

MuleSoft

Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1



NEW QUESTION 1

A system API EmployeeSAPI is used to fetch employee's data from an underlying SQL database.

The architect must design a caching strategy to query the database only when there is an update to the employees stable or else return a cached response in order to minimize the number of redundant transactions being handled by the database.

What must the architect do to achieve the caching objective?

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

Answer: A

NEW QUESTION 2

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 3

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

Explanation:

* Clustering is a group of servers or mule runtime which acts as a single unit.

* Mulesoft Enterprise Edition supports scalable clustering to provide high availability for the Mulesoft application.

* In simple terms, virtual servers composed of multiple nodes and they communicate and share information through a distributed shared memory grid.

* By default, Mulesoft ensures the High availability of applications if clustering is implemented.

* Let's consider the scenario one of the nodes in cluster crashed or goes down and under maintenance. In such cases, Mulesoft will ensure that requests are processed by other nodes in the cluster. Mulesoft clustering also ensures that the request is load balanced between all the nodes in a cluster.

* Clustering is only supported by on-premise Mule runtime and it is not supported in Cloudhub.

Correct answer is External monitoring tools or log aggregators must be configured to recognize the new nodes

* Rest of the options are automatically taken care of when a new node is added in cluster.

NEW QUESTION 4

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

Explanation:

Correct answer is Client Management

* Anypoint Platform acts as a client provider by default, but you can also configure external client providers to authorize client applications.

* As an API owner, you can apply an OAuth 2.0 policy to authorize client applications that try to access your API. You need an OAuth 2.0 provider to use an OAuth 2.0 policy.

* You can configure more than one client provider and associate the client providers with different environments. If you configure multiple client providers after you have already created environments, you can associate the new client providers with the environment.

* You should review the existing client configuration before reassigning client providers to avoid any downtime with existing assets or APIs.

* When you delete a client provider from your master organization, the client provider is no longer available in environments that used it.

* Also, assets or APIs that used the client provider can no longer authorize users who want to access them.

-----MuleSoft

Reference: <https://docs.mulesoft.com/access-management/managing-api-clients>

<https://www.folkstalk.com/2019/11/mulesoft-integration-and-platform.html>

NEW QUESTION 5

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

Explanation:

As the stakeholders are semitechnical users, preferred option is 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

NEW QUESTION 6

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

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

Answer: A

Explanation:

1- way SSL : The server presents its certificate to the client and the client adds it to its list of trusted certificate. And so, the client can talk to the server.

2- way SSL: The same principle but both ways. i.e. both the client and the server has to establish trust between themselves using a trusted certificate. In this way of a digital handshake, the server needs to present a certificate to authenticate itself to client and client has to present its certificate to server.

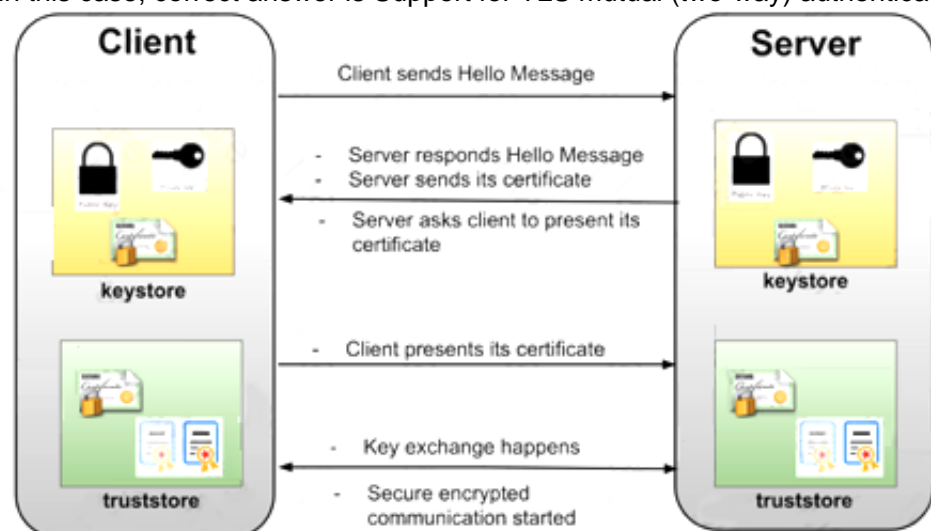
* TLS is a cryptographic protocol that provides communications security for your Mule app.

* TLS offers many different ways of exchanging keys for authentication, encrypting data, and guaranteeing message integrity. Keystores and Truststores. Truststore and keystore contents differ depending on whether they are used for clients or servers:

For servers: the truststore contains certificates of the trusted clients, the keystore contains the private and public key of the server. For clients: the truststore contains certificates of the trusted servers, the keystore contains the private and public key of the client.

Adding both a keystore and a truststore to the configuration implements two-way TLS authentication also known as mutual authentication.

* in this case, correct answer is Support for TLS mutual (two-way) authentication with HTTP clients.



NEW QUESTION 7

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: A

NEW QUESTION 8

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 9

A global organization operates datacenters in many countries. There are private network links between these datacenters because all business data (but NOT metadata) must be exchanged over these private network connections.

The organization does not currently use AWS in any way.

The strategic decision has just been made to rigorously minimize IT operations effort and investment going forward.

What combination of deployment options of the Anypoint Platform control plane and runtime plane(s) best serves this organization at the start of this strategic journey?

- A. MuleSoft-hosted Anypoint Platform control plane CloudHub Shared Worker Cloud in multiple AWS regions
 B. Anypoint Platform - Private Cloud Edition Customer-hosted runtime plane in each datacenter
 C. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in multiple AWS regions
 D. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in each datacenter

Answer: D

Explanation:

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

Hybrid is the best choice to start. Mule hosted Control plane and Customer hosted Runtime to start with. Once they mature in cloud migration, everything can be in Mule hosted.

NEW QUESTION 10

A mule application is being designed to perform product orchestration. The Mule application needs to join together the responses from an inventory API and a Product Sales History API with the least latency.

To minimize the overall latency. What is the most idiomatic (used for its intended purpose) design to call each API request in the Mule application?

- A. Call each API request in a separate lookup call from Dataweave reduce operator
 B. Call each API request in a separate route of a Scatter-Gather
 C. Call each API request in a separate route of a Parallel For Each scope
 D. Call each API request in a separate Async scope

Answer: B

Explanation:

Scatter-Gather sends a request message to multiple targets concurrently. It collects the responses from all routes, and aggregates them into a single message.

NEW QUESTION 10

What is true about the network connections when a Mule application uses a JMS connector to interact with a JMS provider (message broker)?

- A. To complete sending a JMS message, the JMS connector must establish a network connection with the JMS message recipient
 B. To receive messages into the Mule application, the JMS provider initiates a network connection to the JMS connector and pushes messages along this connection
 C. The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider
 D. The AMQP protocol can be used by the JMS connector to portably establish connections to various types of JMS providers

Answer: C

Explanation:

* To send message or receive JMS (Java Message Service) message no separate network connection need to be established. So option A, C and D are ruled out. Correct Answer The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider.

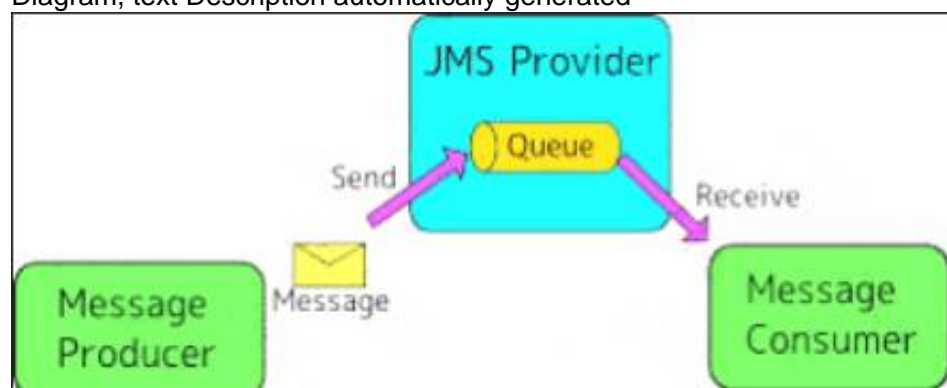
* JMS Connector enables sending and receiving messages to queues and topics for any message service that implements the JMS specification.

* JMS is a widely used API for message-oriented middleware.

* It enables the communication between different components of a distributed application to be loosely coupled, reliable, and asynchronous.

MuleSoft Doc Reference: <https://docs.mulesoft.com/jms-connector/1.7/>

Diagram, text Description automatically generated



NEW QUESTION 15

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

- A. Implement a scatter-gather scope to invoke the two vendor applications on two different routes. Use the Until-Successful scope to implement the retry mechanism for timeout errors on each route.
- B. Implement a Choice scope to invoke the two vendor applications on two different routes. Use the try-catch scope to implement the retry mechanism for timeout errors on each route.
- C. Implement a For-Each scope to invoke the two vendor applications. Use the until successful scope to implement the retry mechanism for the timeout errors.
- D. Implement Round-Robin scope to invoke the two vendor applications on two different routes. Use the Try-Catch scope to implement the retry mechanism for timeout errors on each route.

Answer: A

NEW QUESTION 19

An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.

What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

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

Answer: B

NEW QUESTION 20

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

- A. View object store entries
- B. Replay an unsuccessful message
- C. Set a particular log4j2 log level to TRACE
- D. Deploy a Mule application

Answer: C

Explanation:

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

<https://docs.mulesoft.com/mule-runtime/3.9/jmx-management>

NEW QUESTION 22

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

Explanation:

Set the Anypoint MQ connector operation to publish or consume messages, or to accept (ACK) or not accept (NACK) a message.

