Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1

https://www.2passeasy.com/dumps/MCIA-Level-1/

NEW QUESTION 1
Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improved throughput of the system. But it was observed over the period of time that few of the important exception log messages which were used to rollback transactions are not working as expected causing huge loss to the Organization. Organization wants to avoid these losses. Application also has constraints due to which they cant compromise on throughput much. What is the possible option in this case?

A. Logging needs to be changed from asynchronous to synchronous
B. External log appender needs to be used in this case
C. Persistent memory storage should be used in such scenarios
D. Mixed configuration of asynchronous or synchronous loggers should be used to log exceptions via synchronous way

Answer: D

Explanation:
Correct approach is to use Mixed configuration of asynchronous or synchronous loggers shoud be used to log exceptions via synchronous way Asynchronous logging poses a performance-reliability trade-off. You may lose some messages if Mule crashes before the logging buffers flush to the disk. In this case, consider that you can have a mixed configuration of asynchronous or synchronous loggers in your app. Best practice is to use asynchronous logging over synchronous with a minimum logging level of WARN for a production application. In some cases, enable INFO logging level when you need to confirm events such as successful policy installation or to perform troubleshooting. Configure your logging strategy by editing your application's src/main/resources/log4j2.xml file

NEW QUESTION 2
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 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 3
As a part of design , Mule application is required call the Google Maps API to perform a distance computation. The application is deployed to cloudhub.
At the minimum what should be configured in the TLS context of the HTTP request configuration to meet these requirements?

A. The configuration is built-in and nothing extra is required for the TLS context
B. Request a private key from Google and create a PKCS12 file with it and add it in keyStore as a part of TLS context
C. Download the Google public certificate from a browser, generate JKS file from it and add it in key store as a part of TLS context
D. Download the Google public certificate from a browser, generate a JKS file from it and add it in Truststore as part of the TLS context

Answer: A

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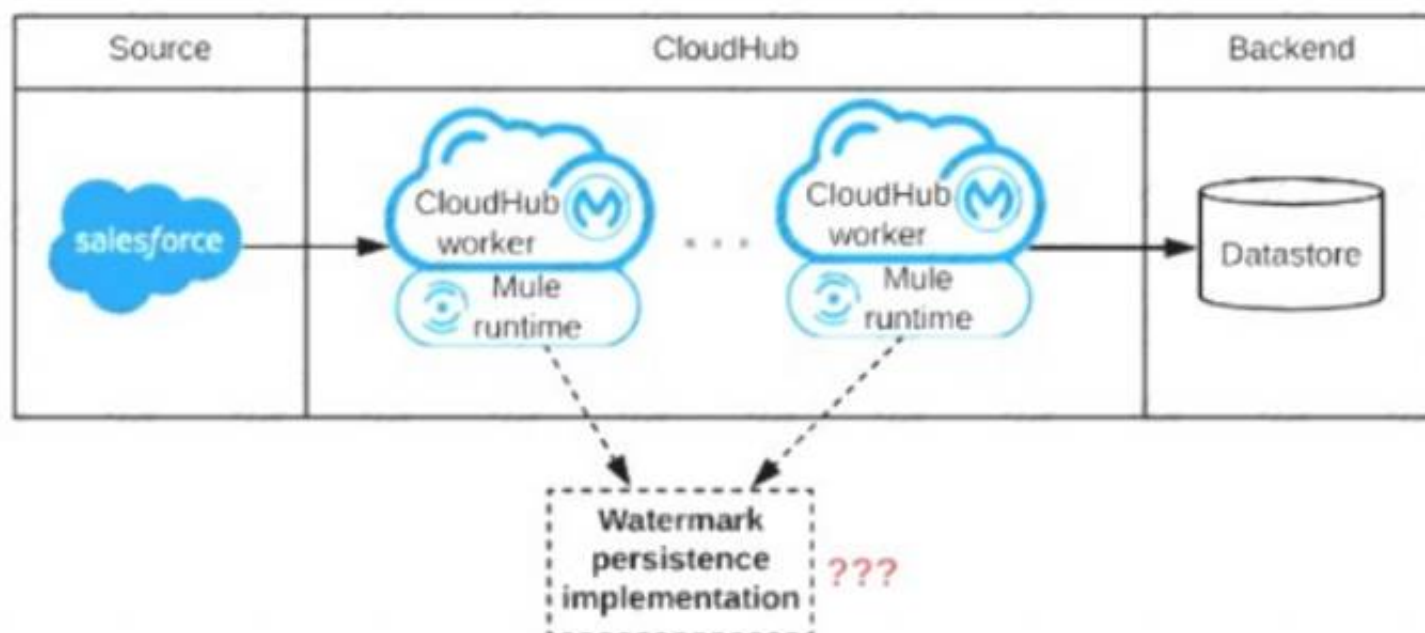Which method should be used to meet the company's requirement for its system?

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

Answer: C

NEW QUESTION 6
Refer to the exhibit.

A Mule application is being designed to be deployed to several CloudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.
A watermark will be used to only retrieve those Satesforce Accounts that have been modified since the last time the integration logic ran.
What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?

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

Answer: B

Explanation:
* An object store is a facility for storing objects in or across Mule applications. Mule uses object stores to persist data for eventual retrieval.
* Mule provides two types of object stores:
1) In-memory store – stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime.
2) Persistent store – Mule persists data when an object store is explicitly configured to be persistent.
In a standalone Mule runtime, Mule creates a default persistent store in the file system. If you do not specify an object store, the default persistent object store is used.
MuleSoft Reference: https://docs.mulesoft.com/mule-runtime/3.9/mule-object-stores

NEW QUESTION 7
Refer to the exhibit.
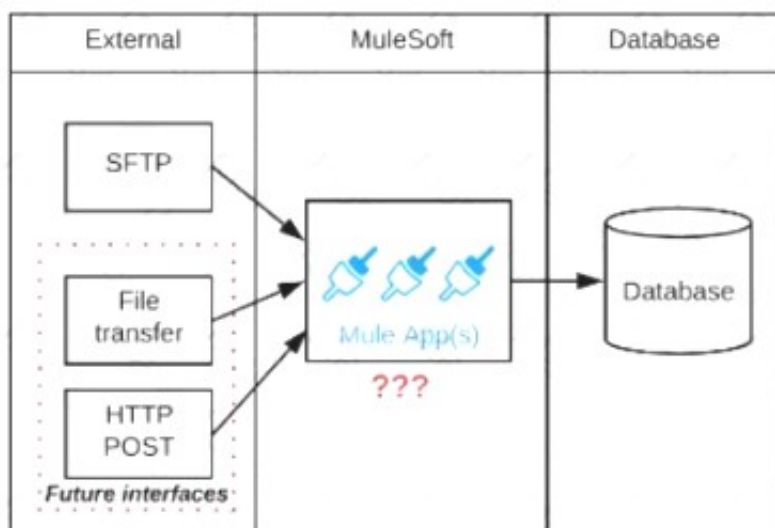
A business process involves the receipt of a file from an external vendor over SFTP. The file needs to be parsed and its content processed, validated, and ultimately persisted to a database. The delivery mechanism is expected to change in the future as more vendors send similar files using other mechanisms such as file transfer or HTTP POST.
What is the most effective way to design for these requirements in order to minimize the impact of future change?

A. Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
B. Create a Process API to receive the file and process it using a MuleSoft Batch Job while delegating the data save process to a System API
C. Create an API that receives the file and invokes a Process API with the data contained In the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed
D. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing

