MCIA-Level-1 Dumps

MuleSoft Certified Integration Architect - Level 1

https://www.certleader.com/MCIA-Level-1-dumps.html

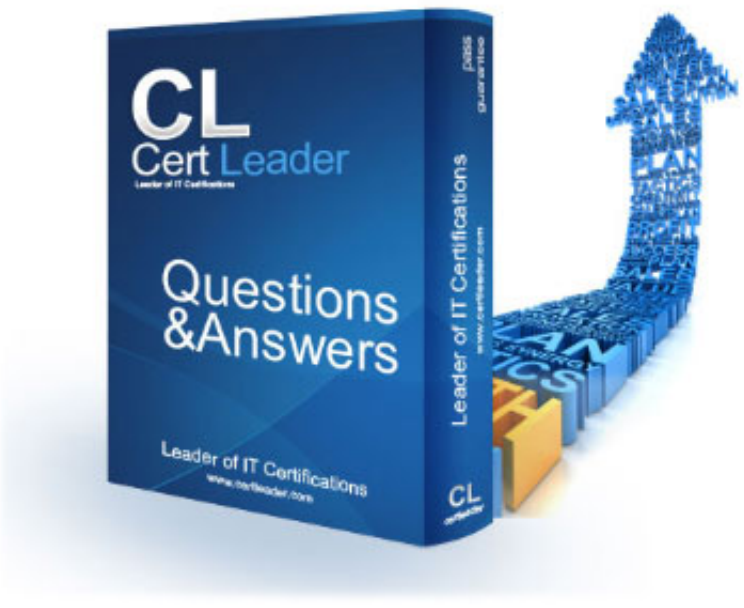

NEW QUESTION 1
A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction.
What is the behavior of the Scatter-Gather when running within this local transaction?

A. Execution of each route within the Scatter-Gather occurs sequentiallyAny error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
B. Execution of all routes within the Scatter-Gather occurs in parallelAny error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
C. Execution of each route within the Scatter-Gather occurs sequentiallyAny error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations
D. Execution of each route within the Scatter-Gather occurs in parallelAny error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations

Answer: A

NEW QUESTION 2
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 3
An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to custom hosted Mule runtimes (on-premises) while others execute in the MuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems.
How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access the backend systems?

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

Answer: D

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

NEW QUESTION 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 Mule application is being designed for deployment to a single CloudHub worker. The Mule application will have a flow that connects to a SaaS system to perform some operations each time the flow is invoked.
The SaaS system connector has operations that can be configured to request a short-lived token (fifteen minutes) that can be reused for subsequent connections within the fifteen minute time window. After the token expires, a new token must be requested and stored.
What is the most performant and idiomatic (used for its intended purpose) Anypoint Platform component or service to use to support persisting and reusing tokens in the Mule application to help speed up reconnecting the Mule application to the SaaS application?

A. Nonpersistent object store
B. Persistent object store
C. Variable
D. Database

Answer: D

NEW QUESTION 6
Refer to the exhibit.

