


What is the most efficient way (using the least amount of custom coding or configuration) for the web store backend and the implementations of the Experience API and Process API to participate in end-to-end correlation of the API invocations for each checkout instance?

- A) The web store backend, being a Java EE application, automatically makes use of the thread-local correlation ID generated by the Java EE application server and automatically transmits that to the Experience API using HTTP-standard headers

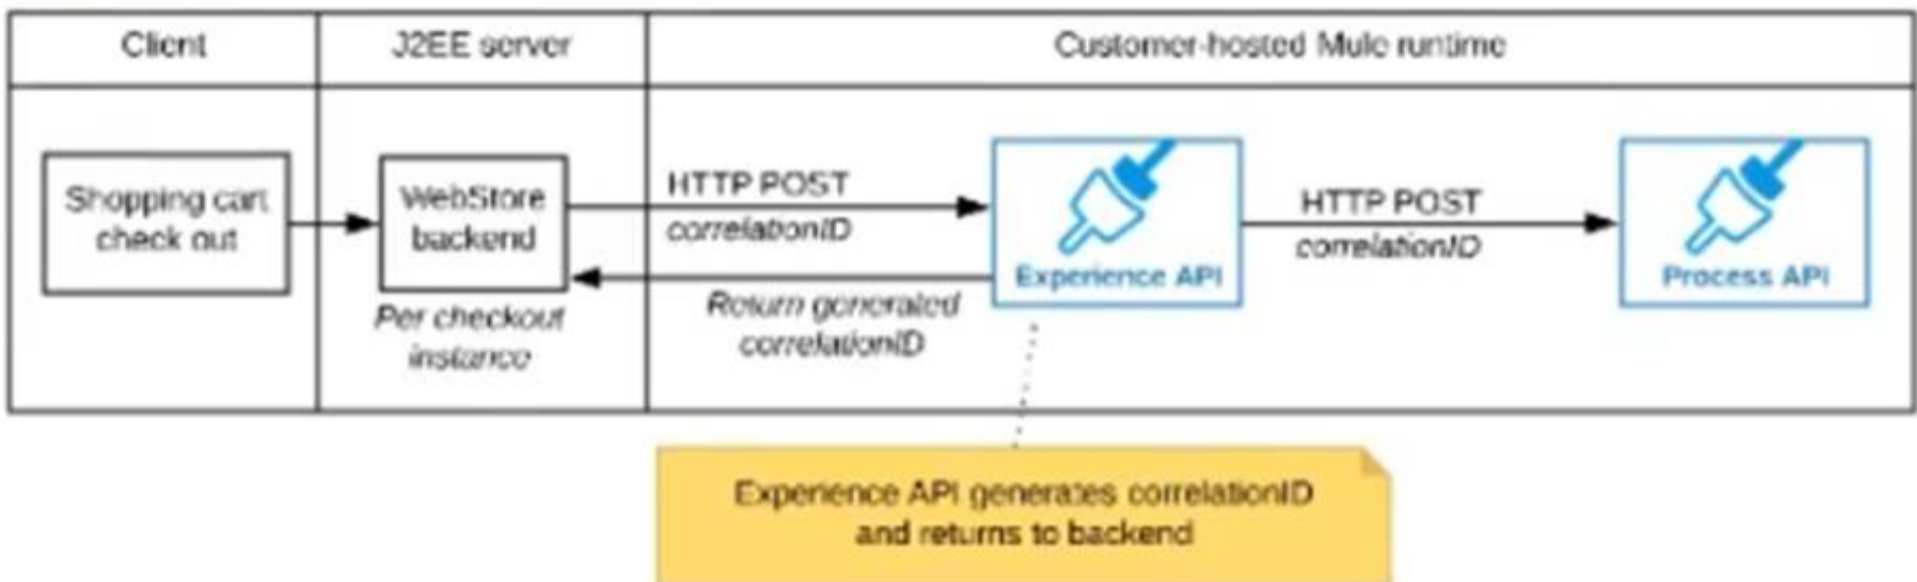
No special code or configuration is included in the web store backend, Experience API, and Process API implementations to generate and manage the correlation ID




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



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



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



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

D. Option D

Answer: B

NEW QUESTION 201

A corporation has deployed multiple mule applications implementing various public and private API's to different cloudhub workers. These API's are 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

Answer: A

NEW QUESTION 202

An integration architect is designing an API that must accept requests from API clients for both XML and JSON content over HTTP/1.1 by default. Which API architectural style, when used for its intended and typical purposes, should the architect choose to meet these requirements?