Answer: C

Explanation:
* Scatter-Gather is used for parallel processing, to improve performance. In this scenario, input files are coming from different vendors so mostly at different times. Goal here is to minimize the impact of future change. So scatter Gather is not the correct choice.
* If we use 1 API to receive all files from different Vendors, any new vendor addition will need changes to that 1 API to accommodate new requirements. So Option A and C are also ruled out.
* Correct answer is Create an API that receives the file and invokes a Process API with the data contained in the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed. Answer to this question lies in the API led connectivity approach.
* API-led connectivity is a methodical way to connect data to applications through a series of reusable and purposeful modern APIs that are each developed to play a specific role – unlock data from systems, compose data into processes, or deliver an experience. System API : System API tier, which provides consistent, managed, and secure access to backend systems. Process APIs : Process APIs take core assets and combines them with some business logic to create a higher level of value. Experience APIs : These are designed specifically for consumption by a specific end-user app or device.
So in case of any future plans , organization can only add experience API on addition of new Vendors, which reuse the already existing process API. It will keep impact minimal.
Diagram Description automatically generated
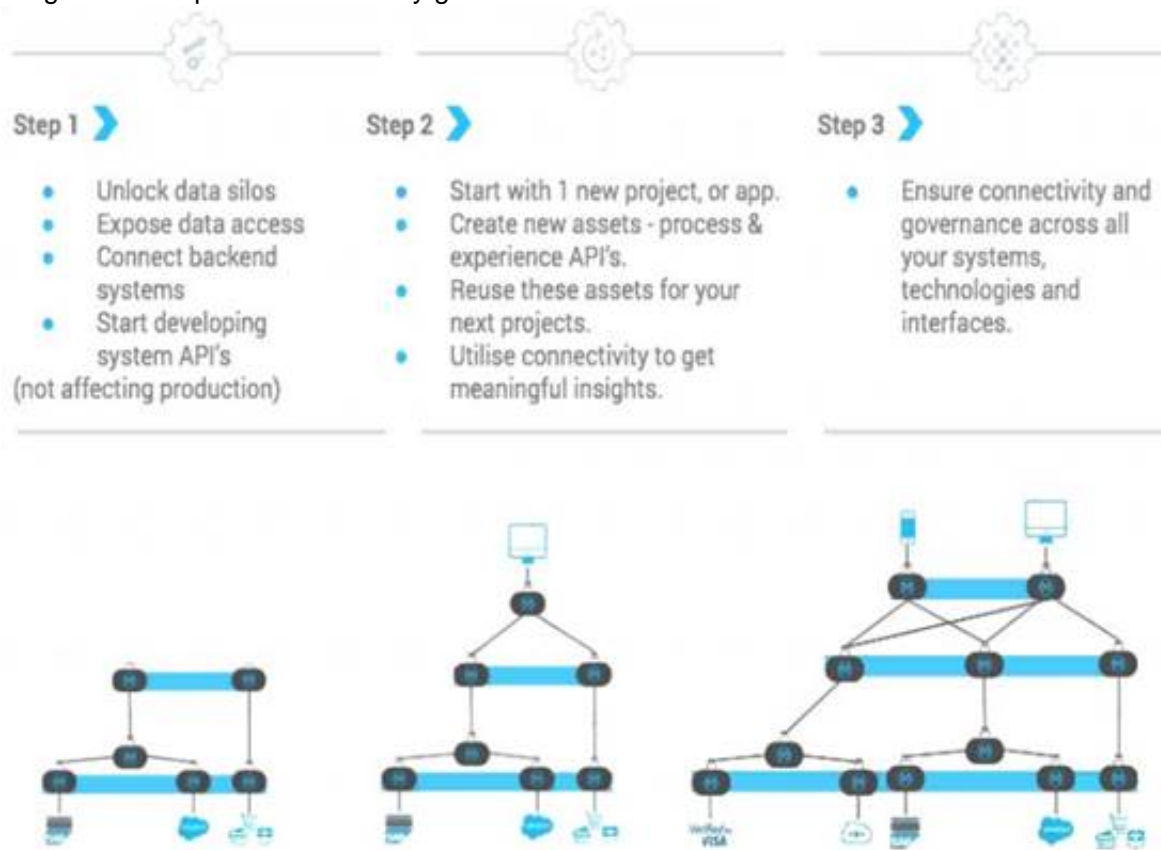
NEW QUESTION 8
A mule application is being designed to perform product orchestration. The Mule application needs to join together the responses from an inventory API and a Product Sales History API with the least latency.
To minimize the overall latency. What is the most idiomatic (used for its intended 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 9
An organization designing a hybrid, load balanced, single cluster production environment. Due to performance service level agreement goals, it is looking into running the Mule applications in an active-active multi node cluster configuration.
What should be considered when running its Mule applications in this type of environment?

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

Answer: B

NEW QUESTION 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
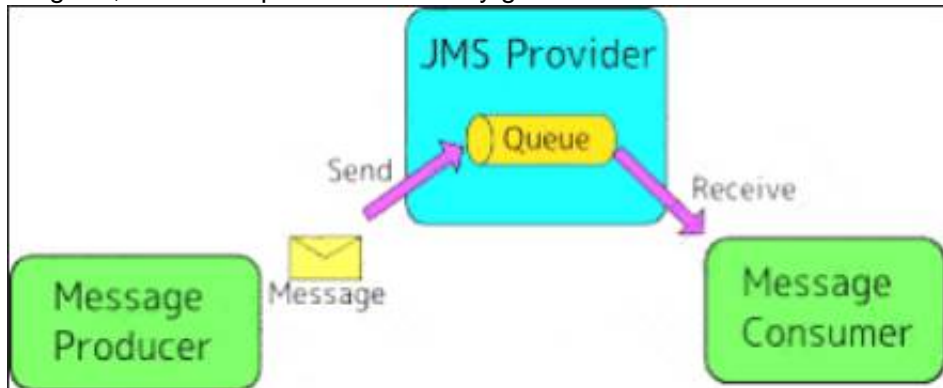
NEW QUESTION 10
A company is planning to migrate its deployment environment from on-premises cluster to a Runtime Fabric (RTF) cluster. It also has a requirement to enable Mule applications deployed to a Mule runtime instance to store and share data across application replicas and restarts.
How can these requirements be met?

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

Answer: D

NEW QUESTION 12
A Mule application is being designed To receive nightly a CSV file containing millions of records from an external vendor over SFTP, The records from the file need to be validated, transformed. And then written to a database. Records can be inserted into the database in any order.
In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

A. Use a Parallel for Each scope to Insert records one by one into the database
B. Use a Scatter-Gather to bulk insert records into the database
C. Use a Batch job scope to bulk insert records into the database.
D. Use a DataWeave map operation and an Async scope to insert records one by one into the database.

Answer: C

Explanation:
Correct answer is Use a Batch job scope to bulk insert records into the database
* Batch Job is most efficient way to manage millions of records. A few points to note here are as follows :
Reliability: If you want reliabilty while processing the records, i.e should the processing survive a runtime crash or other unhappy scenarios, and when restarted process all the remaining records, if yes then go for batch as it uses persistent queues.
Error Handling: In Parallel for each an error in a particular route will stop processing the remaining records in that route and in such case you'd need to handle it using on error continue, batch process does not stop during such error instead you can have a step for failures and have a dedicated handling in it.
Memory footprint: Since question said that there are millions of records to process, parallel for each will aggregate all the processed records at the end and can possibly cause Out Of Memory.
Batch job instead provides a BatchResult in the on complete phase where you can get the count of failures and success. For huge file processing if order is not a concern definitely go ahead with Batch Job

NEW QUESTION 13
What operation can be performed through a JMX agent enabled in a Mule application?

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

Answer: C

Explanation:
JMX Management Java Management Extensions (JMX) is a simple and standard way to manage applications, devices, services, and other resources. JMX is dynamic, so you can use it to monitor and manage resources as they are created, installed, and implemented. You can also use JMX to monitor and manage the Java Virtual Machine (JVM). Each resource is instrumented by one or more Managed Beans, or MBeans. All MBeans are registered in an MBean Server. The JMX

server agent consists of an MBean Server and a set of services for handling Mbeans. There are several agents provided with Mule for JMX support. The easiest way to configure JMX is to use the default JMX support agent. Log4J Agent The log4j agent exposes the configuration of the Log4J instance used by Mule for JMX management. You enable the Log4J agent using the <jmx-log4j> element. It does not take any additional properties MuleSoft Reference: https://docs.mulesoft.com/mule-runtime/3.9/jmx-management

NEW QUESTION 17

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

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

Answer: D

Explanation:
The mocking service is a feature of Anypoint Platform and runs continuously. You can run the mocking service from the text editor, the visual editor, and from Anypoint Exchange. You can simulate calls to the API in API Designer before publishing the API specification to Exchange or in Exchange after publishing the API specification.

NEW QUESTION 21

A Mule application is deployed to a cluster of two(2) cusomter-hosted Mule runtimes. Currently the node name Alice is the primary node and node named bob is the secondary node. The mule application has a flow that polls a directory on a file system for new files.
The primary node Alice fails for an hour and then restarted.
After the Alice node completely restarts, from what node are the files polled, and what node is now the primary node for the cluster?

A. Files are polled from Alice node Alice is now the primary node
B. Files are polled form Bob node Alice is now the primary node
C. Files are polled from Alice node Bob is the now the primary node
D. Files are polled form Bob node Bob is now the primary node

Answer: D

Explanation:
* Mule High Availability Clustering provides basic failover capability for Mule. * When the primary Mule Runtime becomes unavailable, for example, because of a fatal JVM or hardware failure or it's taken offline for maintenance, a backup Mule Runtime immediately becomes the primary node and resumes processing where the failed instance left off. * After a system administrator recovers a failed Mule Runtime server and puts it back online, that server automatically becomes the backup node. In this case, Alice, once up, will become backup
--Reference: https://docs.mulesoft.com/mule-runtime/4.3/hadr-guide So correct choice is : Files are polled form Bob node Bob is now the primary node

NEW QUESTION 22

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

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

Answer: C

NEW QUESTION 26

A Mule application is being designed to do the following:
Step 1: Read a SalesOrder message from a JMS queue, where each SalesOrder consists of a header and a list of SalesOrderLineItems.
Step 2: Insert the SalesOrder header and each SalesOrderLineItem into different tables in an RDBMS.
Step 3: Insert the SalesOrder header and the sum of the prices of all its SalesOrderLineItems into a table In a different RDBMS.
No SalesOrder message can be lost and the consistency of all SalesOrder-related information in both RDBMSs must be ensured at all times.
What design choice (including choice of transactions) and order of steps addresses these requirements?