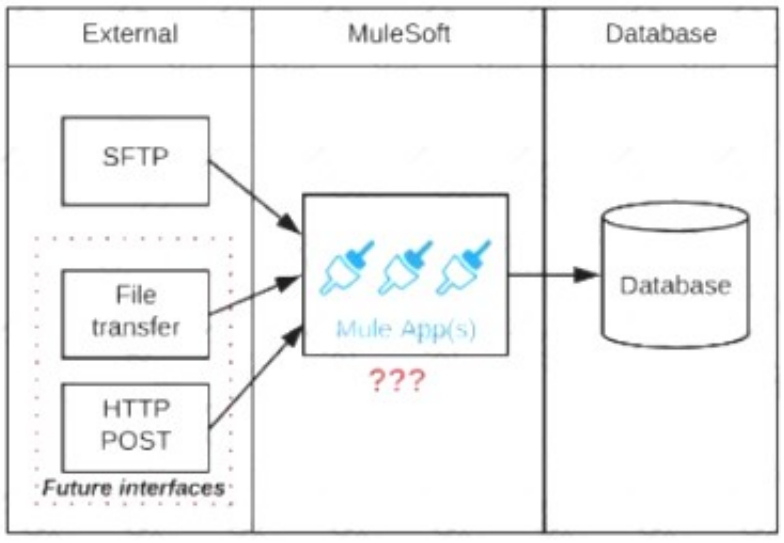

A business process involves the receipt of a file from an external vendor over SFTP. The file needs to be parsed and its content processed, validated, and ultimately persisted to a database. The delivery mechanism is expected to change in the future as more vendors send similar files using other mechanisms 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

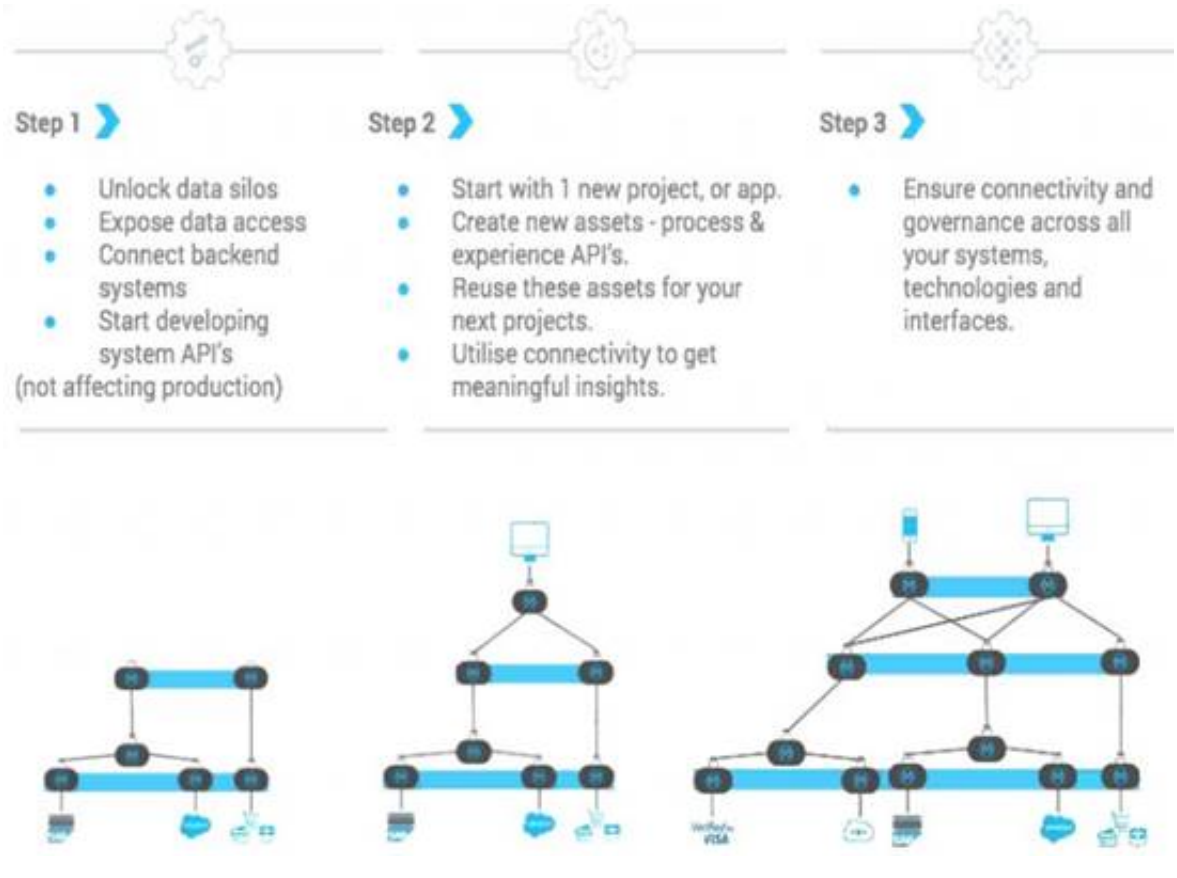

NEW QUESTION 7
Refer to the exhibit.