- A. SOAP
- B. GraphQL
- C. REST
- D. gRPC

Answer: C

NEW QUESTION 204


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

Answer: CD

NEW QUESTION 207

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



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

Answer: C

NEW QUESTION 209

An auto mobile company want to share inventory updates with dealers D1 and D2 asynchronously and concurrently via queues Q1 and Q2. Dealer D1 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

Answer: D

NEW QUESTION 210

A stock trading company handles millions of trades a day and requires excellent performance and reliability within its stock trading system. The company operates a number of event-driven APIs implemented as Mule applications that are hosted on various customer-hosted Mule clusters and needs to enable message exchanges between the APIs within their internal network using shared message queues.

What is an effective way to meet the cross-cluster messaging requirements of its event-driven APIs?

- A. Non-transactional JMS operations with a reliability pattern and manual acknowledgements
- B. Persistent VM queues with automatic acknowledgements
- C. JMS transactions with automatic acknowledgements
- D. extended Architecture (XA) transactions and XA connected components with manual acknowledgements

Answer: C

NEW QUESTION 211

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

Answer: D

NEW QUESTION 213

Which Exchange asset type represents a complete API specification in RAML or OAS format?

- A. Connectors
- B. REST APIs
- C. API Spec Fragments
- D. SOAP APIs

Answer: B

NEW QUESTION 217

Mule application A receives a request Anypoint MQ message REQU with a payload containing a variable-length list of request objects. Application A uses the For Each scope to split the list into individual objects and sends each object as a message to an Anypoint MQ queue.

Service S listens on that queue, processes each message independently of all other messages, and sends a response message to a response queue.

Application A listens on that response queue and must in turn create and publish a response Anypoint MQ message RESP with a payload containing the list of responses sent by service S in the same order as the request objects originally sent in REQU.

Assume successful response messages are returned by service S for all request messages.

What is required so that application A can ensure that the length and order of the list of objects in RESP and REQU match, while at the same time maximizing message throughput?

- A. Use a Scatter-Gather within the For Each scope to ensure response message order. Configure the Scatter-Gather with a persistent object store.
- B. Perform all communication involving service S synchronously from within the For Each scope, so objects in RESP are in the exact same order as request objects in REQU.
- C. Use an Async scope within the For Each scope and collect response messages in a second For Each scope in the order in which they arrive, then send RESP using this list of responses.
- D. Keep track of the list length and all object indices in REQU, both in the For Each scope and in all communication involving service S. Use persistent storage when creating RESP.

Answer: D

NEW QUESTION 218

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

Answer: D

NEW QUESTION 223

In preparation for a digital transformation initiative, an organization is reviewing related IT integration projects that failed for various reasons. According to MuleSoft's surveys of global IT leaders, what is a common cause of IT project failure that this organization may likely discover in its assessment?

- A. Following an Agile delivery methodology
- B. Reliance on an Integration-Platform-as-a-Service (iPaaS)
- C. Spending too much time on enablement
- D. Lack of alignment around business outcomes

Answer: D

NEW QUESTION 224

As a part of a project, existing Java implementation is being migrated to MuleSoft. Business is very tight on the budget and wishes to complete the project in the 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

Answer: A

NEW QUESTION 228

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 integration 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 load the properties at startup by invoking that credentials service

Answer: D

NEW QUESTION 230

An organization plans to migrate its deployment environment from an on-premises cluster to a Runtime Fabric (RTF) cluster. The on-premises Mule applications are currently configured with persistent object stores.

There is a requirement to enable Mule applications deployed to the RTF cluster to store and share data across application replicas and through restarts of the entire RTF cluster.

How can these reliability requirements be met?

- A. Replace persistent object stores with persistent VM queues in each Mule application deployment
- B. Install the Object Store pod on one of the cluster nodes
- C. Configure Anypoint Object Store v2 to share data between replicas in the RTF cluster
- D. Configure the Persistence Gateway in the RTF installation

Answer: C

NEW QUESTION 234

An organization is building out a test suite for their application using MUnit.

The Integration Architect has recommended using Test Recorder in Anypoint Studio to record the processing flows and then configure unit tests based on the captured events.

What is a core consideration that must be kept in mind while using Test Recorder?

- A. The Recorder supports loops where the structure of the data being tested changes inside the iteration
- B. Mocking values resulting from parallel processes are possible and will not affect the execution of the processors that follow in the test
- C. The Recorder supports mocking a message before or inside a Foreach processor
- D. Tests for flows cannot be created if Mule errors are raised inside the flows, even if the errors are handled by On-Error Continue error handlers

Answer: D

NEW QUESTION 237

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 table 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 send new employees data to cache. Use an object store caching strategy and set expiration interval to 1-hour

Answer: A

NEW QUESTION 239

As a part of project requirement, client will send a stream of data to mule application. Payload size can vary between 10mb to 5GB. Mule application is required to transform the data and send across multiple sftp servers. Due to the cost cuttings in the organization, mule application can only be allocated one worker with size of 0.2 vCore.

As an integration architect, which streaming strategy you would suggest to handle this scenario?