NEW QUESTION 26

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

Explanation:

High availability is about up-time of your application

A) High availability can be achieved only in CloudHub isn't correct statement. It can be achieved in customer hosted runtime planes as well
 B) An object store is a facility for storing objects in or across Mule applications. Mule runtime engine (Mule) uses object stores to persist data for eventual retrieval. It can be used for disaster recovery but not for High Availability. Using object store can't guarantee that all instances won't go down at once. So not an appropriate choice.

NEW QUESTION 30

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

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

Answer: B

Explanation:

Types of logging:

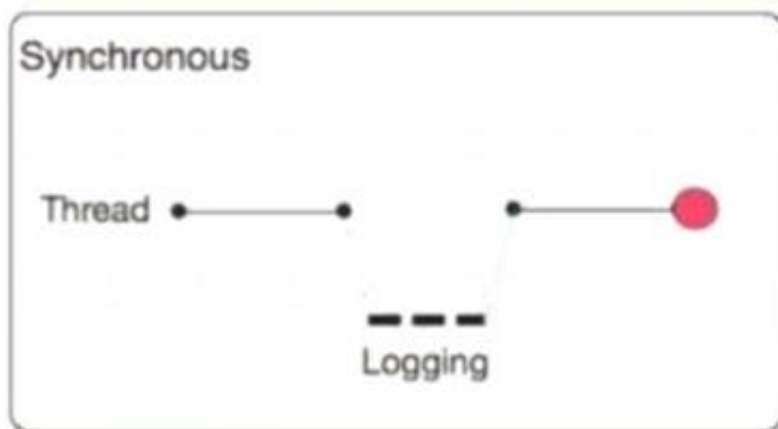
A) Synchronous: The execution of thread that is processing messages is interrupted to wait for the log message to be fully handled before it can continue. The execution of the thread that is processing your message is interrupted to wait for the log message to be fully output before it can continue

Performance degrades because of synchronous logging

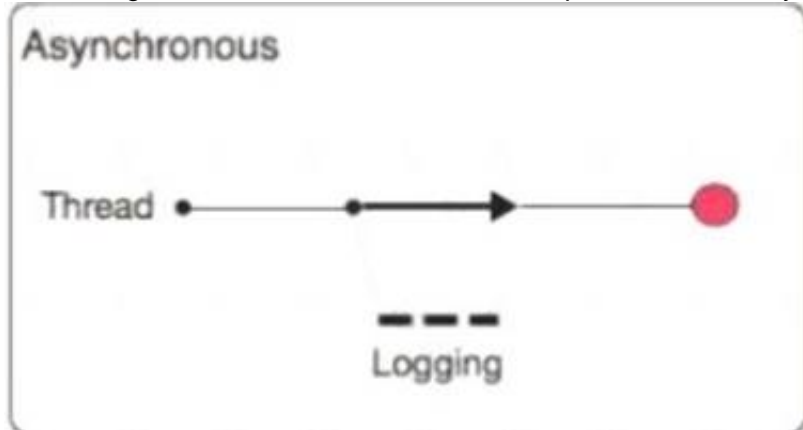
Used when the log is used as an audit trail or when logging ERROR/CRITICAL messages

If the logger fails to write to disk, the exception would raise on the same thread that's currently processing the Mule event. If logging is critical for you, then you can rollback the transaction.

Chart, diagram Description automatically generated



Chart, diagram, box and whisker chart Description automatically generated



B) Asynchronous:

The logging operation occurs in a separate thread, so the actual processing of your message won't be delayed to wait for the logging to complete

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

If the logger fails to write to disk, the thread doing the processing won't be aware of any issues writing to the disk, so you won't be able to rollback anything.

Because the actual writing of the log gets deferred, there's a chance that log messages might never make it to disk and get lost, if Mule were to crash before the buffers are flushed.

So Correct answer is: Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event

NEW QUESTION 31

An organization uses a set of customer-hosted Mule runtimes that are managed using the Mulesoft-hosted control plane. What is a condition that can be alerted on from Anypoint Runtime Manager without any custom components or custom coding?

- A. When a Mule runtime on a given customer-hosted server is experiencing high memory consumption during certain periods
- B. When an SSL certificate used by one of the deployed Mule applications is about to expire
- C. When the Mule runtime license installed on a Mule runtime is about to expire
- D. When a Mule runtime's customer-hosted server is about to run out of disk space

Answer: A

Explanation:

Correct answer is When a Mule runtime on a given customer-hosted server is experiencing high memory consumption during certain periods Using Anypoint Monitoring, you can configure two different types of alerts: Basic alerts for servers and Mule apps Limit per organization: Up to 50 basic alerts for users who do not have a Titanium subscription to Anypoint Platform You can set up basic alerts to trigger email notifications when a metric you are measuring passes a specified threshold. You can create basic alerts for the following metrics for servers or Mule apps: For on-premises servers and CloudHub apps: * CPU utilization * Memory utilization * Thread count Advanced alerts for graphs in custom dashboards in Anypoint Monitoring. You must have a Titanium subscription to use this feature. Limit per organization: Up to 20 advanced alerts

NEW QUESTION 32

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

- A. Compile, package, unit test, deploy, create associated API instances in API ManagerB Import from API designer, compile, package, unit test, deploy, publish to Anypoint Exchange
- B. Compile, package, unit test, validate unit test coverage, deploy
- C. Compile, package, unit test, deploy, integration test

Answer: C

NEW QUESTION 37

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

Explanation:

As the question says that data needs to be sent across multiple sftp servers, we cannot use non-repeatable streams. The non-repeatable strategy disables repeatable streams, which enables you to read an input stream only once.

You cannot use in-memory storage because with 0.2 vcore you will get only 1 GB of heap memory. Hence application will error out for file more than 1 GB.

Hence the correct option is file-based repeatable stream

NEW QUESTION 41

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 approaches would you 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

Explanation:

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

NEW QUESTION 44

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

Explanation:

Anypoint Platform provides a scripting and command-line tool for both Anypoint Platform and Anypoint Platform Private Cloud Edition (Anypoint Platform PCE).

The command-line interface (CLI) supports both the interactive shell and standard CLI modes and works with: Anypoint Exchange Access management Anypoint Runtime Manager

NEW QUESTION 48

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

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

Answer: B

NEW QUESTION 53

A company is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTP5 POST and must be acknowledged immediately.

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

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

Answer: B

NEW QUESTION 58

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

Explanation:

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

NEW QUESTION 60

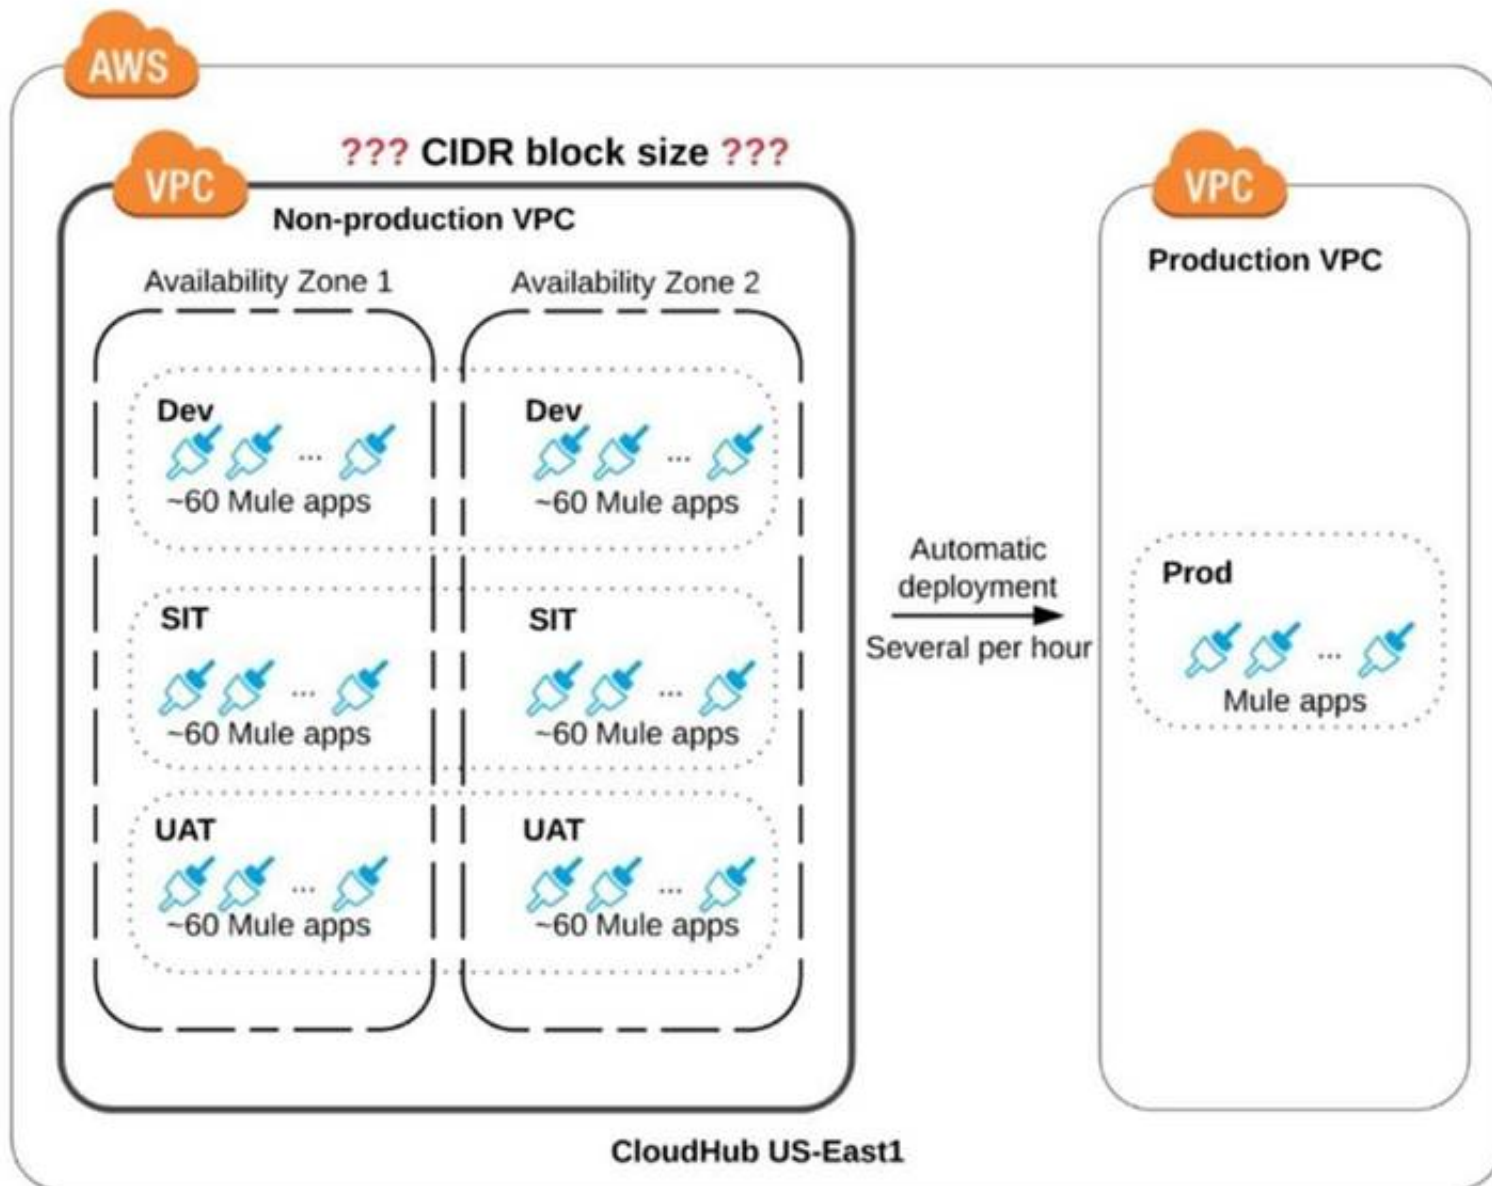
As a part of project requirement, Java Invoke static connector in a mule 4 application needs to invoke a static method in a dependency jar file. What are two ways to add the dependency to be visible by the connectors class loader?
(Choose two answers)