A. 1) Read the JMS message (NOT in an XA transaction)2) Perform BOTH DB inserts in ONE DB transaction3) Acknowledge the JMS message
B. 1) Read the JMS message (NOT in an XA transaction)2) Perform EACH DB insert in a SEPARATE DB transaction3) Acknowledge the JMS message
C. 1) Read the JMS message in an XA transaction2) In the SAME XA transaction, perform BOTH DB inserts but do NOT acknowledge the JMS message
D. 1) Read and acknowledge the JMS message (NOT in an XA transaction)2) In a NEW XA transaction, perform BOTH DB inserts

Answer: A

Explanation:
Option A says "Perform EACH DB insert in a SEPARATE DB transaction". In this case if first DB insert is successful and second one fails then first insert won't be rolled back causing inconsistency. This option is ruled out.
Option D says Perform BOTH DB inserts in ONE DB transaction.

Rule of thumb is when one or more DB connections are required we must use XA transaction as local transactions support only one resource. So this option is also ruled out.

Option B acknowledges the before DB processing, so message is removed from the queue. In case of system failure at later point, message can't be retrieved.

Option C is Valid: Though it says "do not ack JMS message", message will be auto acknowledged at the end of transaction. Here is how we can ensure all components are part of XA transaction: https://docs.mulesoft.com/jms-connector/1.7/jms-transactions

Additional Information about transactions:

XA Transactions - You can use an XA transaction to group together a series of operations from multiple transactional resources, such as JMS, VM or JDBC resources, into a single, very reliable, global transaction.

The XA (eXtended Architecture) standard is an X/Open group standard which specifies the interface between a global transaction manager and local transactional resource managers.

The XA protocol defines a 2-phase commit protocol which can be used to more reliably coordinate and sequence a series of "all or nothing" operations across multiple servers, even servers of different types

Use JMS ack if
– Acknowledgment should occur eventually, perhaps asynchronously
– The performance of the message receipt is paramount
– The message processing is idempotent
– For the choreography portion of the SAGA pattern Use JMS transactions
– For all other times in the integration you want to perform an atomic unit of work
– When the unit of work comprises more than the receipt of a single message
– To simply and unify the programming model (begin/commit/rollback)

NEW QUESTION 29
An organization is designing multiple new applications to run on CloudHub in a single Anypoint VPC and that must share data using a common persistent Anypoint object store V2 (OSv2).
Which design gives these mule applications access to the same object store instance?

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

Answer: D

NEW QUESTION 34
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 36
What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

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

Answer: A

Explanation:
* Shared load balancers don't allow you to configure custom SSL certificates or proxy rules
* Dedicated Load Balancer are optional but you need to purchase them additionally if needed.
* 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.
* The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
* Only a DLB allows the configuration of a custom TLS server certificate
* DLB enables you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
* To use a DLB in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.
* MuleSoft Reference: https://docs.mulesoft.com/runtime-manager/dedicated-load-balancer-tutorial Additional Info on SLB Vs DLB:
Table Description automatically generated

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

NEW QUESTION 39

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

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

Answer: C

NEW QUESTION 43
Refer to the exhibit.
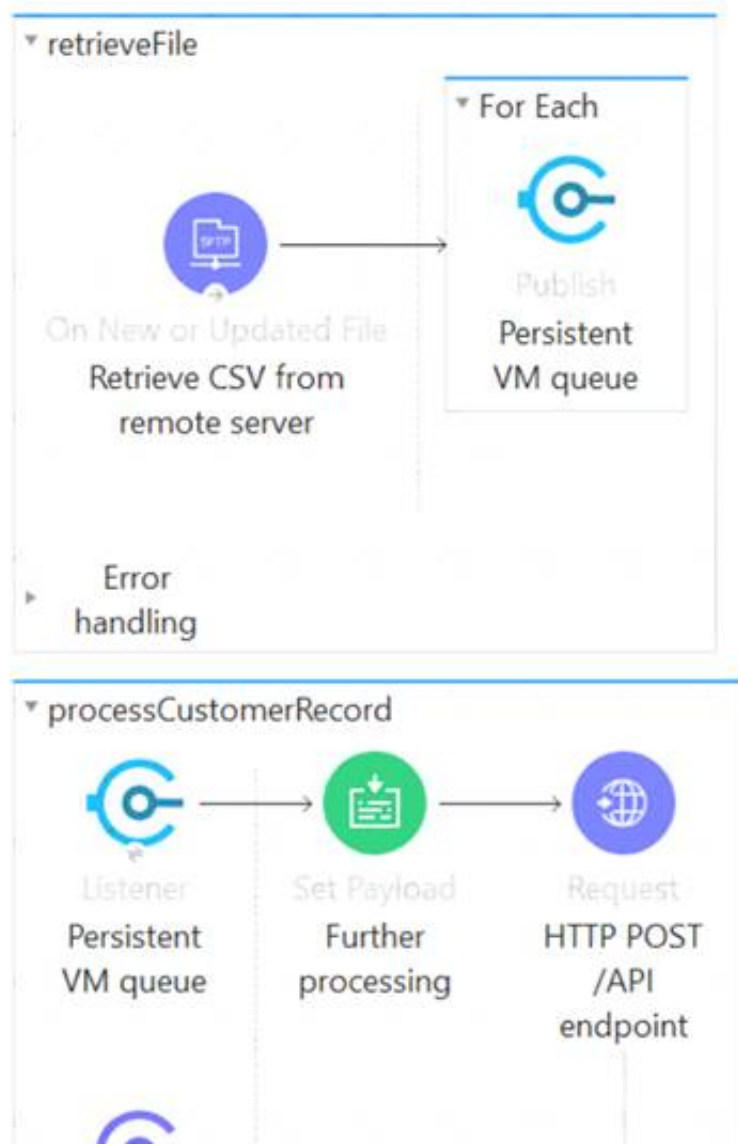
This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow's VM Listner receives messages from the same VM queue and then processes each message separately.
How are messages routed to the cloudhub workers as messages are received by the VM Listener?

A. Each message is routed to ONE of the Cloudhub workers in a DETERMINSTIC round robin fashion thereby EXACTLY BALANCING messages among the cloudhub workers
B. Each messages routes to ONE of the available Clouhub workers in a NON- DETERMINSTIC non round-robin fashion thereby APPROXIMATELY BALANCING messages among the cloudhub workers
C. Each message is routed to the SAME Cloudhub worker that retrieved the file, thereby BINDING ALLmessages to ONLY that ONE Cloudhub worker
D. Each message is duplicated to ALL of the Cloudhub workers, thereby SHARING EACH message with ALL the Cloudhub workers.

Answer: B

NEW QUESTION 48
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 50
In Anypoint Platform, a company wants to configure multiple identity providers (IdPs) for multiple lines of business (LOBs). Multiple business groups, teams, and environments have been defined for these LOBs.
What Anypoint Platform feature can use multiple IdPs across the company's business groups, teams, and environments?

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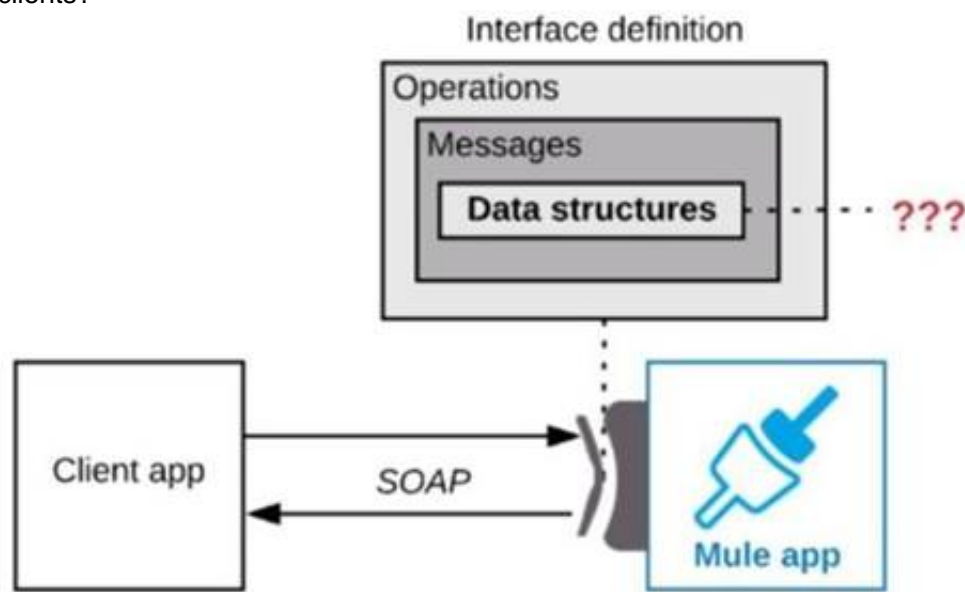
Answer: A

Explanation:
To use a dedicated load balancer in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.