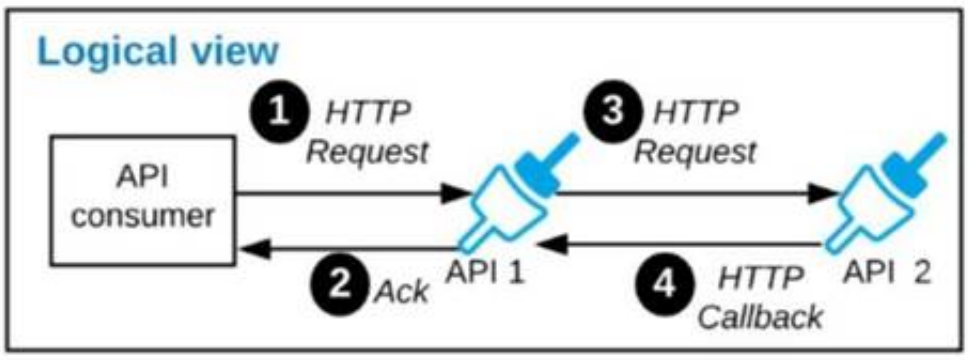

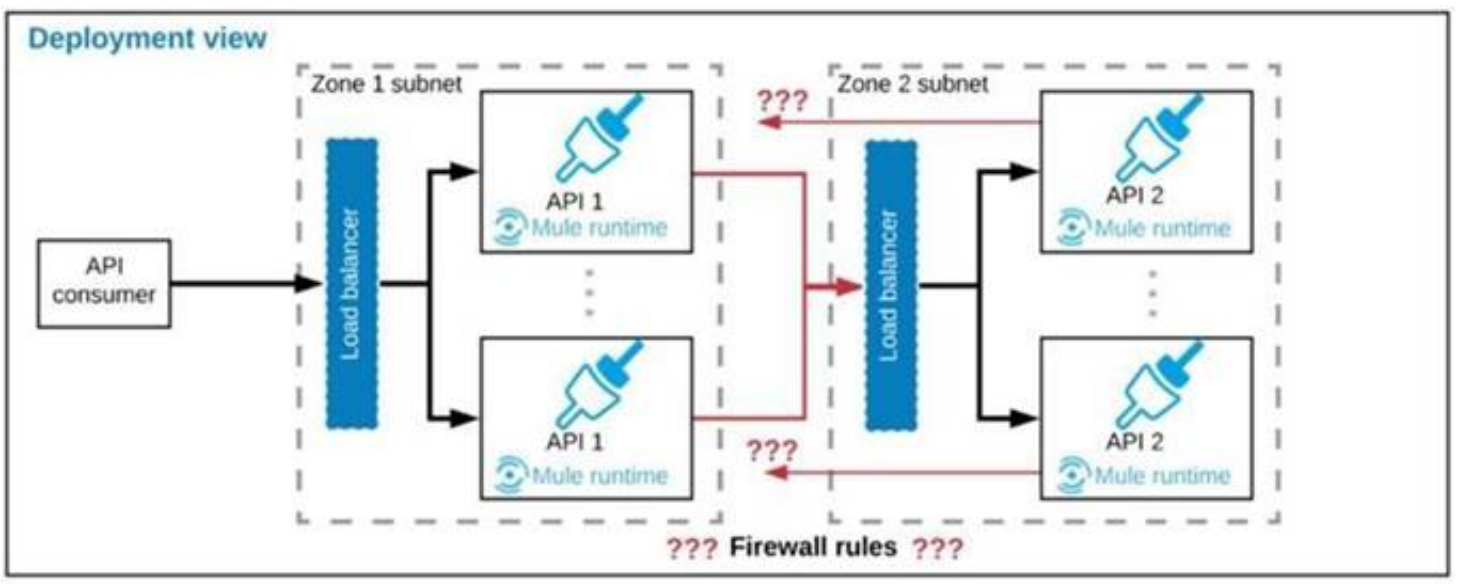

A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API l to notify about completion of the asynchronous process.
Each API Is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.
In the network architecture, how must the firewall rules be configured to enable the above Interaction between API 1 and API 2?