- 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: BD

NEW QUESTION 64

Refer to the exhibit.



An organization is sizing an Anypoint VPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization currently has three non-production environments (DEV, SIT and UAT) that share this VPC. The AWS region of the VPC has two AZs.

The organization has a very mature DevOps approach which automatically progresses each application through all non-production environments before automatically deploying to production. This process results in several Mule application deployments per hour, using CloudHub's normal zero-downtime deployment feature.

What is a CIDR block for this VPC that results in the smallest usable private IP address range?

- A. 10.0.0.0/26 (64 IPs)
- B. 10.0.0.0/25 (128 IPs)
- C. 10.0.0.0/24 (256 IPs)
- D. 10.0.0.0/22 (1024 IPs)

Answer: D

Explanation:

Mule applications are deployed in CloudHub workers and each worker is assigned with a dedicated IP • For zero downtime deployment, each worker in CloudHub needs additional IP addresses • A few IPs in a VPC are reserved for infrastructure (generally 2 IPs) • The IP addresses are usually in a private range with a subnet block specifier, such as 10.0.0.1/24 • The smallest CIDR network subnet block you can assign for your VPC is /24 (256 IP addresses) (60*3 env * 2 worker per application) + 50% of (total) for zero downtime = 540 In this case correct answer is 10.0.0.0/22 as this provided 1024 IP's . Other IP's are insufficient.

NEW QUESTION 67

An insurance company is using a CloudHub runtime plane. As a part of requirement, email alert should be sent to internal operations team every time of policy applied to an API instance is deleted As an integration architect suggest on how this requirement be met?

- A. Use audit logs in Anypoint platform to detect a policy deletion and configure the Audit logs alert feature to send an email to the operations team
- B. Use Anypoint monitoring to configure an alert that sends an email to the operations team every time a policy is deleted in API manager
- C. Create a custom connector to be triggered every time of policy is deleted in API manager
- D. Implement a new application that uses the Audit log REST API to detect the policy deletion and send an email to operations team the SMTP connector

Answer: D

NEW QUESTION 70

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

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

Answer: C

NEW QUESTION 74

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

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

Answer: C

NEW QUESTION 77

A company is planning to extend its Mule APIs to the Europe region. Currently all new applications are deployed to Cloudhub in the US region following this naming convention

{API name}-{environment}. for example, Orders-SAPI-dev, Orders-SAPI-prod etc.

Considering there is no network restriction to block communications between API's, what strategy should be implemented in order to apply the same new API's running in the EU region of CloudHub as well to minimize latency between API's and target users and systems in Europe?

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

Answer: D

NEW QUESTION 79

A leading bank implementing new mule API.

The purpose of API to fetch the customer account balances from the backend application and display them on the online platform the online banking platform. The online banking platform will send an array of accounts to Mule API get the account balances.

As a part of the processing the Mule API needs to insert the data into the database for auditing purposes and this process should not have any performance related implications on the account balance retrieval flow

How should this requirement be implemented to achieve better throughput?

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

Answer: D

NEW QUESTION 80

A Mule application name Pub uses a persistence object store. The Pub Mule application is deployed to Cloudhub and it configured to use Object Store v2.

Another Mule application name sub is being developed to retrieve values from the Pub Mule application persistence object Store and will also be deployed to cloudhub.

What is the most direct way for the Sub Mule application to retrieve values from the Pub Mule application persistence object store with the least latency?

- A. Use an object store connector configured to access the Pub Mule application persistence object store
- B. Use a VM connector configured to directly access the persistence queue of the Pub Mule application persistence object store.
- C. Use an Anypoint MQ connector configured to directly access the Pub Mule application persistence object store
- D. Use the Object store v2 REST API configured to access the Pub Mule application persistence object store.

Answer: D

Explanation:

* The Object Store V2 API enables API access to Anypoint Platform Object Store v2.

* You can configure a Mule app to use the Object Store REST API to store and retrieve values from an object store in another Mule app. However, Object Store v2 is not designed for app-to-app communication. To share data between two Mule4 apps, use a queue in Anypoint MQ.

* The Object Store v2 APIs enable you to use REST to perform the following:

- Retrieve a list of object stores and keys associated with an application.
- Store and retrieve key-value pairs in an object store.
- Delete key-value pairs from an object store.
- Retrieve Object Store usage statistics for your organization.
- Object Store provides these APIs: Object Store API

Object Store Stats API

NEW QUESTION 82

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 recordsEach Batch Step scope is invoked with one record in the payload of the received Mule event For each Batch Step scope, all 25 records within a block

are processed in parallel. All the records in a block must be completed before the block of 25 records is available to the next Batch Step scope.

B. The Batch Job scope processes each record block sequentially, one at a time. Each Batch Step scope is invoked with one record in the payload of the received Mule event. For each Batch Step scope, all 25 records within a block are processed sequentially, one at a time. All 4000 records must be completed before the blocks of records are available to the next Batch Step scope.

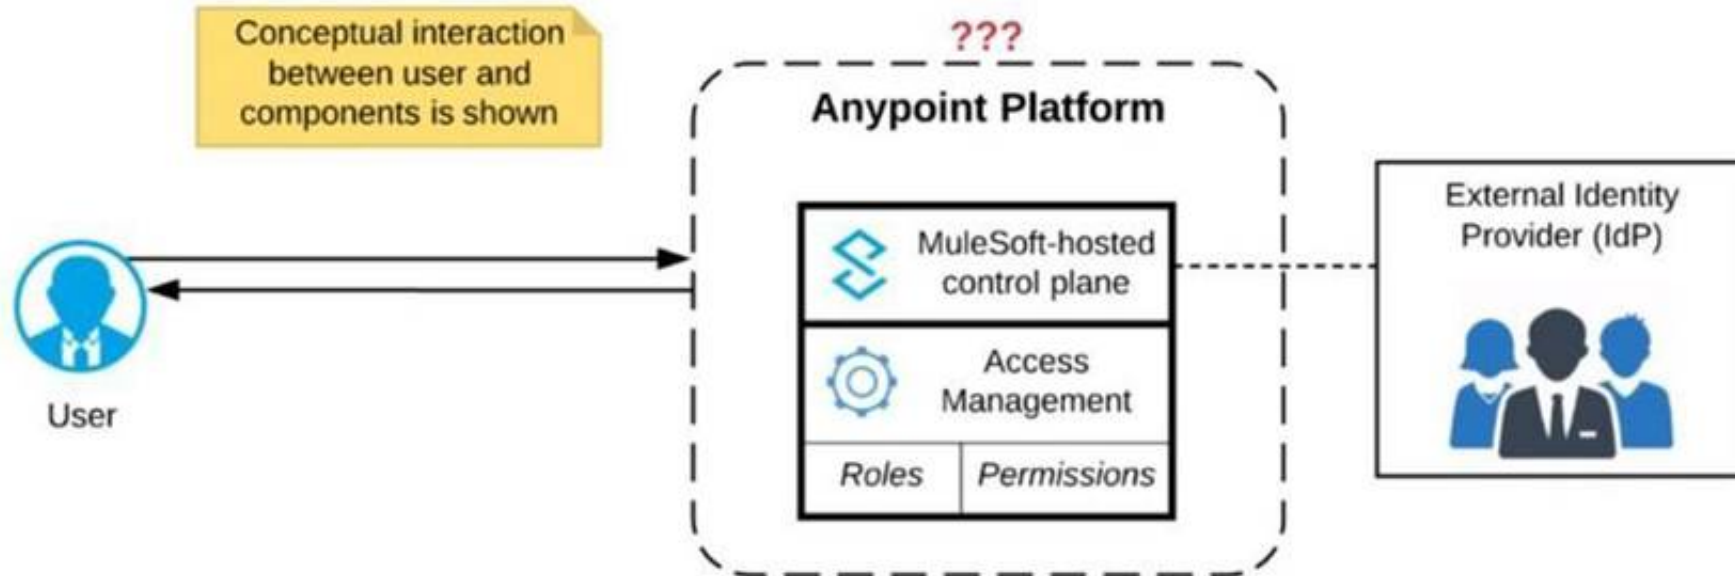
C. The Batch Job scope processes multiple record blocks in parallel, and a block of 25 records can jump ahead to the next Batch Step scope over an earlier block of records. Each Batch Step scope is invoked with one record in the payload of the received Mule event. For each Batch Step scope, all 25 records within a block are processed sequentially, one record at a time. All the records in a block must be completed before the block of 25 records is available to the next Batch Step scope.