NEW QUESTION 51
Refer to the exhibit.
A Mule application is being designed to expose a SOAP web service to its clients.

What language is typically used inside the web service's interface definition to define the data structures that the web service Is expected to exchange with its clients?

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

Answer: B

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

NEW QUESTION 52
When designing an upstream API and its implementation, the development team has been advised to not set timeouts when invoking downstream API. Because the downstream API has no SLA that can be relied upon. This is the only donwstream API dependency of that upstream API. Assume the downstream API runs uninterrupted without crashing. What is the impact of this advice?

A. The invocation of the downstream API will run to completion without timing out.
B. An SLA for the upstream API CANNOT be provided.
C. A default timeout of 500 ms will automatically be applied by the Mule runtime in which the upstream API implementation executes.
D. A load-dependent timeout of less than 1000 ms will be applied by the Mule runtime in which the downstream API implementation executes.

Answer: B

Explanation:
An SLA for the upstream API CANNOT be provided.

NEW QUESTION 55
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 parallelAll 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 timeEach Batch Step scope is invoked with one record in the payload of the received Mule eventFor 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 Stepscope
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 recordsEach Batch Step scope is invoked with one record in the payload of the received Mule eventFor 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 nextBatch Step scope
D. The Batch Job scope processes multiple record blocks in parallelEach 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 parallelIndividual 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 56
What metrics about API invocations are available for visualization in custom charts using Anypoint Analytics?

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

Answer: C

Explanation:

Correct answer is Request size, number of requests, response size, response time Analytics API Analytics can provide insight into how your APIs are being used and how they are performing. From API Manager, you can access the Analytics dashboard, create a custom dashboard, create and manage charts, and create reports. From API Manager, you can get following types of analytics: - API viewing analytics - API events analytics - Charted metrics in API Manager
It can be accessed using: http://anypoint.mulesoft.com/analytics
API Analytics provides a summary in chart form of requests, top apps, and latency for a particular duration. The custom dashboard in Anypoint Analytics contains a set of charts for a single API or for all APIs Each
chart displays various API characteristics
– Requests size: Line chart representing size of requests in KBs
– Requests : Line chart representing number of requests over a period
– Response size : Line chart representing size of response in KBs
– Response time :Line chart representing response time in ms
* To check this, You can go to API Manager > Analytics > Custom Dashboard > Edit Dashboard > Create Chart > Metric
Graphical user interface, chart Description automatically generated
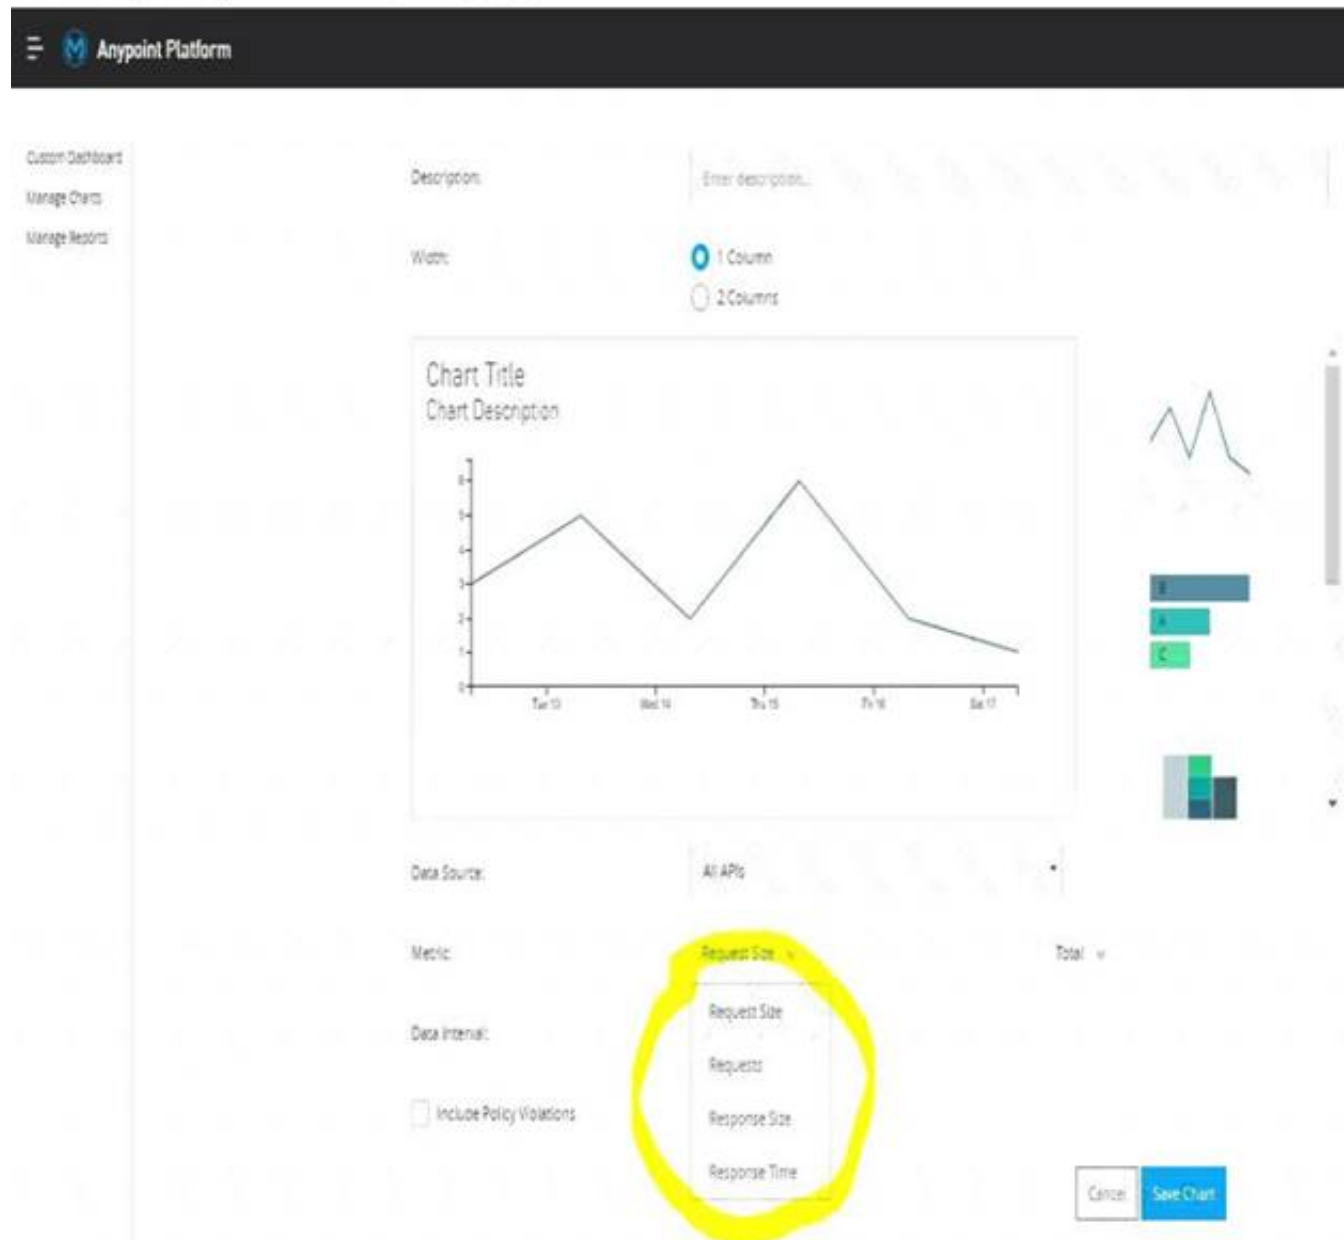
NEW QUESTION 58
When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