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

Answer: B

Explanation:
* If your API implementation involves putting a load balancer in front of your APIkit application, configure the load balancer to redirect URLs that reference the baseUri of the application directly. If the load balancer does not redirect URLs, any calls that reach the load balancer looking for the application do not reach their destination.
* When you receive incoming traffic through the load balancer, the responses will go out the same way. However, traffic that is originating from your instance will not pass through the load balancer. Instead, it is sent directly from the public IP address of your instance out to the Internet. The ELB is not involved in that scenario.
* The question says "each API is deployed to multiple redundant Mule runtimes", that seems to be a hint for self hosted Mule runtime cluster. Set Inbound allowed for the LB, outbound allowed for runtime to request out.
* Hence correct way is to enable communication from each API's Mule Runtimes and Network zone to the load balancer of the other API. Because communication is asynchronous one

NEW QUESTION 8
An insurance company is implementing a MuleSoft API to get inventory details from the two vendors. Due to network issues, the invocations to vendor applications are getting timed-out intermittently. But the transactions are successful upon reprocessing
What is the most performant way of implementing this requirement?

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

Answer: A

NEW QUESTION 9
An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.
What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

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

Answer: B

NEW QUESTION 10
A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should be published in a way that is guaranteed in normal situations and also minimizes duplicate delivery in less frequent failure scenarios.
The organizational firewall is configured to only allow outbound traffic on ports 80 and 443. Some external event consumers are within the organizational network, while others are located outside the firewall.
What Anypoint Platform service is most idiomatic (used for its intended purpose) for publishing these Mule events to all external consumers while addressing the desired reliability goals?

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

Answer: B

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

NEW QUESTION 10
Organization wants to achieve high availability goal for Mule applications in customer hosted runtime plane. Due to the complexity involved, data cannot be shared among of different instances of same Mule application. What option best suits to this requirement considering high availability is very much critical to the organization?

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

Answer: B

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

NEW QUESTION 13
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 18
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 21
What is a key difference between synchronous and asynchronous logging from Mule applications?

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

Answer: B

Explanation:
Types of logging:
A) Synchronous: The execution of thread that is processing messages is interrupted to wait for the log message to be fully handled before it can continue.
The execution of the thread that is processing your message is interrupted to wait for the log message to be fully output before it can continue
Performance degrades because of synchronous logging
Used when the log is used as an audit trail or when logging ERROR/CRITICAL messages
If the logger fails to write to disk, the exception would raise on the same thread that's currently processing the Mule event. If logging is critical for you, then you can rollback the transaction.
Chart, diagram Description automatically generated