D. The Batch Job scope processes multiple record blocks in parallel. Each Batch Step scope is invoked with a batch of 25 records in the payload of the received Mule event. For each Batch Step scope, all 4000 records are processed in parallel. Individual records can jump ahead to the next Batch Step scope before the rest of the records finish processing in the current Batch Step scope.

Answer: A

NEW QUESTION 86

Refer to the exhibit.



Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.

What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled via the external Identity Provider?

- A. Controlling the business group within Anypoint Platform to which the user belongs
- B. Assigning Anypoint Platform permissions to a role
- C. Assigning Anypoint Platform role(s) to a user
- D. Removing a user's access to Anypoint Platform when they no longer work for the organization

Answer: B

Explanation:

* By default, Anypoint Platform performs its own user management

– For user management, one external IdP can be integrated with the Anypoint Platform organization (note: not at business group level)

– Permissions and access control are still enforced inside Anypoint Platform and CANNOT be controlled via the external Identity Provider * As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). * You can map users in a federated organization's group to a role which also gives the flexibility of controlling the business group within Anypoint Platform to which the user belongs to.

Also user can nbe removed from external identity

management system when they no longer work for the organization. So they wont be able to authenticate using SSO to login to Anypoint Platform. * Using external identity we can no change permissions of a particular role in Mulesoft Anypoint platform.

* So Correct answer is Assigning Anypoint Platform permissions to a role

NEW QUESTION 88

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

Explanation:

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

NEW QUESTION 92

An integration Mute application consumes and processes a list of rows from a CSV file. Each row must be read from the CSV file, validated, and the row data sent to a JMS queue, in the exact order as in the CSV file.

If any processing step for a row falls, then a log entry must be written for that row, but processing of other rows must not be affected.

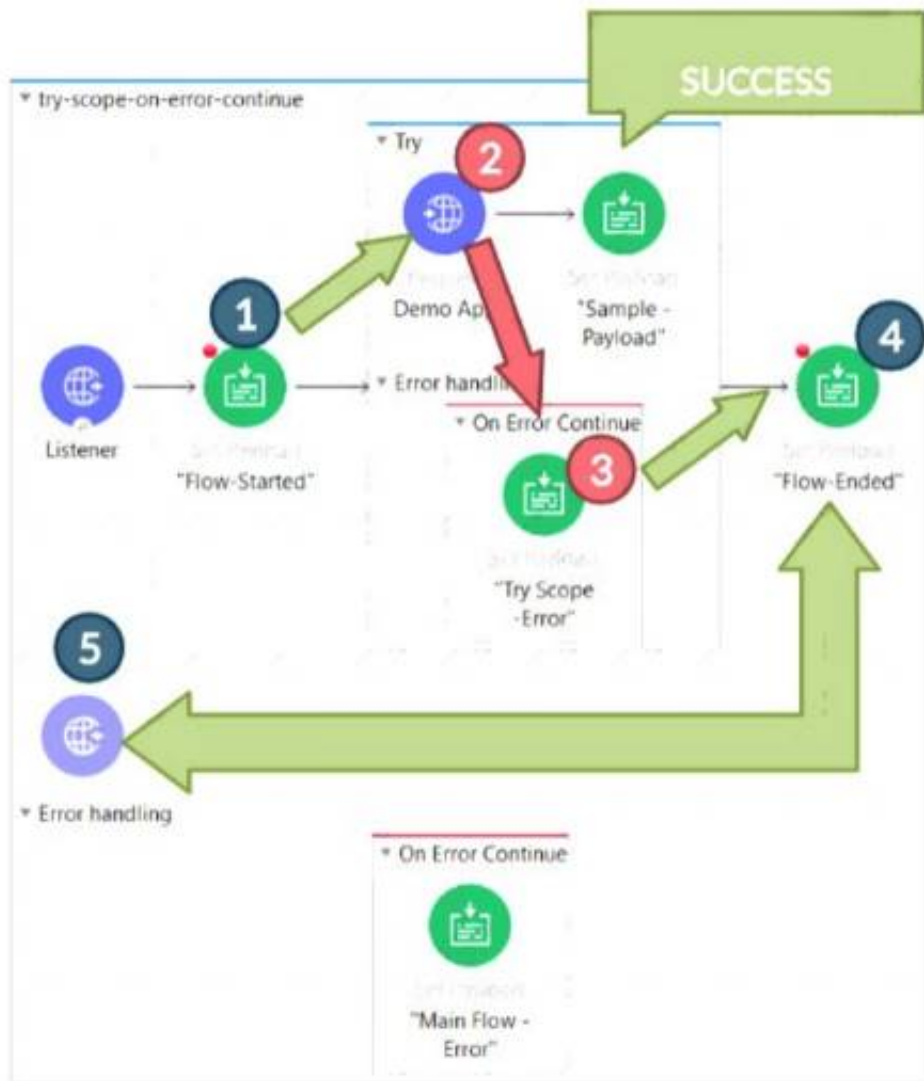
What combination of Mute components is most idiomatic (used according to their intended purpose) when Implementing the above requirements?

- A. Scatter-Gather component On Error Continue scope
- B. VM connector first Successful scope On Error Propagate scope
- C. For Each scope On Error Continue scope
- D. Async scope On Error Propagate scope

Answer: C

Explanation:

- * On Error Propagate halts execution and sends error to the client. In this scenario it's mentioned that "processing of other rows must not be affected" so Option B and C are ruled out.
- * Scatter gather is used to club multiple responses together before processing. In this scenario, we need sequential processing. So option A is out of choice.
- * Correct answer is For Each scope & On Error Continue scope Below requirement can be fulfilled in the below way
- 1) Using For Each scope , which will send each row from csv file sequentially. each row needs to be sent sequentially as requirement is to send the message in exactly the same way as it is mentioned in the csv file
- 2) Also other part of requirement is if any processing step for a row fails then it should log an error but should not affect other record processing . This can be achieved using On error Continue scope on these set of activities. so that error will not halt the processing. Also logger needs to be added in error handling section so that it can be logged.
- * Attaching diagram for reference. Here it's try scope, but similar would be the case with For Each loop. Diagram Description automatically generated



NEW QUESTION 97

An Order microservice and a Fulfillment microservice are being designed to communicate with their clients through message-based integration (and NOT through API invocations).

The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mule application, the Fulfillment microservice.

The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message (a kind of event message). Each OrderFulfilled message can be consumed by any interested Mule application, and the Order microservice is one such Mule application.

What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

- A. Order messages are sent to an Anypoint MQ exchange OrderFulfilled messages are sent to an Anypoint MQ queue Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices
- B. Order messages are sent to a JMS queue
- C. OrderFulfilled messages are sent to a JMS topic Both microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices
- D. Order messages are sent directly to the Fulfillment microservice
- E. OrderFulfilled messages are sent directly to the Order microservice The Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load of each microservice
- F. Order messages are sent to a JMS queue
- G. OrderFulfilled messages are sent to a JMS topic The Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice

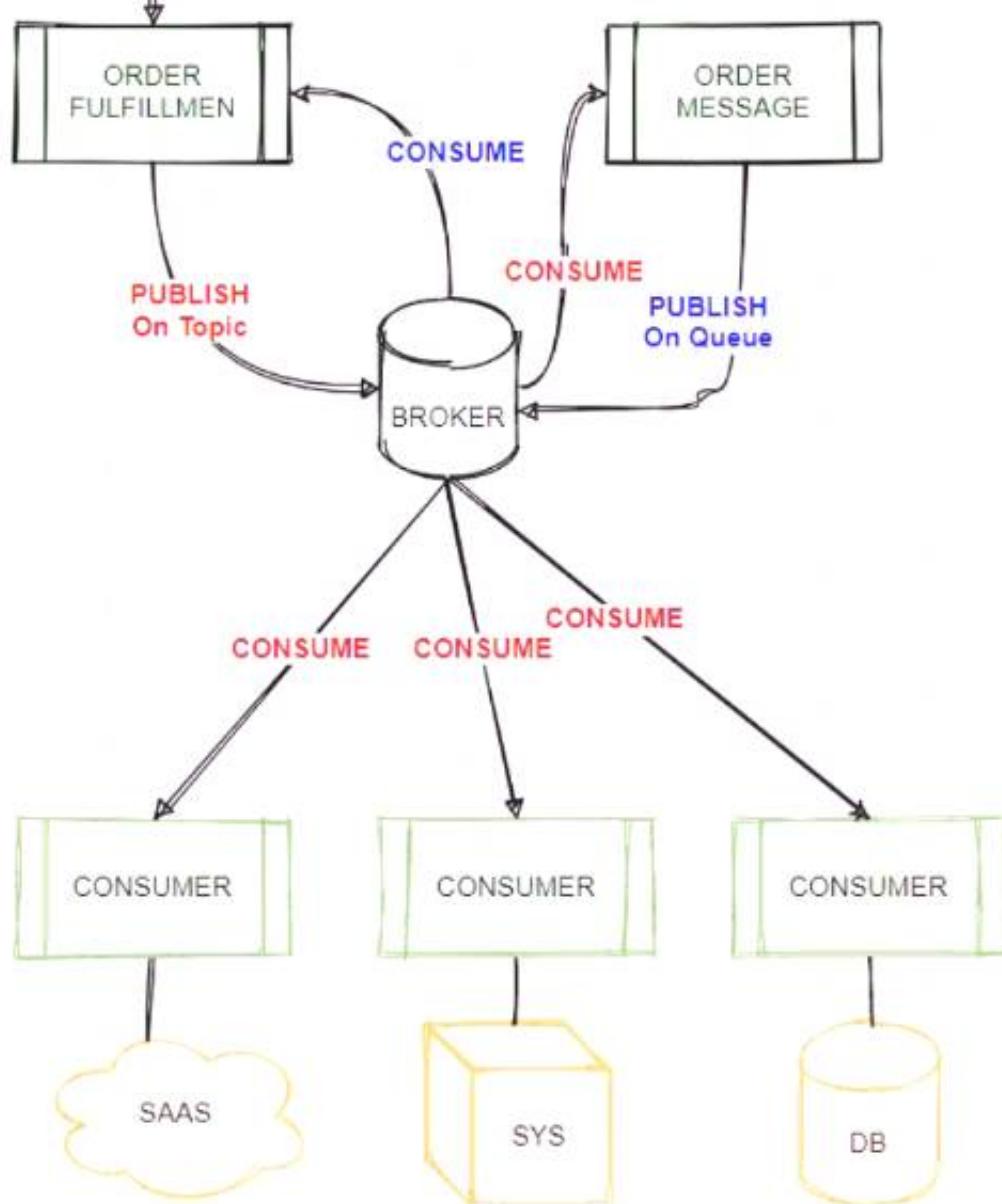
Answer: B

Explanation:

- * If you need to scale a JMS provider/ message broker, - add nodes to scale it horizontally or - add memory to scale it vertically
- * Cons of adding another JMS provider/ message broker: - adds cost. - adds complexity to use two JMS brokers - adds Operational overhead if we use two brokers, say, ActiveMQ and IBM MQ
- * So Two options that mention to use two brokers are not best choice. * It's mentioned that "The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message. Each OrderFulfilled message can be consumed by any interested Mule application." - When you publish a message on a topic, it goes to all the subscribers who are interested - so zero to many subscribers will receive a copy of the message. - When

you send a message on a queue, it will be received by exactly one consumer. * As we need multiple consumers to consume the message below option is not valid choice: "Order messages are sent to an Anypoint MQ exchange. OrderFulfilled messages are sent to an Anypoint MQ queue. Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices" * Order messages are only consumed by one Mule application, the Fulfillment microservice, so we will publish it on queue and OrderFulfilled message can be consumed by any interested Mule application so it need to be published on Topic using same broker. * Correct Answer Best choice in this scenario is: "Order messages are sent to a JMS queue. OrderFulfilled messages are sent to a JMS topic. Both microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices" Tried to depict scenario in diagram:

Diagram Description automatically generated



NEW QUESTION 102

An organization is struggling frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects. How can these best practices be achieved with the applications having the least amount of responsibility?

- A. Create a Mule plugin project with all the dependencies and add it as a dependency in each application's POM.xml file
- B. Create a mule domain project with all the dependencies define in its POM.xml file and add each application to the domain Project
- C. Add all dependencies in each application's POM.xml file
- D. Create a parent POM of all the required dependencies and reference each in each application's POM.xml file

Answer: D

NEW QUESTION 103

A mule application is deployed to a Single Cloudhub worker and the public URL appears in Runtime Manager as the APP URL. Requests are sent by external web clients over the public internet to the mule application App url. Each of these requests routed to the HTTPS Listener event source of the running Mule application. Later, the DevOps team edits some properties of this running Mule application in Runtime Manager. Immediately after the new property values are applied in runtime manager, how is the current Mule application deployment affected and how will future web client requests to the Mule application be handled?

- A. Cloudhub will redeploy the Mule application to the OLD Cloudhub workerNew web client requests will RETURN AN ERROR until the Mule application is redeployed to the OLD Cloudhub worker
- B. CloudHub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests will RETURN AN ERROR until the NEW Cloudhub worker is available
- C. Cloudhub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker until the NEW Cloudhub worker is available.
- D. Cloudhub will redeploy the mule application to the OLD Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.

Answer: C

Explanation:

CloudHub supports updating your applications at runtime so end users of your HTTP APIs experience zero downtime. While your application update is deploying, CloudHub keeps the old version of your application running. Your domain points to the old version of your application until the newly uploaded version is fully started. This allows you to keep servicing requests from your old application while the new version of your application is starting.

NEW QUESTION 108

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

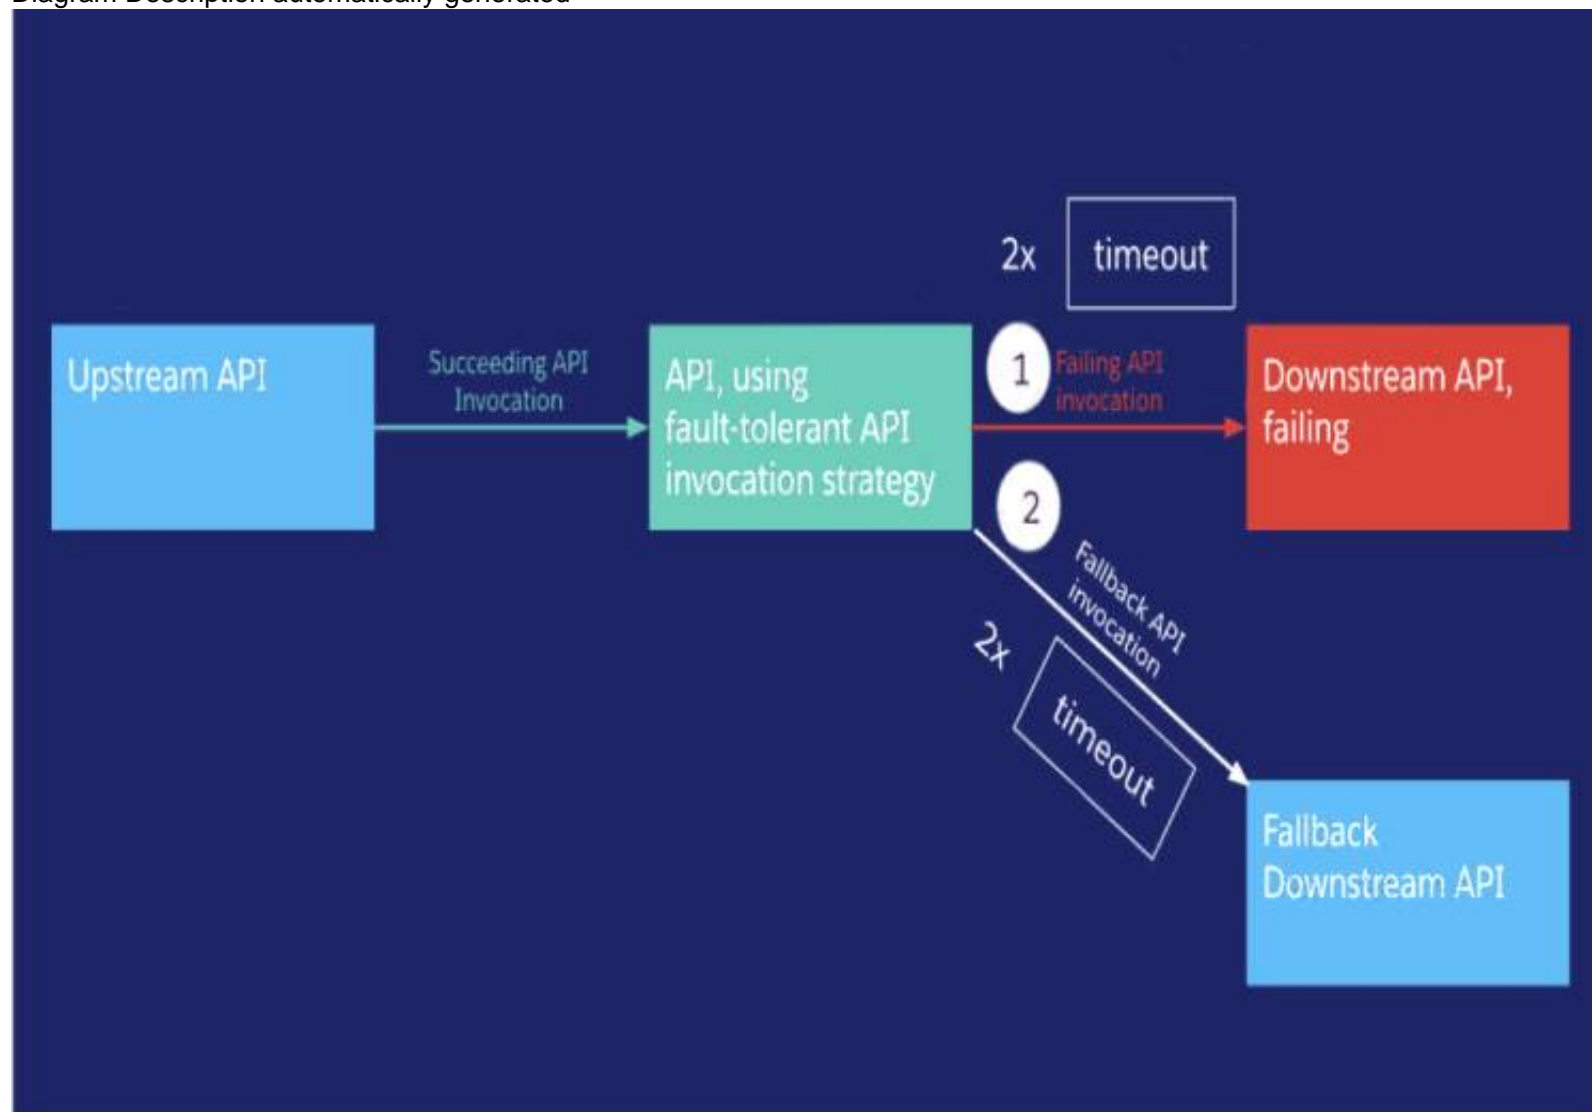
- A. Redirect client requests through an HTTP 303 temporary redirect status code to the fallback API whenever the Order API is unavailable
- B. Set an option in the HTTP Requester component that invokes the order API to instead invoke a fallback API whenever an HTTP 4XX or 5XX response status code is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

Answer: A

Explanation:

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

Diagram Description automatically generated



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

NEW QUESTION 110

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

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

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 112

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 tosupport at Mule application deployments.
- D. Keep the default logging configuration in Runtime ManagerDefine the Splunk appender in EACH Mule application log4j2.xml file

Answer: B

Explanation:

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

NEW QUESTION 115

A stock broking company makes use of CloudHub VPC to deploy Mule applications. Mule application needs to connect to a database application in the customers on-premises corporate data center and also to a Kafka cluster running in AWS VPC.