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 61
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
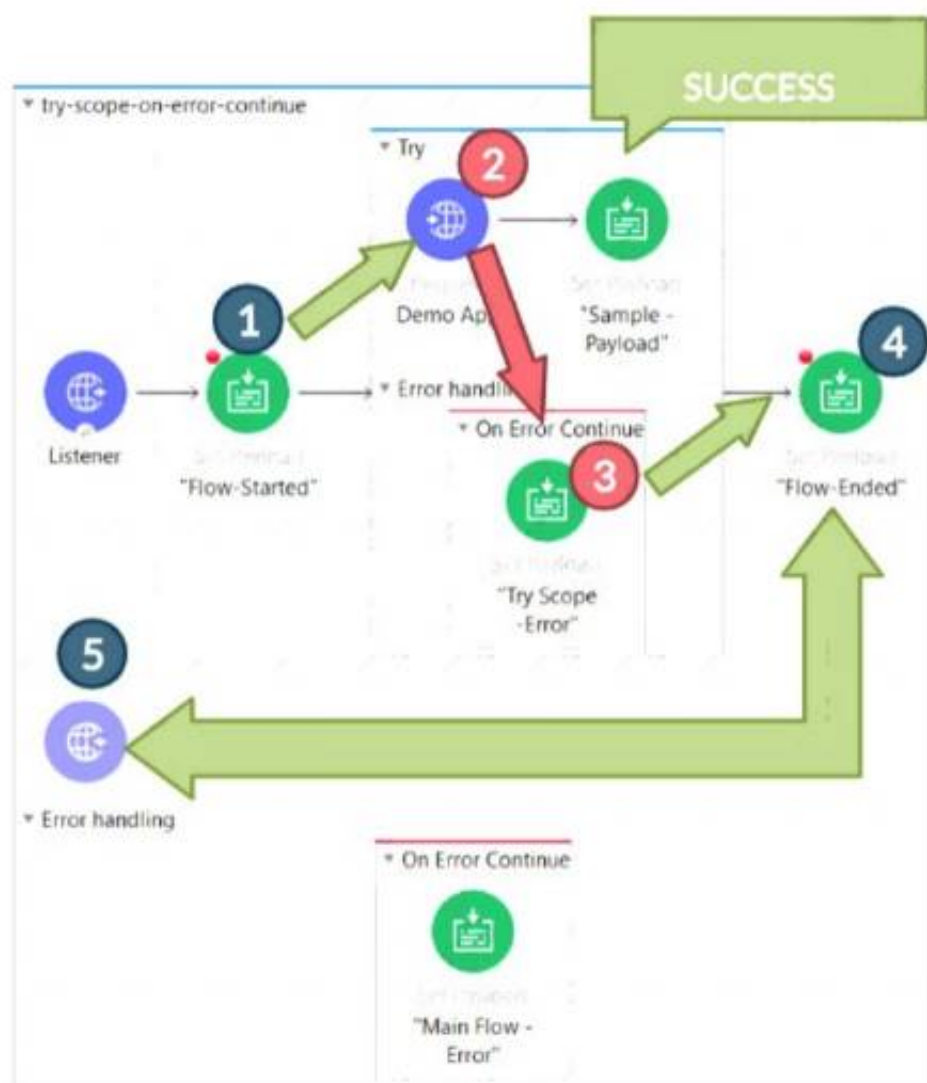
NEW QUESTION 66
An organization if struggling frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects.
How can these best practices be achieved with the applications having the least amount of responsibility?

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

Answer: D

NEW QUESTION 71
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 73
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 78
Which of the below requirements prevent the usage of Anypoint MQ in a company's network? (Choose two answers)

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

Answer: CD

NEW QUESTION 81
An organization has deployed runtime fabric on an eight note cluster with performance profile. An API uses and non persistent object store for maintaining some of its state data. What will be the impact to the stale data if server crashes?

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

Answer: D

NEW QUESTION 86
A retailer is designing a data exchange interface to be used by its suppliers. The interface must support secure communication over the public internet. The interface must also work with a wide variety of programming languages and IT systems used by suppliers.
What are suitable interface technologies for this data exchange that are secure, cross-platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

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

Answer: C

Explanation:
As per definition of API by Mulesoft , it is Application Programming Interface using HTTP-based protocols. Non-HTTP-based programmatic interfaces are not APIs.
* HTTP-based programmatic interfaces are APIs even if they don't use REST or JSON. Hence implementation based on Java RMI, CORBA/IIOP, raw TCP/IP interfaces are not API's as they are not using HTTP.
* One more thing to note is FTP was not built to be secure. It is generally considered to be an insecure protocol because it relies on clear-text usernames and passwords for authentication and does not use encryption.
* Data sent via FTP is vulnerable to sniffing, spoofing, and brute force attacks, among other basic attack methods.
Considering the above points only correct option is
-XML over ActiveMQ
- XML over SFTP
- XML/REST over HTTPS

NEW QUESTION 89

An organization has decided on a cloud migration strategy to minimize the organization's own IT resources. Currently the organization has all of its new applications running on its own premises and uses an on-premises load balancer that exposes all APIs under the base URL (https://api.rutujar.com).

As part of migration strategy, the organization is planning to migrate all of its new applications and load balancer CloudHub.

What is the most straightforward and cost-effective approach to Mule application deployment and load balancing that preserves the public URL's?

A. Deploy the Mule application to CloudhubCreate a CNAME record for base URL(httpsr://api.rutujar.com) in the Cloudhub shared load balancer that points to the A record of theon-premises load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications

B. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL (https://api.rutujar.com) in the organization's DNS server to point to the A record of the Cloudhub dedicated load balancerApply mapping rules in DLB to map URLto their corresponding Mule applications

C. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL (https://api.rutujar.com) in the organization's DNS server to point to the A record of the CloudHub shared load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications

D. For each migrated Mule application, deploy an API proxy application to Cloudhub with all traffic to themule applications routed through a Cloud Hub Dedicated load balancer (DLB)Update a CNAME record for base URL (https://api.rutujar.com) in the organization's DNS server to point to the A record of the CloudHub dedicated load balancerApply mapping rules in DLB to map each API proxy application who is responding new application

Answer: B

NEW QUESTION 94

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 98

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 103

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 104

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 107
Which Mulesoft feature helps users to delegate their access without sharing sensitive credentials or giving full control of accounts to 3rd parties?

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

Answer: C

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

NEW QUESTION 110
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 112
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-SAMLTokenProfile.pdf

NEW QUESTION 116
Refer to the exhibit.
An organization is designing a Mule application to receive data from one external business partner. The two companies currently have no shared IT infrastructure and do not want to establish one. Instead, all communication should be over the public internet (with no VPN).
What Anypoint Connector can be used in the organization's Mule application to securely receive data from this external business partner?
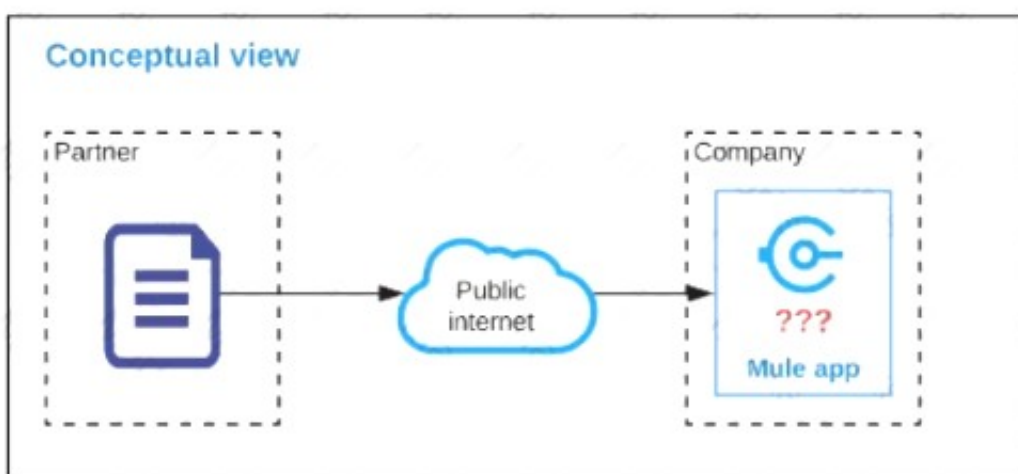
A. File connector
B. VM connector
C. SFTP connector
D. Object Store connector

Answer: C

Explanation:
* Object Store and VM Store is used for sharing data inter or intra mule applications in same setup. Can't be used with external Business Partner
* Also File connector will not be useful as the two companies currently have no shared IT infrastructure. It's specific for local use.
* Correct answer is SFTP connector. The SFTP Connector implements a secure file transport channel so that your Mule application can exchange files with external resources. SFTP uses the SSH security protocol to transfer messages. You can implement the SFTP endpoint as an inbound endpoint with a one-way exchange pattern, or as an outbound endpoint configured for either a one-way or request-response exchange pattern.

NEW QUESTION 120
A manufacturing company is planning to deploy Mule applications to its own Azure Kubernetes Service infrastructure.
The organization wants to make the Mule applications more available and robust by deploying each Mule application to an isolated Mule runtime in a Docker container while managing all the Mule applications from the MuleSoft-hosted control plane.
What is the most idiomatic (used for its intended purpose) choice of runtime plane to meet these organizational requirements?

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

**Answer:** B

NEW QUESTION 121
Mule applications need to be deployed to CloudHub so they can access on-premises database systems. These systems store sensitive and hence tightly protected data, so are not accessible over the internet.
What network architecture supports this requirement?