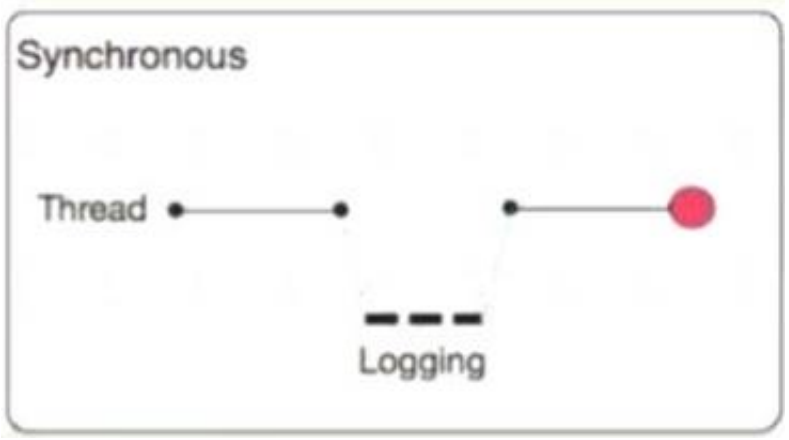

Chart, diagram, box and whisker chart Description automatically generated

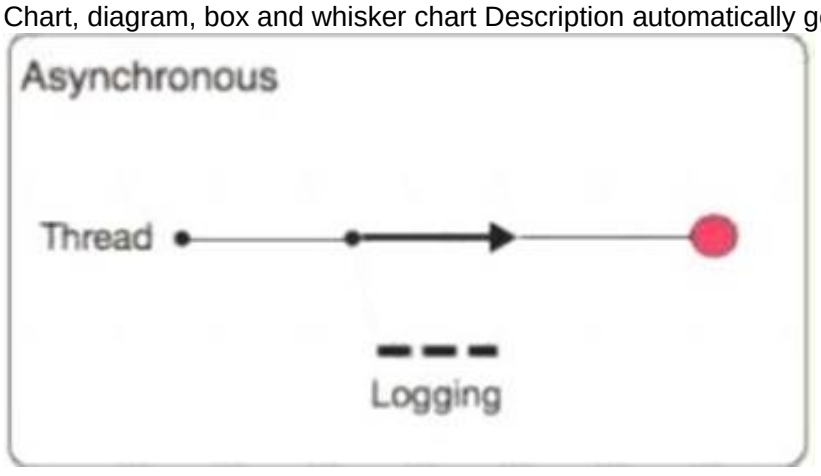

B) Asynchronous:
The logging operation occurs in a separate thread, so the actual processing of your message won't be delayed to wait for the logging to complete
Substantial improvement in throughput and latency of message processing Mule runtime engine (Mule) 4 uses Log4j 2 asynchronous logging by default The disadvantage of asynchronous logging is error handling.
If the logger fails to write to disk, the thread doing the processing won't be aware of any issues writing to the disk, so you won't be able to rollback anything.
Because the actual writing of the log gets differed, there's a chance that log messages might never make it to disk and get lost, if Mule were to crash before the buffers are flushed.

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

NEW QUESTION 23
Insurance organization is planning to deploy Mule application in MuleSoft Hosted runtime plane. As a part of requirement , application should be scalable . highly available. It also has regulatory requirement which demands logs to be retained for at least 2 years. As an Integration Architect what step you will recommend in order to achieve this?

A. It is not possible to store logs for 2 years in CloudHub deploymen
B. External log management system is required.
C. When deploying an application to CloudHub , logs retention period should be selected as 2 years
D. When deploying an application to CloudHub, worker size should be sufficient to store 2 years data
E. Logging strategy should be configured accordingly in log4j file deployed with the application.

Answer: A

Explanation:
Correct answer is It is not possible to store logs for 2 years in CloudHub deployment. External log management system is required. CloudHub has a specific log retention policy, as described in the documentation: the platform stores logs of up to 100 MB per app & per worker or for up to 30 days, whichever limit is hit first. Once this limit has been reached, the oldest log information is deleted in chunks and is irretrievably lost. The recommended approach is to persist your logs to a external logging system of your choice (such as Splunk, for instance) using a log appender. Please note that this solution results in the logs no longer being stored on our platform, so any support cases you lodge will require for you to provide the appropriate logs for review and case resolution

NEW QUESTION 25
As a part of project requirement, client will send a stream of data to mule application. Payload size can vary between 10mb to 5GB. Mule application is required to transform the data and send across multiple sftp servers. Due to the cost cuttings in the organization, mule application can only be allocated one worker with size of 0.2 vCore.
As an integration architect , which streaming strategy you would suggest to handle this scenario?

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

Answer: D

Explanation:
As the question says that data needs to be sent across multiple sftp serves , we cannot use non-repeatable streams. The non-repeatable strategy disables repeatable streams, which enables you to read an input stream only once.
You cant use in memory storage because with 0.2 vcore you will get only 1 GB of heap memory. Hence application will error out for file more than 1 GB.
Hence the correct option is file base repeatable stream

NEW QUESTION 26
A mule application uses an HTTP request operation to involve an external API. The external API follows the HTTP specification for proper status code usage.
What is possible cause when a 3xx status code is returned to the HTTP Request operation from the external API?