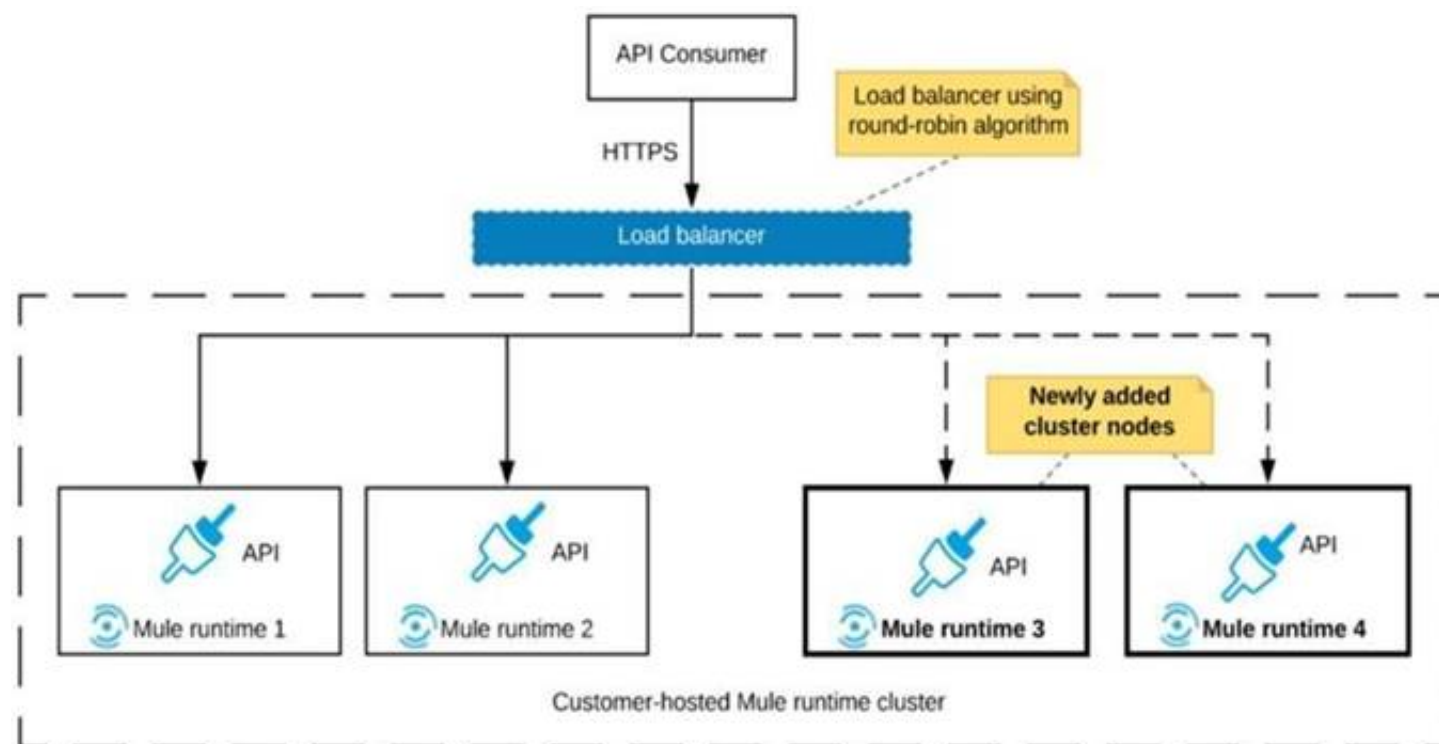
How is access enabled for the API to connect to the database application and Kafka cluster securely?

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

Answer: B

NEW QUESTION 120

Refer to the exhibit.



An organization uses a 2-node Mule runtime cluster to host one stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution.

Two additional nodes have been added to the cluster and the load balancer has been configured to recognize the new nodes with no other change to the load balancer.

What average performance change is guaranteed to happen, assuming all cluster nodes are fully operational?

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

Answer: D

NEW QUESTION 121

An API client is implemented as a Mule application that includes an HTTP Request operation using a default configuration. The HTTP Request operation invokes an external API that follows standard HTTP status code conventions, which causes the HTTP Request operation to return a 4xx status code. What is a possible cause of this status code response?

- A. An error occurred inside the external API implementation when processing the HTTP request that was received from the outbound HTTP Request operation of the Mule application
- B. The external API reported that the API implementation has moved to a different external endpoint
- C. The HTTP response cannot be interpreted by the HTTP Request operation of the Mule application after it was received from the external API
- D. The external API reported an error with the HTTP request that was received from the outbound HTTP Request operation of the Mule application

Answer: D

Explanation:

Correct choice is: "The external API reported an error with the HTTP request that was received from the outbound HTTP Request operation of the Mule application"

Understanding HTTP 4XX Client Error Response Codes : A 4XX Error is an error that arises in cases where there is a problem with the user's request, and not with the server.

Such cases usually arise when a user's access to a webpage is restricted, the user misspells the URL, or when a webpage is nonexistent or removed from the public's view.

In short, it is an error that occurs because of a mismatch between what a user is trying to access, and its availability to the user — either because the user does not have the right to access it, or because what the user is trying to access simply does not exist. Some of the examples of 4XX errors are

400 Bad Request The server could not understand the request due to invalid syntax. 401 Unauthorized Although the HTTP standard specifies "unauthorized", semantically this response means "unauthenticated". That is, the client must authenticate itself to get the requested response. 403 Forbidden The client does not have access rights to the content; that is, it is unauthorized, so the server is refusing to give the requested resource. Unlike 401, the client's identity is known to the server. 404 Not Found The server can not find the requested resource. In the browser, this means the URL is not recognized. In an API, this can also mean that the endpoint is valid but the resource itself does not exist. Servers may also send this response instead of 403 to hide the existence of a resource from an unauthorized client. This response code is probably the most famous one due to its frequent occurrence on the web. 405 Method Not Allowed The request method is known by the server but has been disabled and cannot be used. For example, an API may forbid DELETE-ing a resource. The two mandatory methods, GET and HEAD, must never be disabled and should not return this error code. 406 Not Acceptable This response is sent when the web server, after performing server-driven content negotiation, doesn't find any content that conforms to the criteria given by the user agent. The external API reported that the API implementation has moved to a different external endpoint cannot be the correct answer as in this situation 301 Moved Permanently The URL of the requested resource has been changed permanently. The new URL is given in the response.

-----In Lay man's term the scenario would be: API CLIENT —> MuleSoft API - HTTP request "Hey, API.. process this" —> External API API CLIENT <- MuleSoft API - http response "I'm sorry Client.. something is wrong with that request" <- (4XX) External API

NEW QUESTION 125

A project team uses RAML specifications to document API functional requirements and deliver API definitions. As per the current legal requirement, all designed API definitions to be augmented with an additional non-functional requirement to protect the services from a high rate of requests according to define service level agreements.

Assuming that the project is following Mulesoft API governance and policies, how should the project team convey the necessary non-functional requirement to stakeholders?

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

Answer: D

NEW QUESTION 130

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

The organization wants to monitor these APIs by running health checks: for example, to determine if an API can properly accept and process requests. The organization does not have subscriptions to any external monitoring tools and also does not want to extend its IT footprint.

What Anypoint Platform feature provides the most idiomatic (used for its intended purpose) way to monitor the availability of both the Mule and the non-Mule API implementations?

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

Answer: D

NEW QUESTION 134

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

Explanation:

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

create WSDL.

NEW QUESTION 138

An organization is creating a set of new services that are critical for their business. The project team prefers using REST for all services but is willing to use SOAP with common WS-" standards if a particular service requires it.

What requirement would drive the team to use SOAP/WS-* for a particular service?

- A. Must use XML payloads for the service and ensure that it adheres to a specific schema
- B. Must publish and share the service specification (including data formats) with the consumers of the service
- C. Must support message acknowledgement and retry as part of the protocol
- D. Must secure the service, requiring all consumers to submit a valid SAML token

Answer: D

Explanation:

Security Assertion Markup Language (SAML) is an open standard that allows identity providers (IdP) to pass authorization credentials to service providers (SP). SAML transactions use Extensible Markup Language (XML) for standardized communications between the identity provider and service providers.

SAML is the link between the authentication of a user's identity and the authorization to use a service. WS-Security is the key extension that supports many authentication models including: basic

username/password credentials, SAML, OAuth and more.

A common way that SOAP API's are authenticated is via SAML Single Sign On (SSO). SAML works by facilitating the exchange of authentication and authorization credentials across applications. However, there is no specification that describes how to add SAML to REST web services.

Reference: <https://www.oasis-open.org/committees/download.php/16768/wss-v1.1-spec-os-SAMLSecurityTokenProfile.pdf>

NEW QUESTION 140

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

Explanation:

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

NEW QUESTION 144

What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

- A. To send Mule application log entries to Splunk
- B. To change log4j2 log levels in Anypoint Runtime Manager without having to restart the Mule application
- C. To log certain messages to a custom log category
- D. To directly reference one shared and customized log4j2.xml file from multiple Mule applications

Answer: D

Explanation:

* Correct answer is To directly reference one shared and customized log4j2.xml file from multiple Mule applications. Key word to note in the answer is directly.