A. An Anypoint VPC connected to the on-premises network using an IPsec tunnel or AWS DirectConnect, plus matching firewall rules in the VPC and on-premises network
B. Static IP addresses for the Mule applications deployed to the CloudHub Shared Worker Cloud, plus matching firewall rules and IPwhitelisting in the on-premises network
C. An Anypoint VPC with one Dedicated Load Balancer fronting each on-premises database system, plus matching IP whitelisting in the load balancer and firewall rules in the VPC and on-premises network
D. Relocation of the database systems to a DMZ in the on-premises network, with Mule applications deployed to the CloudHub Shared Worker Cloud connecting only to the DMZ

Answer: A

Explanation:
* "Relocation of the database systems to a DMZ in the on-premises network, with Mule applications deployed
to the CloudHub Shared Worker Cloud connecting only to the DMZ" is not a feasible option
* "Static IP addresses for the Mule applications deployed to the CloudHub Shared Worker Cloud, plus matching firewall rules and IP whitelisting in the on-premises network" - It is risk for sensitive data. - Even if you whitelist the database IP on your app, your app wont be able to connect to the database so this is also not a feasible option
* "An Anypoint VPC with one Dedicated Load Balancer fronting each on-premises database system, plus matching IP whitelisting in the load balancer and firewall rules in the VPC and on-premises network" Adding one VPC with a DLB for each backend system also makes no sense, is way too much work. Why would you add a LB for one system.
* Correct Answer "An Anypoint VPC connected to the on-premises network using an IPsec tunnel or AWS DirectConnect, plus matching firewall rules in the VPC and on-premises network"
IPsec Tunnel You can use an IPsec tunnel with network-to-network configuration to connect your on-premises data centers to your Anypoint VPC. An IPsec VPN tunnel is generally the recommended solution for VPC to on-premises connectivity, as it provides a standardized, secure way to connect. This method also integrates well with existing IT infrastructure such as routers and appliances.
Reference: https://docs.mulesoft.com/runtime-manager/vpc-connectivity-methods-concept
Diagram Description automatically generated
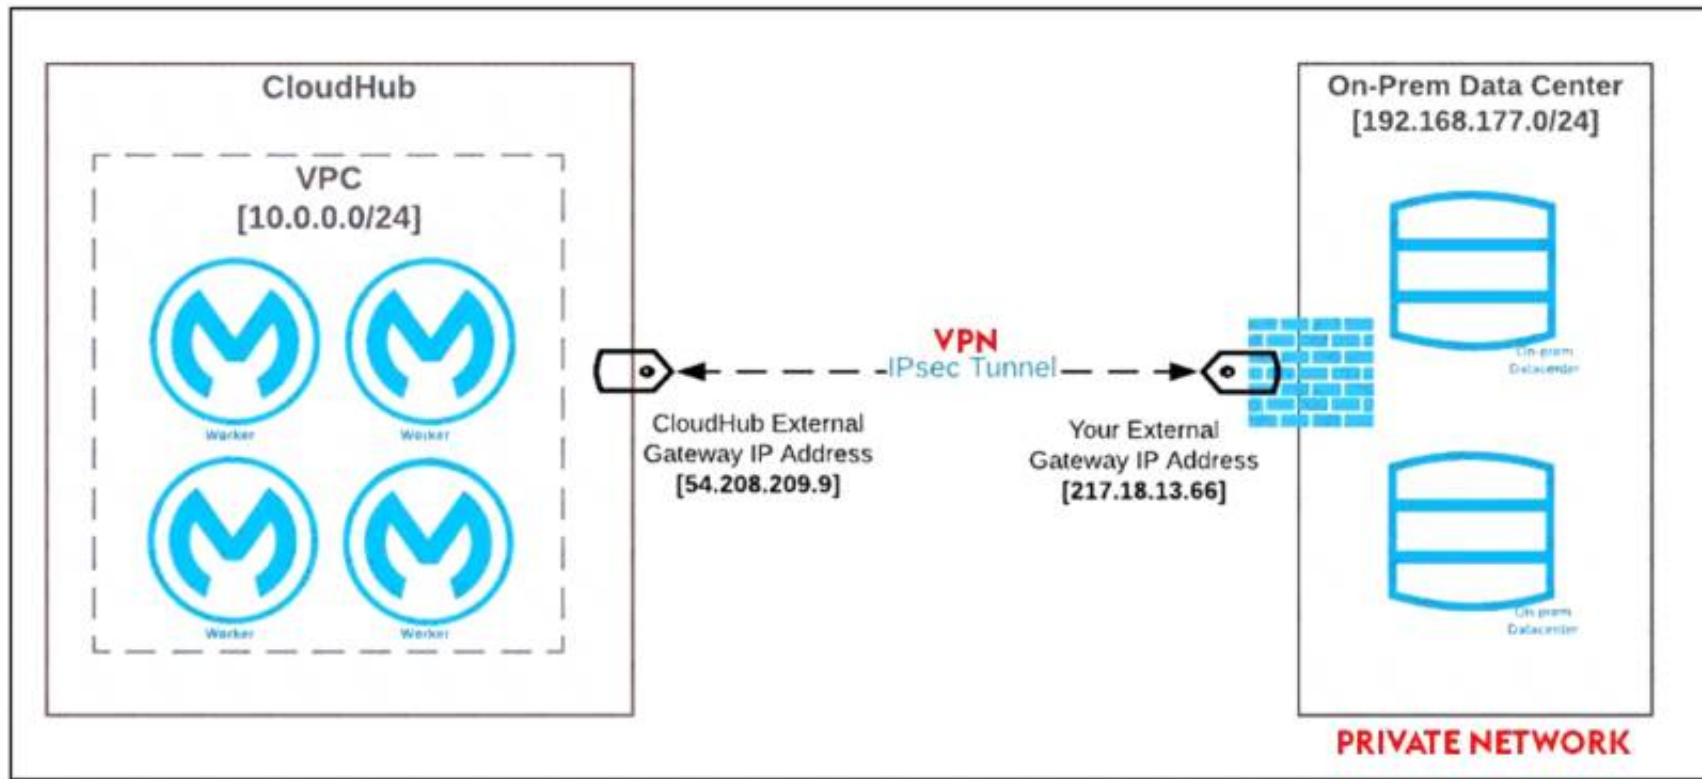
NEW QUESTION 126
An external REST client periodically sends an array of records in a single POST request to a Mule application API endpoint.
The Mule application must validate each record of the request against a JSON schema before sending it to a downstream system in the same order that it was received in the array
Record processing will take place inside a router or scope that calls a child flow. The child flow has its own error handling defined. Any validation or communication failures should not prevent further processing of the remaining records.
To best address these requirements what is the most idiomatic(used for it intended purpose) router or scope to used in the parent flow, and what type of error handler should be used in the child flow?

A. First Successful router in the parent flowOn Error Continue error handler in the child flow
B. For Each scope in the parent flowOn Error Continue error handler in the child flow
C. Parallel For Each scope in the parent flowOn Error Propagate error handler in the child flow
D. Until Successful router in the parent flowOn Error Propagate error handler in the child flow

Answer: B

Explanation:
Correct answer is For Each scope in the parent flow On Error Continue error handler in the child flow. You can extract below set of requirements from the question
a) Records should be sent to downstream system in the same order that it was received in the array b) Any validation or communication failures should not prevent further processing of the remaining records First requirement can be met using For Each scope in the parent flow and second requirement can be met using On Error Continue scope in child flow so that error will be suppressed.

NEW QUESTION 127
The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure.
Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).
Given the company's current setup and requirements, what is the most idiomatic (used for its intended purpose) way to send Mule application logs to the external LMT?

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

Answer: A

NEW QUESTION 128
An integration Mute application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mute application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system.
Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system).
The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed.
The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages.
What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

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

Answer: D

Explanation:
Correct answer is using below set of activities Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing We will

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

NEW QUESTION 130
What Anypoint Connectors support transactions?

A. Database, JMS, VM
B. Database, 3MS, HTTP
C. Database, JMS, VM, SFTP
D. Database, VM, File

Answer: A

Explanation:
Below Anypoint Connectors support transactions JMS – Publish – Consume VM – Publish – Consume Database – All operations

NEW QUESTION 131
49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58.
https://api.acme.com/partnefs/vl. The API implementation is estimated to require deployment to 5 CloudHub workers.
The retailer has obtained a public X.509 certificate for the name apl.acme.com, signed by a reputable CA, to be used as the server certificate.
Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

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

Answer: C

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

NEW QUESTION 132
One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make.
Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments.
The production environment however does not have any rate limiting.
What is the cost-effective approach to conduct performance test of the API implementation in the non-production staging environment?

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

Answer: C

NEW QUESTION 134
An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events
What are the two considerations that must be kept in mind while using test recorder (Choose two answers)

A. Tests for flows cannot be created with Mule errors raised inside the flow or already existing in the incoming event

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

Answer: AD

NEW QUESTION 135
An API has been updated in Anypoint Exchange by its API producer from version 3.1.1 to 3.2.0 following accepted semantic versioning practices and the changes have been communicated via the API's public portal. The API endpoint does NOT change in the new version. How should the developer of an API client respond to this change?