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

Answer: D

NEW QUESTION 240

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 `https://api.acme.com/partners/v1`. The API implementation is estimated to require deployment to 5 CloudHub workers.

The retailer has obtained a public X.509 certificate for the name `api.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 workers
Create 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 workers
Create a CNAME for `api.acme.com` pointing to the SLB's A record

Answer: C

NEW QUESTION 243

A stock broking company makes use of CloudHub VPC to deploy Mule applications. Mule application needs to connect to a database application in the customer's 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 customer's 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 customer's devices corporate data center
- D. Setup VPC peering with the customer's on-premise corporate data center and Anypoint VPN to AWS VPC

Answer: B

NEW QUESTION 247

When a Mule application using VM queues is deployed to a customer-hosted cluster or multiple CloudHub v1.0 workers/replicas, how are messages consumed across the nodes?

- A. Sequentially, from a dedicated Anypoint MQ queue
- B. Sequentially, only from the primary node
- C. In a non-deterministic way
- D. Round-robin, within an XA transaction

Answer: C

NEW QUESTION 248

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

Answer: A

NEW QUESTION 251

Which key DevOps practice and associated Anypoint Platform component should a MuleSoft integration team adopt to improve delivery quality?

- A. A Continuous design with API Designer
- B. Automated testing with MUnit

- C. Passive monitoring with Anypoint Monitoring
- D. Manual testing with Anypoint Studio

Answer: B

NEW QUESTION 255

An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform. What feature of Anypoint Platform should be used to meet this requirement?

- A. Enterprise Security Module
- B. Client ID Secret
- C. Federated Identity Management
- D. Federated Client Management

Answer: C

NEW QUESTION 260

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 application
- 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

NEW QUESTION 264

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

Answer: B

NEW QUESTION 265

During a planning session with the executive leadership, the development team director presents plans for a new API to expose the data in the company's order database. An earlier effort to build an API on top of this data failed, so the director is recommending a design-first approach.

Which characteristics of a design-first approach will help make this API successful?

- A. Building MUnit tests so administrators can confirm code coverage percentage during deployment
- B. Publishing the fully implemented API to Exchange so all developers can reuse the API
- C. Adding global policies to the API so all developers automatically secure the implementation before coding anything
- D. Developing a specification so consumers can test before the implementation is built

Answer: D

NEW QUESTION 266

A company is modernizing its legal systems to accelerate access to applications and data while supporting the adoption of new technologies. The key to achieving this business goal is unlocking the companies' key systems and data including microservices running 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 APIs in mule runtime that are able to scale, connect to on premises systems and migrate as needed.

Which runtime deployment option supports company's goals?

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

Answer: C

NEW QUESTION 271

A mule application must periodically process a large dataset which varies from 6 GB to 8 GB from a back-end database and write transform data to an FTPS server using a properly configured batch 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 repeatable store strategy for reading records from the database and batch aggregator with streaming to write to FTPS
- C. Use streaming with a file store repeatable strategy for reading records from the database and batch aggregator with an optimal size

D. Use streaming with a file store repeatable strategy reading records from the database and batch aggregator without any required configuration

Answer: A

NEW QUESTION 276

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.