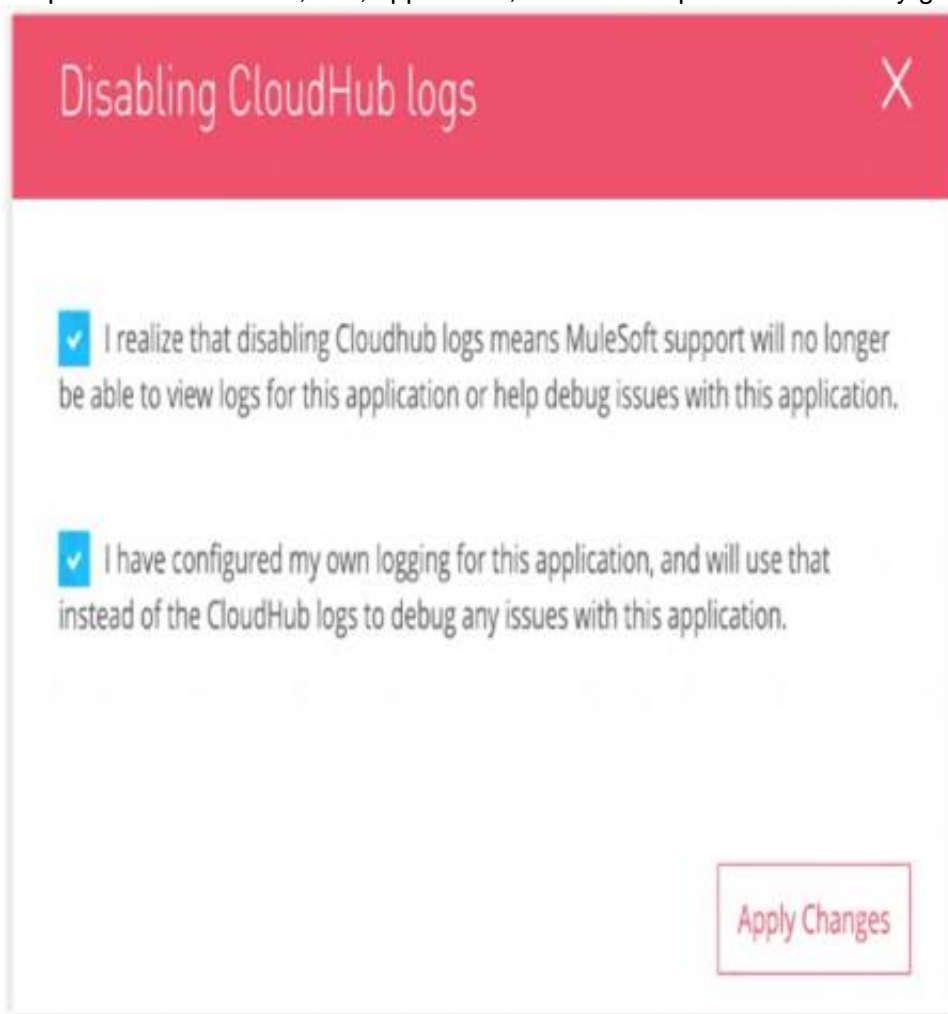
* By default, CloudHub replaces a Mule application's log4j2.xml file with a CloudHub log4j2.xml file. This specifies the CloudHub appender to write logs to the CloudHub logging service.

* You cannot modify CloudHub log4j2.xml file to add any custom appender. But there is a process in order to achieve this. You need to raise a request on support portal to disable CloudHub provided Mule application log4j2 file.

Graphical user interface, application, Word Description automatically generated



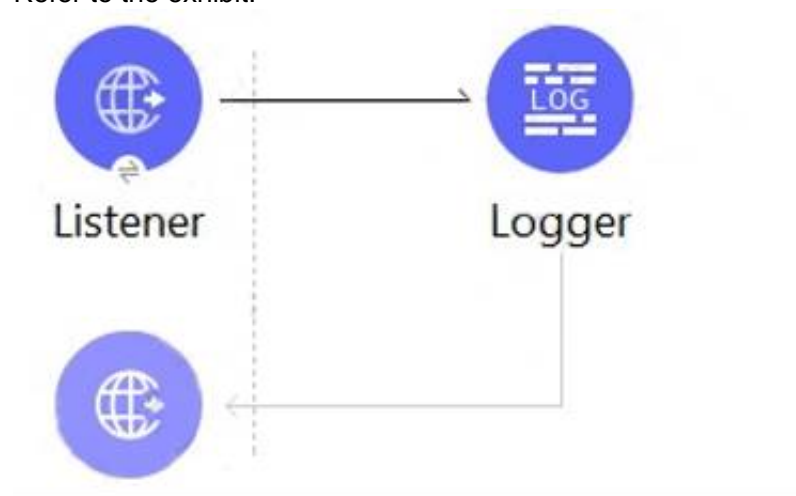
* Once this is done , Mule application's log4j2.xml file is used which you can use to send/export application logs to other log4j2 appenders, such as a custom logging system MuleSoft does not own any responsibility for lost logging data due to misconfiguration of your own log4j appender if it happens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



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

NEW QUESTION 147

Refer to the exhibit.



► Error handling

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 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 153

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 mocking a message before or inside a ForEach processor
- C. The recorder support loops where the structure of the data being 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 156

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

Explanation:

* Sharing Mule applications as templates is a great way to share your work with other people who are in your organization in Anypoint Platform. When they need to build a similar application they can create the mule application using the template project from Anypoint studio.

* Anypoint Templates are designed to make it easier and faster to go from a blank canvas to a production application. They're bit for bit Mule applications requiring only Anypoint Studio to build and design, and are deployable both on-premises and in the cloud.

* Anypoint Templates are based on five common data Integration patterns and can be customized and extended to fit your integration needs. So even if your use case involves different endpoints or connectors than those included in the template, they still offer a great starting point.

Some of the best practices while creating the template project: - Define the common error handler as part of template project, either using pom dependency or mule config file - Define common logger/audit framework as part of the template project - Define the env specific properties and secure properties file as per the requirement - Define global.xml for global configuration - Define the config file for connector configuration like Http,Salesforce,File,FTP etc - Create separate folders to create DWL,Properties,SSL certificates etc - Add the dependency and configure the pom.xml as per the business need - Configure the mule-artifact.json as per the business need

NEW QUESTION 159

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

Explanation:

* Organization should continue to use SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching.

* Reason is that Anypoint exchange is not full fledged version repositories like GitHub.

* But at same time it is tightly coupled with Mule assets

NEW QUESTION 163

What Is a recommended practice when designing an integration Mule 4 application that reads a large XML payload as a stream?

- A. The payload should be dealt with as a repeatable XML stream, which must only be traversed (iterated-over) once and CANNOT be accessed randomly from DataWeave expressions and scripts
- B. The payload should be dealt with as an XML stream, without converting it to a single Java object (POJO)
- C. The payload size should NOT exceed the maximum available heap memory of the Mule runtime on which the Mule application executes
- D. The payload must be cached using a Cache scope If It Is to be sent to multiple backend systems

Answer: C

Explanation:

If the size of the stream exceeds the maximum, a `STREAM_MAXIMUM_SIZE_EXCEEDED` error is raised.

NEW QUESTION 167

A company wants its users to log in to Anypoint Platform using the company's own internal user credentials. To achieve this, the company needs to integrate an external identity provider (IdP) with the company's Anypoint Platform master organization, but SAML 2.0 CANNOT be used. Besides SAML 2.0, what single-sign-on standard can the company use to integrate the IdP with their Anypoint Platform master organization?

- A. SAML 1.0
- B. OAuth 2.0
- C. Basic Authentication
- D. OpenID Connect

Answer: D

Explanation:

As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO).

Configure identity management using one of the following single sign-on standards:

- 1) OpenID Connect: End user identity verification by an authorization server including SSO
- 2) SAML 2.0: Web-based authorization including cross-domain SSO

NEW QUESTION 172

What is not true about Mule Domain Project?

- A. This allows Mule applications to share resources
- B. Expose multiple services within the Mule domain on the same port
- C. Only available Anypoint Runtime Fabric
- D. Send events (messages) to other Mule applications using VM queues

Answer: C

Explanation:

* Mule Domain Project is ONLY available for customer-hosted Mule runtimes, but not for Anypoint Runtime Fabric

* Mule domain project is available for Hybrid and Private Cloud (PCE). Rest all provide application isolation and can't support domain project.

What is Mule Domain Project?