A. The update should be identified as a project risk and full regression testing of the functionality that uses this API should be run.
B. The API producer should be contacted to understand the change to existing functionality.
C. The API producer should be requested to run the old version in parallel with the new one.
D. The API client code ONLY needs to be changed if it needs to take advantage of new features.

Answer: D

Explanation:
* Semantic Versioning is a 3-component number in the format of X.Y.Z, where : X stands for a major version.
Y stands for a minor version:
Z stands for a patch.
So, SemVer is of the form Major.Minor.Patch Coming to our question , minor version of the API has been changed which is backward compatible. Hence there is no change required on API client end. If they want to make use of new featured that have been added as a part of minor version change they may need to change code at their end. Hence correct answer is The API client code ONLY needs to be changed if it needs to take advantage of new features.
Diagram Description automatically generated

NEW QUESTION 137
What aspects of a CI/CD pipeline for Mule applications can be automated using MuleSoft-provided Maven plugins?

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

Answer: A

Explanation:
Correct answer is "Compile, package, unit test, validate unit test coverage, deploy"
Anypoint Platform supports continuous integration and continuous delivery using industry standard tools Mule Maven Plugin The Mule Maven plugin can automate building, packaging and deployment of Mule applications from source projects Using the Mule Maven plugin, you can automate your Mule application deployment to CloudHub, to Anypoint Runtime Fabric, or on-premises, using any of the following deployment strategies • CloudHub deployment • Runtime Fabric deployment • Runtime Manager REST API deployment • Runtime Manager agent deployment MUnit Maven Plugin The MUnit Maven plugin can automate test execution, and ties in with the Mule Maven plugin. It provides a full suite of integration and unit test capabilities, and is fully integrated with Maven and Surefire for integration with your continuous deployment environment. Since MUnit 2.x, the coverage report goal is integrated with the maven reporting section. Coverage Reports are generated during Maven's site lifecycle, during the coverage-report goal. One of the features of MUnit Coverage is to fail the build if a certain coverage level is not reached. MUnit is not used for integration testing Also publishing to Anypoint Exchange or to create associated API instances in API Manager is not a part of CICD pipeline which can ne achieved using mulesoft provided maven plugin
Explanation
Architecture mentioned in the question can be diagrammatically put as below. Persistent Object Store is the correct answer .
* Mule Object Stores: An object store is a facility for storing objects in or across Mule applications. Mule uses object stores to persist data for eventual retrieval.
Mule provides two types of object stores:
1) In-memory store – stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime. So we cant use in memory store in our scenario as we want to share watermark within all cloudhub workers
2) Persistent store – Mule persists data when an object store is explicitly configured to be persistent. Hence this watermark will be available even any of the worker goes down
Diagram Description automatically generated
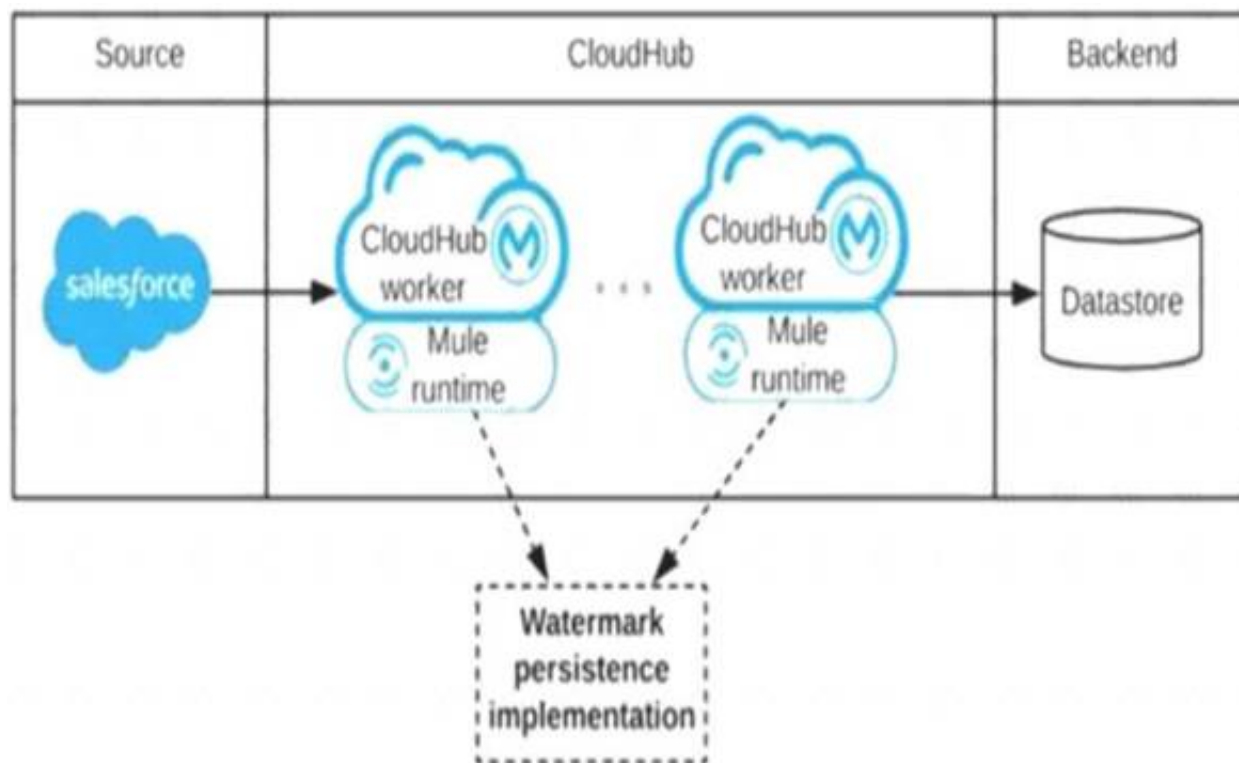
NEW QUESTION 141
Anypoint Exchange is required to maintain the source code of some of the assets committed to it, such as Connectors, Templates, and API specifications.
What is the best way to use an organization's source-code management (SCM) system in this context?

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 143
An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger of flow for processing a high-volume set of records from the source system and synchronize with the SaaS system using the Batch job scope
After processing 1000 records in a periodic synchronization of 1 lakh records, the replicas in which batch job instance was started went down due to unexpected failure in the runtime fabric environment
What is the consequence of losing the replicas that run the Batch job instance?

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

Answer: B

NEW QUESTION 146
A company is building an application network and has deployed four Mule APIs: one experience API, one process API, and two system APIs. The logs from all the APIs are aggregated in an external log aggregation tool. The company wants to trace messages that are exchanged between multiple API implementations. What is the most idiomatic (based on its intended use) identifier that should be used to implement Mule event tracing across the multiple API implementations?

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

Answer: B

Explanation:
Correct answer is Mule correlation ID By design, Correlation Ids cannot be changed within a flow in Mule 4 applications and can be set only at source. This ID is part of the Event Context and is generated as soon as the message is received by the application. When a HTTP Request is received, the request is inspected for "X-Correlation-Id" header. If "X-Correlation-Id" header is present, HTTP connector uses this as the Correlation Id. If "X-Correlation-Id" header is NOT present, a Correlation Id is randomly generated. For Incoming HTTP Requests: In order to set a custom Correlation Id, the client invoking the HTTP request must set "X-Correlation-Id" header. This will ensure that the Mule Flow uses this Correlation Id. For Outgoing HTTP Requests: You can also propagate the existing Correlation Id to downstream APIs. By default, all outgoing HTTP Requests send "X-Correlation-Id" header. However, you can choose to set a different value to "X-Correlation-Id" header or set "Send Correlation Id" to NEVER.

NEW QUESTION 148
A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is 58. proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application.
What non-functional requirement can be expected to be negatively affected when implementing this change?

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

Answer: B

Explanation:
Correct answer is Consistency as XA transactions are implemented to achieve this. XA transactions are added in the implementation to achieve goal of ACID properties. In the context of transaction processing, the acronym ACID refers to the four key properties of a transaction: atomicity, consistency, isolation, and durability. Atomicity : All changes to data are performed as if they are a single operation. That is, all the changes are performed, or none of them are. For example, in an application that transfers funds from one account to another, the atomicity property ensures that, if a debit is made successfully from one account, the corresponding credit is made to the other account. Consistency : Data is in a consistent state when a transaction starts and when it ends.For example, in an application that transfers funds from one account to another, the consistency property ensures that the total value of funds in both the accounts is the same at the start and end of each transaction. Isolation : The intermediate state of a transaction is invisible to other transactions. As a result, transactions that run concurrently appear to be serialized. For example, in an application that transfers funds from one account to another, the isolation property ensures that another transaction sees the transferred funds in one account or the other, but not in both, nor in neither. Durability : After a transaction successfully completes, changes to data persist and are not undone, even in the event of a system failure. For example, in an application that transfers funds from one account to another, the durability property ensures that the changes made to each account will not be reversed. MuleSoft reference: https://docs.mulesoft.com/mule-runtime/4.3/xa-transactions

NEW QUESTION 150
An organization is designing a mule application to support an all or nothing transaction between serval database operations and some other connectors so that they all roll back if there is a problem with any of the connectors
Besides the database connector , what other connector can be used in the transaction.