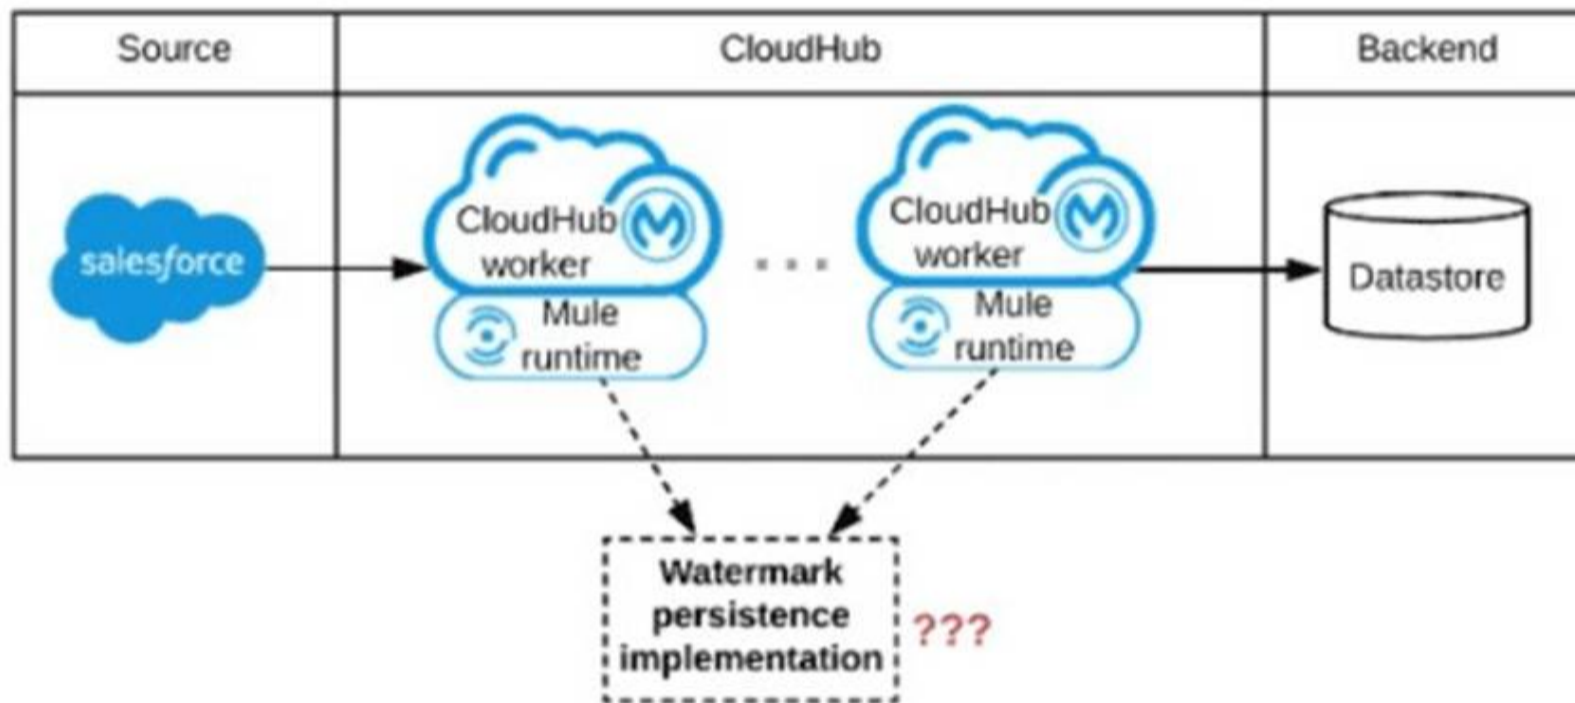
Which steps must the organization take to resolve this error and allow customers to access both the API's?