* A Mule Domain Project is implemented to configure the resources that are shared among different projects. These resources can be used by all the projects associated with this domain. Mule applications can be associated with only one domain, but a domain can be associated with multiple projects. Shared resources allow multiple development teams to work in parallel using the same set of reusable connectors. Defining these connectors as shared resources at the domain level allows the team to:

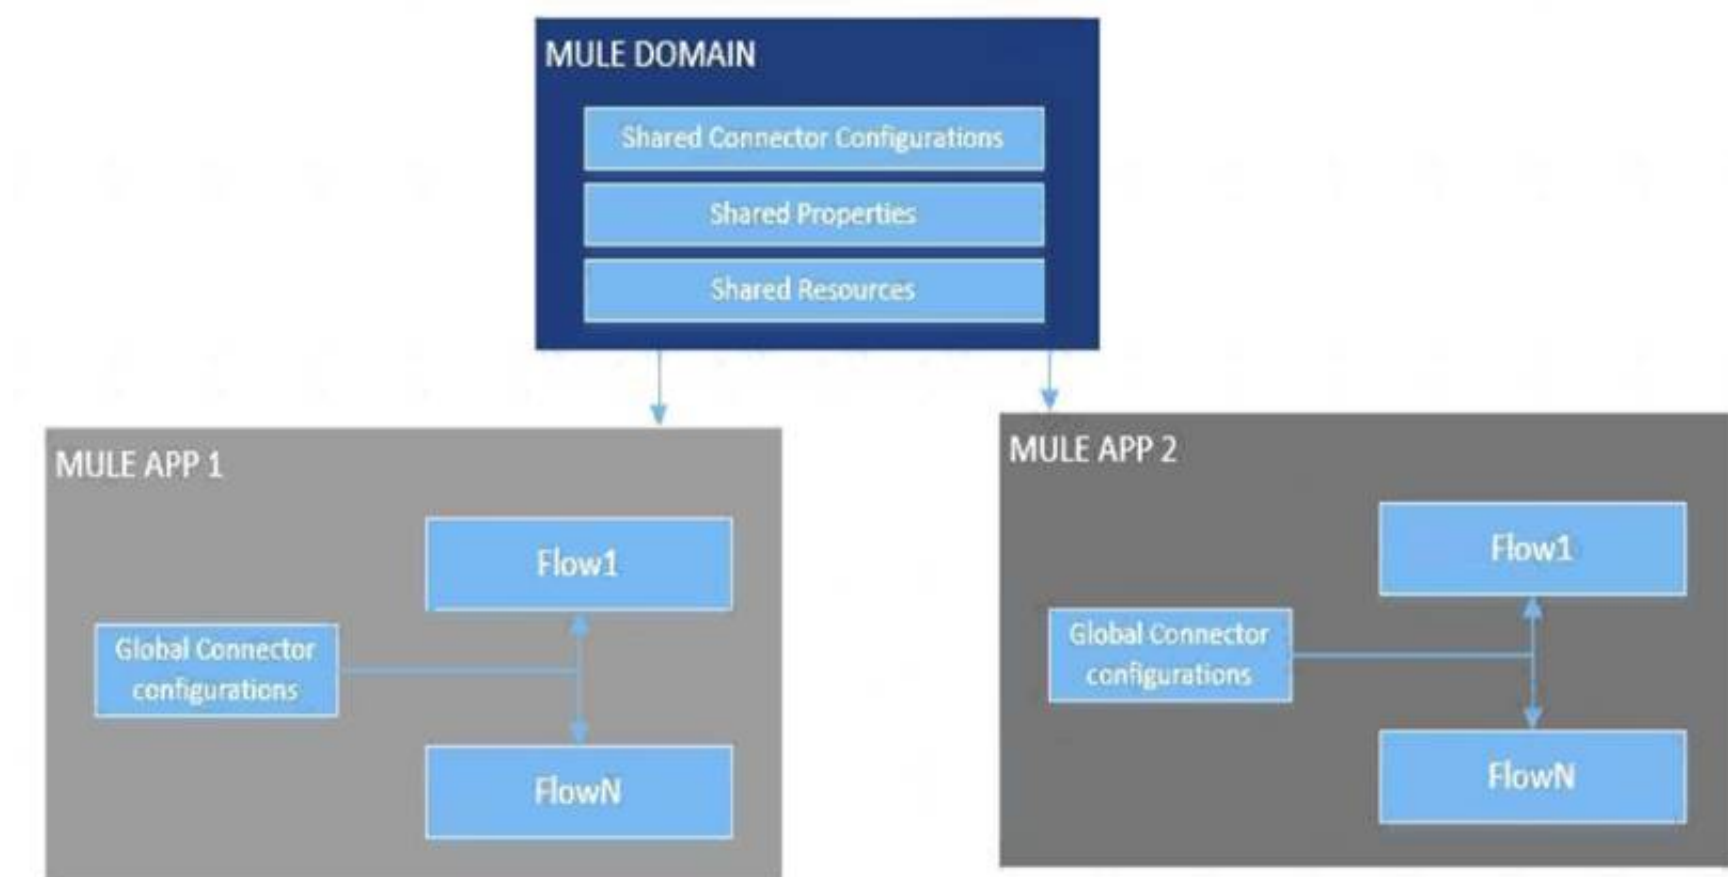
- Expose multiple services within the domain through the same port.
- Share the connection to persistent storage.
- Share services between apps through a well-defined interface.
- Ensure consistency between apps upon any changes because the configuration is only set in one place.

* Use domains Project to share the same host and port among multiple projects. You can declare the http connector within a domain project and associate the domain project with other projects. Doing this also allows to control thread settings, keystore configurations, time outs for all the requests made within multiple applications. You may think that one can also achieve this by duplicating the http connector configuration across all the applications. But, doing this may pose a nightmare if you have to make a change and redeploy all the applications.

* If you use connector configuration in the domain and let all the applications use the new domain instead of a default domain, you will maintain only one copy of the http connector configuration. Any changes will require only the domain to be redeployed instead of all the applications.

You can start using domains in only three steps:

- 1) Create a Mule Domain project
- 2) Create the global connector configurations which needs to be shared across the applications inside the Mule Domain project
- 3) Modify the value of domain in mule-deploy.properties file of the applications Graphical user interface Description automatically generated



NEW QUESTION 174

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

Explanation:

Correct answer is VM. VM supports Transactional Type. When an exception occurs, the transaction rolls back to its original state for reprocessing. This feature is not supported by other connectors.

Here is additional information about Transaction management: Table Description automatically generated

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

NEW QUESTION 177

A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.

The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.

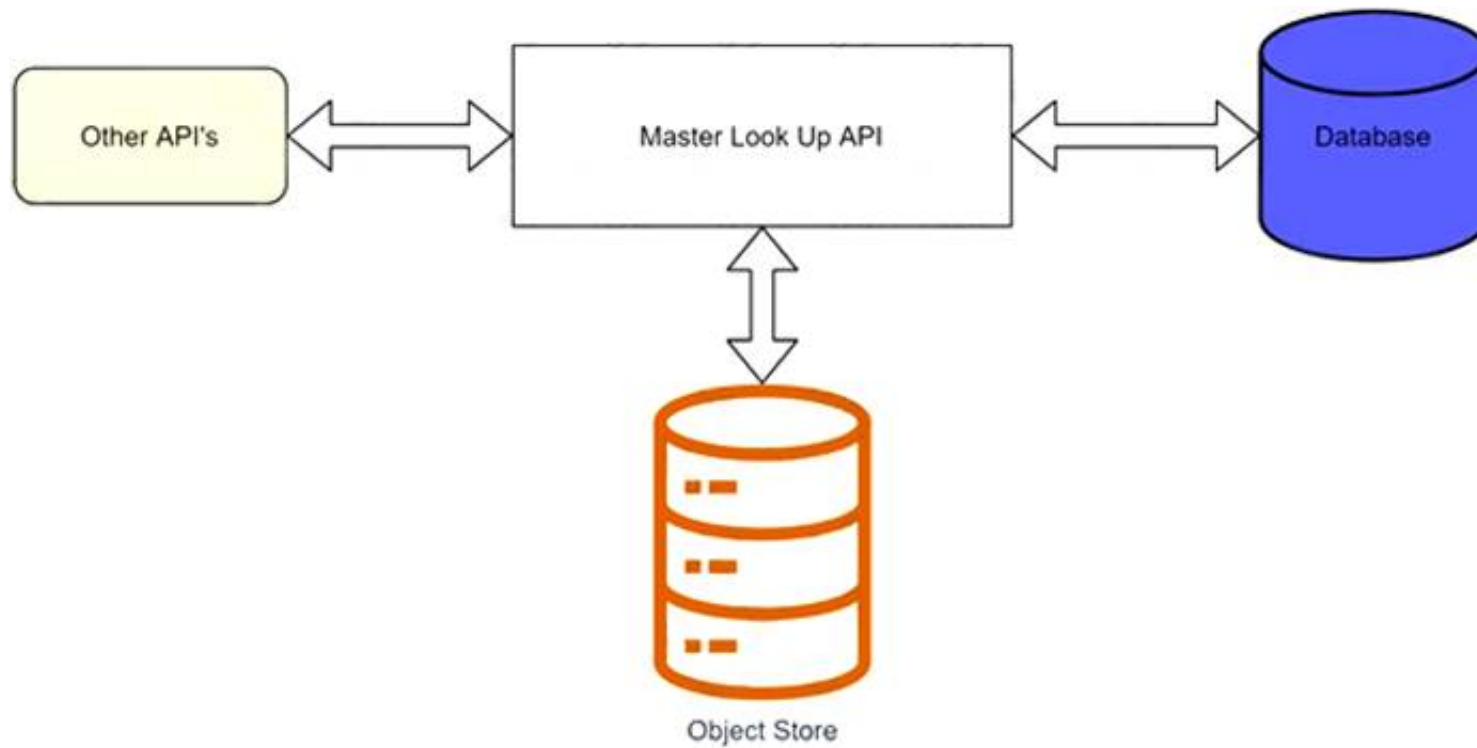
What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

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

Answer: D

NEW QUESTION 182

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 cloudbus 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 the 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 Cloudbus worker to provide additional capacity

Answer: A

NEW QUESTION 185

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 187

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

Explanation:

In Unit testing instead of using actual backends, stubs are used for the backend services. This ensures that developers are not blocked and have no dependency on other systems.

In Unit testing instead of using actual backends, stubs are used for the backend services. This ensures that developers are not blocked and have no dependency on other systems.

Below are the typical characteristics of unit testing.

- Unit tests do not require deployment into any special environment, such as a staging environment
- Unit tests can be run from within an embedded Mule runtime
- Unit tests can/should be implemented using MUnit
- For read-only interactions to any dependencies (such as other APIs): allowed to invoke production endpoints
- For write interactions: developers must implement mocks using MUnit
- Require knowledge of the implementation details of the API implementation under test

NEW QUESTION 189

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

MCIA-Level-1 Practice Exam Features:

- * MCIA-Level-1 Questions and Answers Updated Frequently
- * MCIA-Level-1 Practice Questions Verified by Expert Senior Certified Staff
- * MCIA-Level-1 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * MCIA-Level-1 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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