A. VM
B. Anypoint MQ
C. SFTP
D. ObjectStore

Answer: A

Explanation:
Correct answer is VM VM support Transactional Type. When an exception occur, The transaction rolls back to its original state for reprocessing. This feature is not supported by other connectors.
Here is additional information about Transaction management: Table Description automatically generated

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

NEW QUESTION 153
Refer to the exhibit.
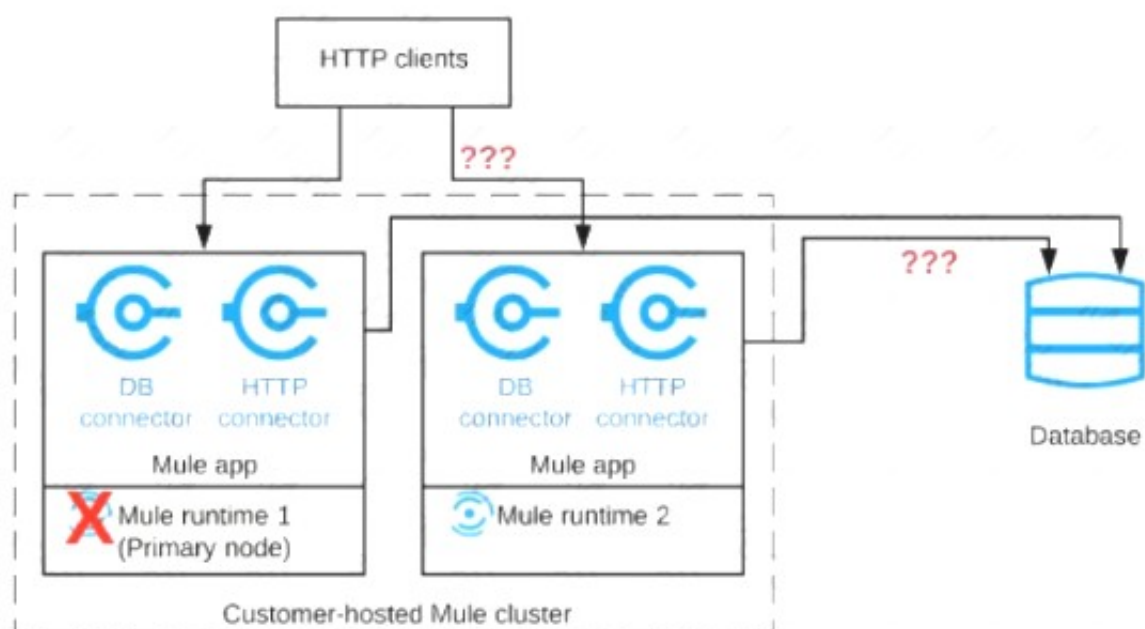
A Mule application is deployed to a cluster of two customer-hosted Mute runtimes. The Mute application has a flow that polls a database and another flow with an HTTP Listener. HTTP clients send HTTP requests directly to individual cluster nodes.
What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has railed, but before that node is restarted?

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

Answer: A

Explanation:
: Architecture descibed in the question could be described as follows.When node 1 is down , DB polling will still continue via node 2 . Also requests which are

coming directly to node 2 will also be accepted and processed in BAU fashion. Only thing that wont work is when requests are sent to Node 1 HTTP connector. The flaw with this architecture is HTTP clients are sending HTTP requests directly to individual cluster nodes. By default, clustering Mule runtime engines ensures high system availability. If a Mule runtime engine node becomes unavailable due to failure or planned downtime, another node in the cluster can assume the workload and continue to process existing events and messages

Diagram Description automatically generated
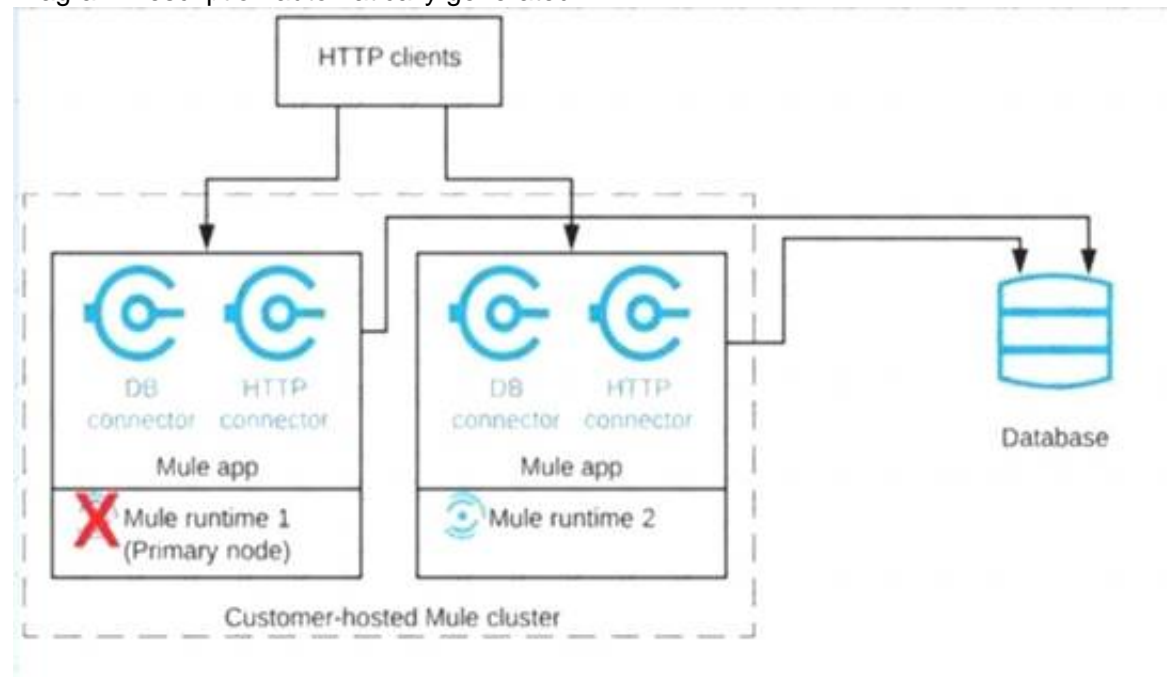
NEW QUESTION 155

A new upstream API Is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity. The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms. If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API's desired SLA?

A. Set a timeout of 100 ms; that leaves 400 ms for the other two downstream APIs to complete
B. Do not set a timeout; the Invocation of this API Is mandatory and so we must wait until it responds
C. Set a timeout of 50 ms; this times out more invocations of that API but gives additional room for retries
D. No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

Answer: D

Explanation:
Before we answer this question , we need to understand what median (50th percentile) and 80th percentile means. If the 50th percentile (median) of a response time is 500ms that means that 50% of my transactions are either as fast or faster than 500ms.
If the 90th percentile of the same transaction is at 1000ms it means that 90% are as fast or faster and only 10% are slower. Now as per upstream SLA , 99th percentile is 800 ms which means 99% of the incoming requests should have response time less than or equal to 800 ms. But as per one of the backend API , their 95th percentile is 1000 ms which means that backend API will take 1000 ms or less than that for 95% of. requests. As there are three API invocation from upstream API , we can not conclude a timeout that can be set to meet the desired SLA as backend SLA's do not support it.
Let see why other answers are not correct.
1) Do not set a timeout --> This can potentially violate SLA's of upstream API
2) Set a timeout of 100 ms; ---> This will not work as backend API has 100 ms as median meaning only 50% requests will be answered in this time and we will get timeout for 50% of the requests. Important thing to note here is, All APIs need to be executed sequentially, so if you get timeout in first API, there is no use of going to second and third API. As a service provider you wouldn't want to keep 50% of your consumers dissatisfied. So not the best option to go with.
*To quote an example: Let's assume you have built an API to update customer contact details.
- First API is fetching customer number based on login credentials
- Second API is fetching Info in 1 table and returning unique key
- Third API, using unique key provided in second API as primary key, updating remaining details
* Now consider, if API times out in first API and can't fetch customer number, in this case, it's useless to call API 2 and 3 and that is why question mentions specifically that all APIs need to be executed sequentially.
3) Set a timeout of 50 ms --> Again not possible due to the same reason as above Hence correct answer is No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

NEW QUESTION 156
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual MCIA-Level-1 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the MCIA-Level-1 Product From:

https://www.2passeasy.com/dumps/MCIA-Level-1/

Money Back Guarantee

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