- 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 279

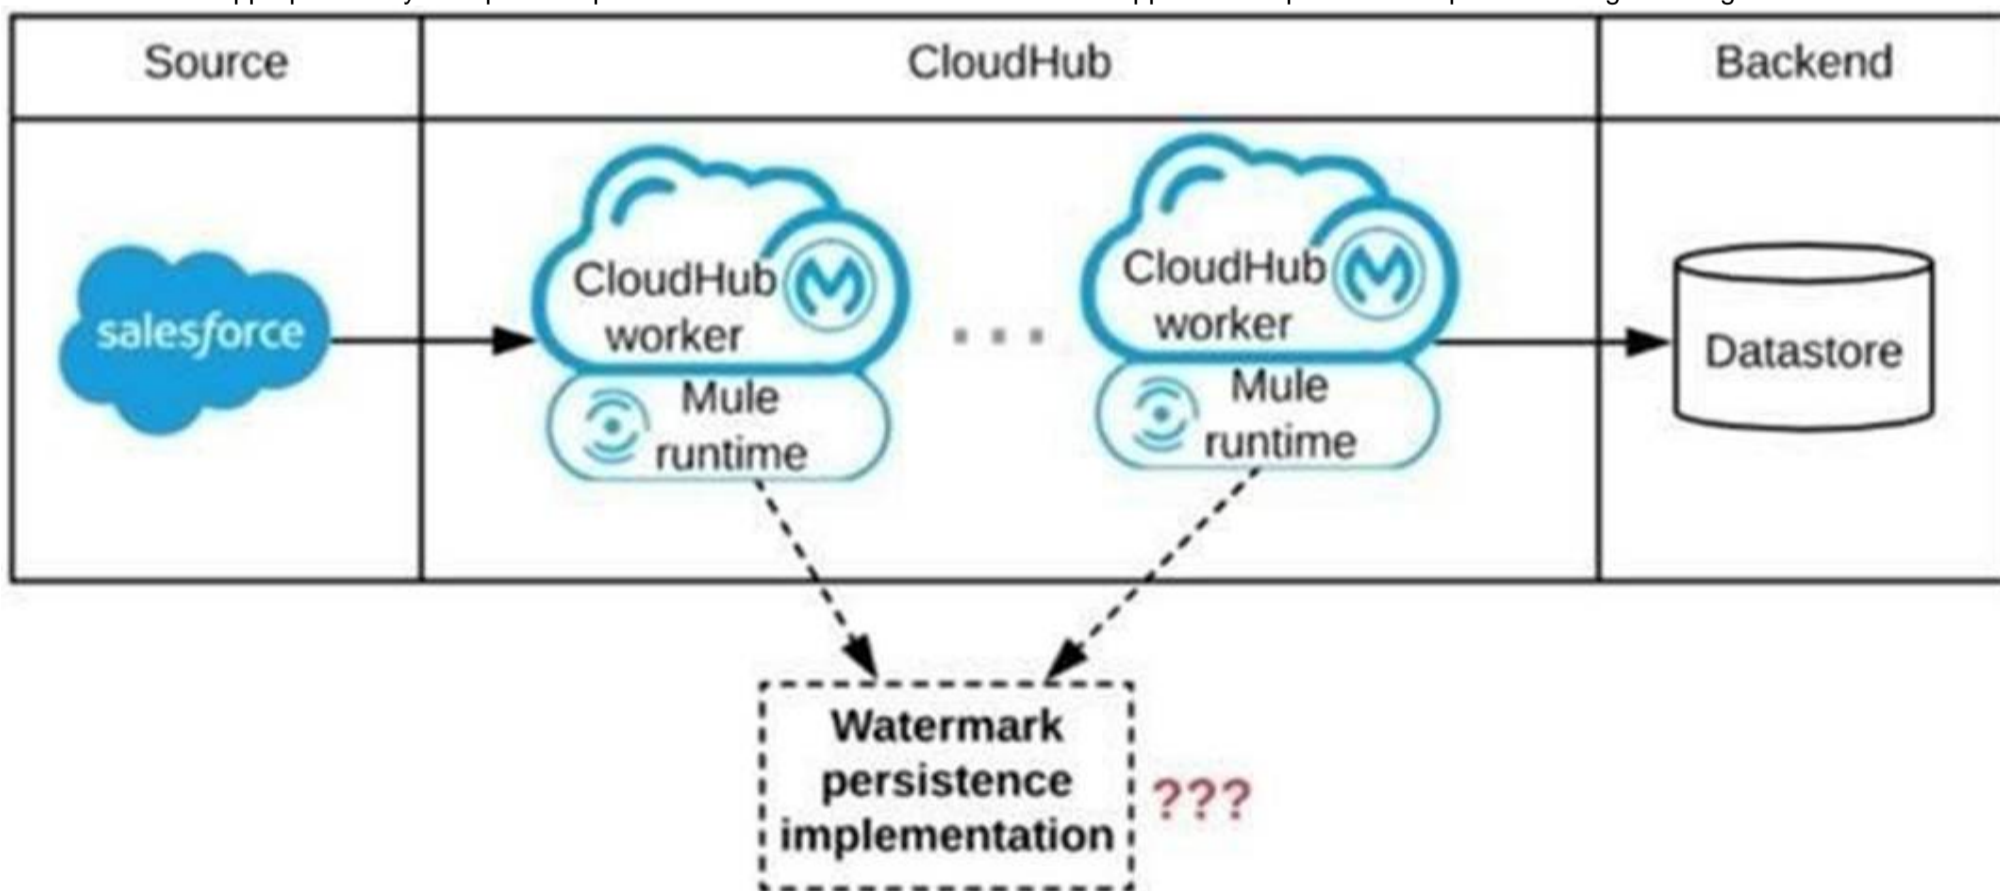
Refer to the exhibit.



A Mule application is being designed to be deployed to several CloudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.

A watermark will be used to only retrieve those Satesforce Accounts that have been modified since the last time the integration logic ran.

What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?



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

Answer: B

NEW QUESTION 283

An automation engineer needs to write scripts to automate the steps of the API lifecycle, including steps to create, publish, deploy and manage APIs and their implementations in Anypoint Platform.

What Anypoint Platform feature can be used to automate the execution of all these actions in scripts in the easiest way without needing to directly invoke the Anypoint Platform REST APIs?

- A. Automated Policies in API Manager
- B. Runtime Manager agent
- C. The Mule Maven Plugin
- D. Anypoint CLI

Answer: D

NEW QUESTION 285

An integration Mule application is being designed to process orders by submitting them to a backend 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 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.

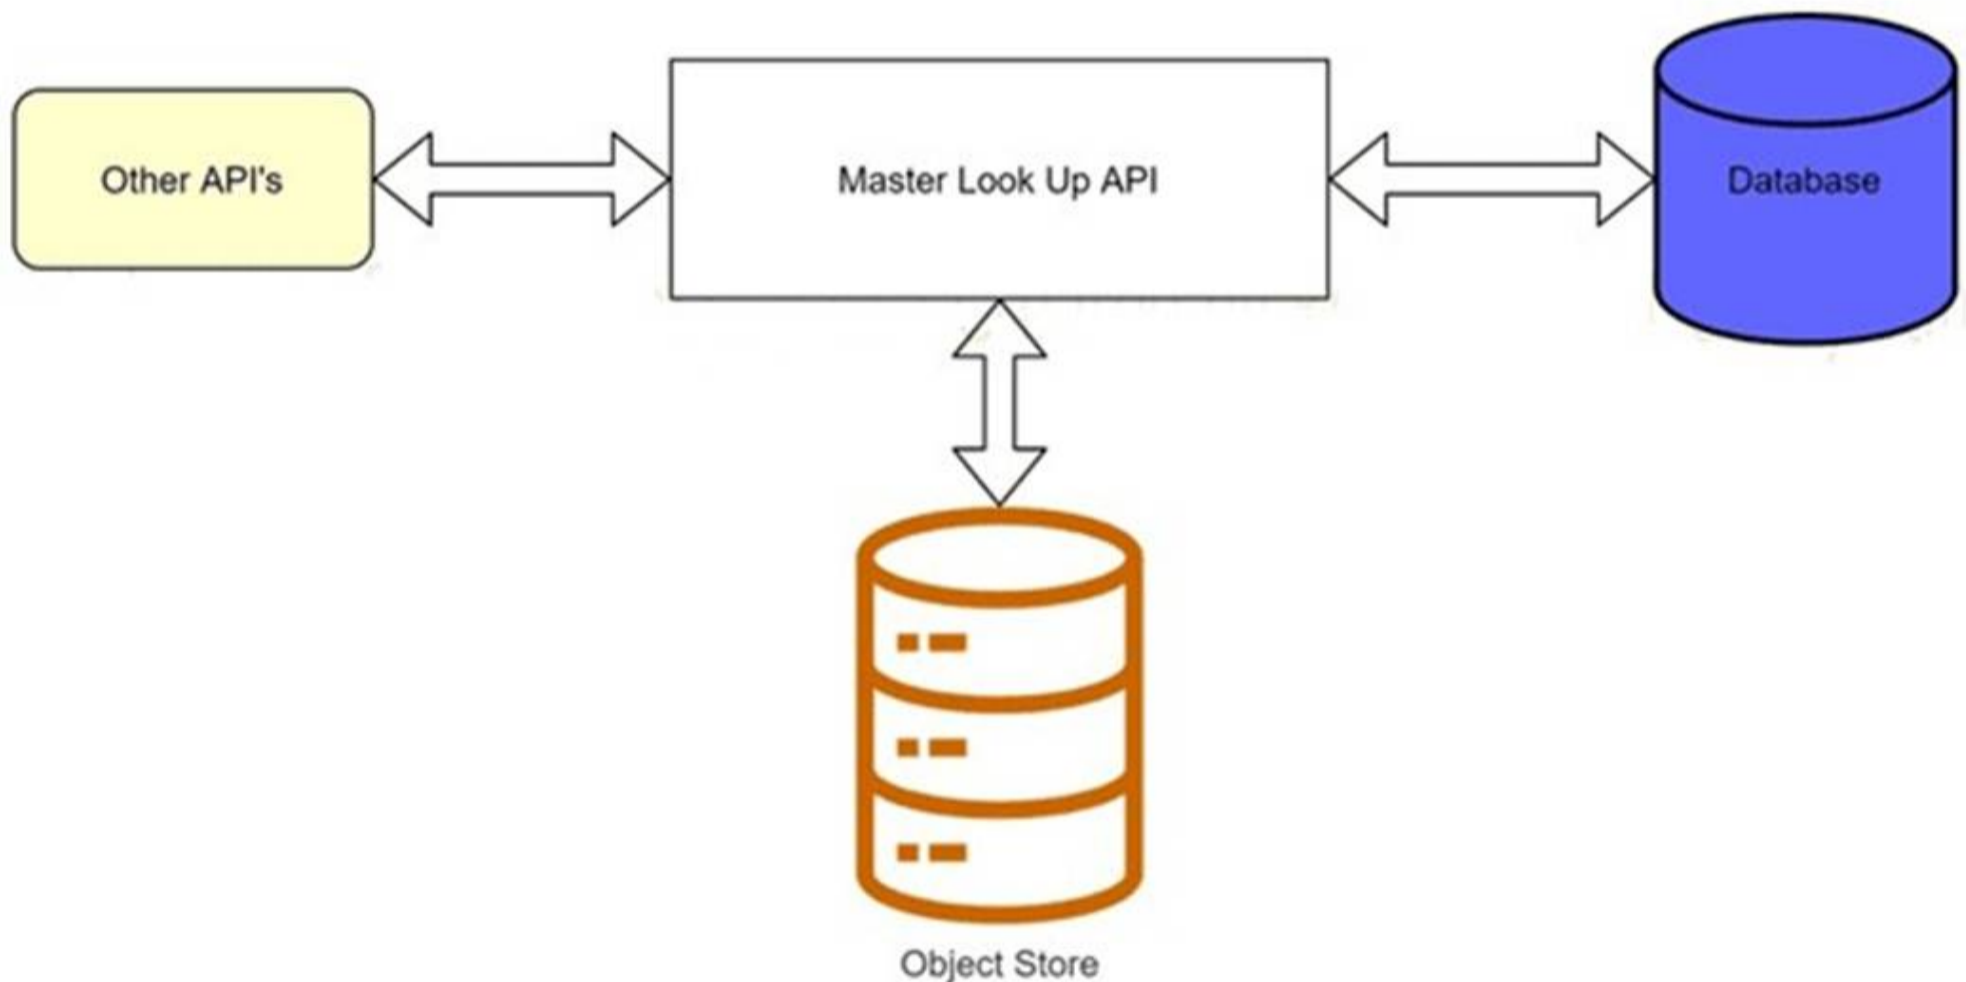
What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

- 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

NEW QUESTION 288

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



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

Doing performance testing it was observed that the Master lookup API has a higher response time due to database queries execution to fetch the master lookup data.

Due to this performance issue, go-live of the online business is on hold which could cause potential financial loss to Bank.

As an integration architect, which of the below option you would suggest to resolve performance issue?

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

Answer: A

NEW QUESTION 289

An application load balancer routes requests to a RESTful web API secured by Anypoint Flex Gateway.

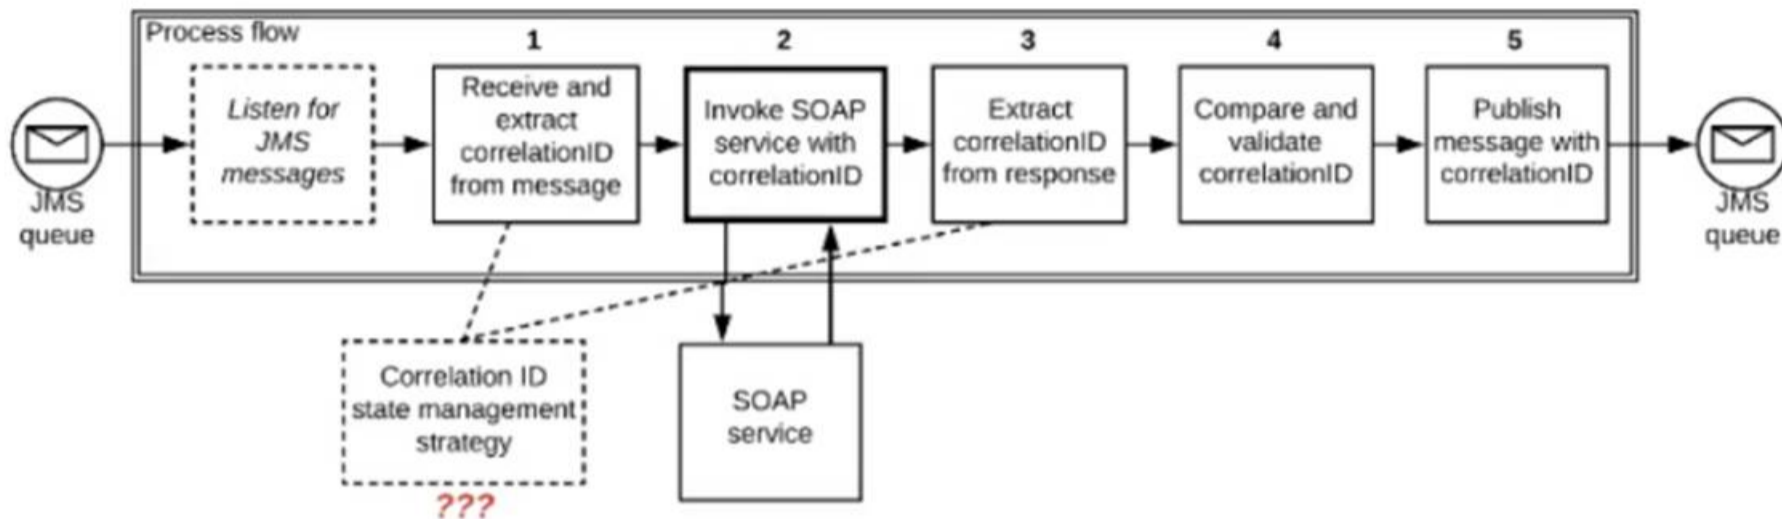
Which protocol is involved in the communication between the load balancer and the Gateway?

- A. SFTP
- B. HTTPS
- C. LDAP
- D. SMTP

Answer: B

NEW QUESTION 290

Refer to the exhibit.



A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to receive JMS messages from a JMS queue. To process each received JMS message, the following steps are performed in a flow:

Step 1: The JMS Correlation ID header is read from the received JMS message. Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS,

passing the JMS Correlation ID as one parameter in the SOAP request.

Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.

Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.

Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.


Where should the Mule application store the JMS Correlation ID values received in Step 1 and Step 3 so that the validation in Step 4 can be performed, while also making the overall Mule application highly available, fault-tolerant, performant, and maintainable?

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

Answer: C

NEW QUESTION 293

Refer to the exhibit.



A customer is running Mule applications on Runtime Fabric for Self-Managed Kubernetes (RTF-BYOKS) in a multi-cloud environment.

Based on this configuration, how do Agents and Runtime Manager communicate, and what is exchanged between them?

- A. BLOCKING_IO, UBER
- B. UBER, Dedicated NIO Selector Pool
- C. CPU_LITE, CPU_INTENSIVE
- D. Shared NIO Selector Pool, CPU_LITE

Answer: D

NEW QUESTION 294

An organization designing a hybrid, load balanced, single cluster production environment. Due to performance service level agreement goals, it is looking into running the Mule applications in an active-active multi node cluster configuration.

What should be considered when running its Mule applications in this type of environment?

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

Answer: B

NEW QUESTION 297

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

Answer: A

NEW QUESTION 298

A developer is examining the responses from a RESTful web service that is compliant with the Hypertext Transfer Protocol (HTTP/1.1) as defined by the Internet Engineering Task Force (IETF).

In this HTTP/1.1-compliant web service, which class of HTTP response status codes should be specified to indicate when client requests are successfully received, understood, and accepted by the web service?

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

Answer: B

NEW QUESTION 303

An organization uses a four(4) node customer hosted Mule runtime cluster to host one(1) stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution. Each node in the cluster has been sized to be able to accept four(4) times the current number of requests.

Two(2) nodes in the cluster experience a power outage and are no longer available. The load balancer directs the outage and blocks the two unavailable nodes from receiving further HTTP requests.

What performance-related consequence is guaranteed to happen to average, assuming the remaining cluster nodes are fully operational?

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

Answer: A

NEW QUESTION 308

According to MuleSoft, what Action should an IT organization take regarding its technology assets in order to close the IT delivery.

- A. Make assets easily discoverable via a central repository
- B. Focus project delivery efforts on custom assets that meet the specific requirements of each individual line of business
- C. Create weekly meetings that all members of IT attend to present justification and request approval to use existing assets
- D. Hire additional staff to meet the demand for asset creation required for approved projects and timelines

Answer: A

NEW QUESTION 311

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 included 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 316

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

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

Answer: D

NEW QUESTION 319

A new Mule application under development must implement extensive data transformation logic. Some of the data transformation functionality is already available as external transformation services that are mature and widely used across the organization; the rest is highly specific to the new Mule application. The organization follows a rigorous testing approach, where every service and application must be extensively acceptance tested before it is allowed to go into production. What is the best way to implement the data transformation logic for this new Mule application while minimizing the overall testing effort?

- A. Implement and expose all transformation logic as mlaoservices using DataWeave, so it can be reused by any application component that needs it, including the new Mule application
- B. Implement transformation logic in the new Mute application using DataWeave, replicating the transformation logic of existing transformation services
- C. Extend the existing transformation services with new transformation logic and Invoke them from the new Mule application
- D. Implement transformation logic in the new Mute application using DataWeave, invoking existing transformation services when possible

Answer: D

NEW QUESTION 324

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

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

Answer: C

NEW QUESTION 328

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

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

Answer: AD

NEW QUESTION 329

What is required before an API implemented using the components of Anypoint Platform can be managed and governed (by applying API policies) on Anypoint Platform?

- A. The API must be published to Anypoint Exchange and a corresponding API instance ID must be obtained from API Manager to be used in the API implementation
- B. The API implementation source code must be committed to a source control management system (such as GitHub)
- C. A RAML definition of the API must be created in API designer so it can then be published to Anypoint Exchange
- D. The API must be shared with the potential developers through an API portal so API consumers can interact with the API

Answer: A

NEW QUESTION 331

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

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

Answer: B

NEW QUESTION 334

Anypoint Exchange is required to maintain the source code of some of the assets committed to it, such as Connectors, Templates, and API specifications. What is the best way to use an organization's source-code management (SCM) system in this context?

- A. Organizations should continue to use an SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching, and merging
- B. Organizations need to use Anypoint Exchange as the main SCM system to centralize versioning and avoid code duplication
- C. Organizations can continue to use an SCM system of their choice for branching and merging, as long as they follow the branching and merging strategy enforced by Anypoint Exchange
- D. Organizations need to point Anypoint Exchange to their SCM system so Anypoint Exchange can pull source code when requested by developers and provide it to Anypoint Studio

Answer: B

NEW QUESTION 339

.....

Relate Links

100% Pass Your MuleSoft-Integration-Architect-I Exam with ExamBible Prep Materials

<https://www.exambible.com/MuleSoft-Integration-Architect-I-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>