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

Answer: B

Explanation:
3xx HTTP status codes indicate a redirection that the user agent (a web browser or a crawler) needs to take further action when trying to access a particular resource.

NEW QUESTION 28
An organization is designing the following two Mule applications that must share data via a common persistent object store instance:
- Mule application P will be deployed within their on-premises datacenter.
- Mule application C will run on CloudHub in an Anypoint VPC.
The object store implementation used by CloudHub is the Anypoint Object Store v2 (OSv2).
what type of object store(s) should be used, and what design gives both Mule applications access to the same object store instance?

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

Answer: C

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

NEW QUESTION 30

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 33
Which Salesforce API is invoked to deploy, retrieve, create or delete customization information such as custom object definitions using a Mule Salesforce connector in a Mule application?

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

Answer: B

NEW QUESTION 36
A company is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTP5 POST and must be acknowledged immediately.
Once acknowledged the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to the rejections from the back-end system will need to be processed manually (outside the banking system).
The mule application will be deployed to a customer hosted runtime and will be able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization's firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.
Which combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

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

Answer: B

NEW QUESTION 37
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/ip-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
Refer to the exhibit.

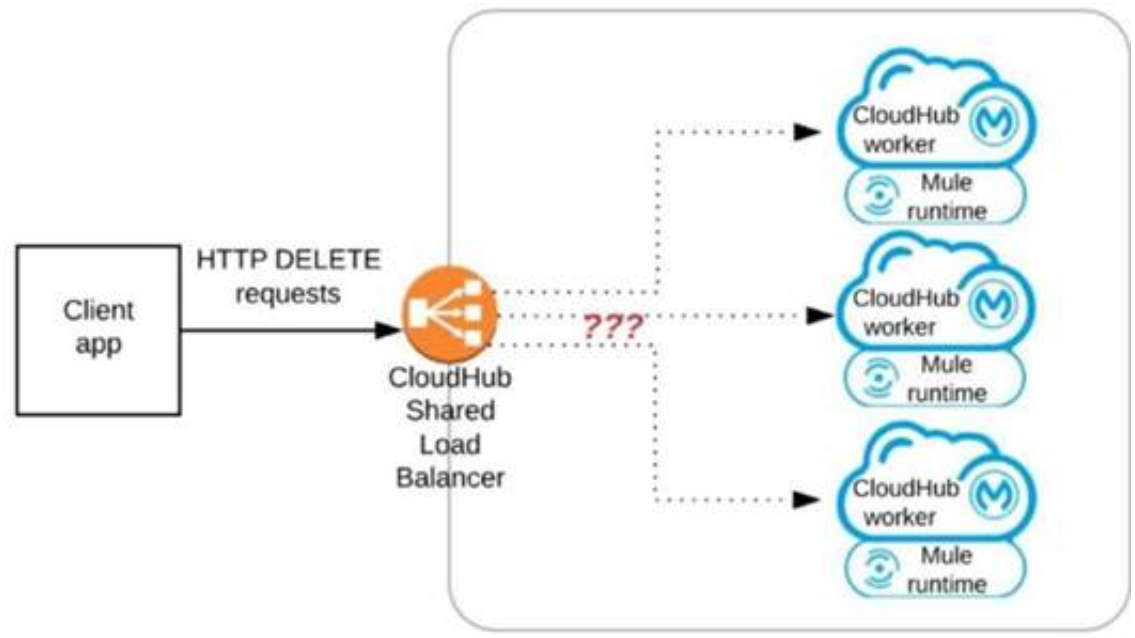

A Mule application has an HTTP Listener that accepts HTTP DELETE requests. This Mule application Is deployed to three CloudHub workers under the control of the CloudHub Shared Load Balancer.
A web client makes a sequence of requests to the Mule application's public URL.
How is this sequence of web client requests distributed among the HTTP Listeners running in the three CloudHub workers?

A. Each request is routed to the PRIMARY CloudHub worker in the PRIMARY Availability Zone (AZ)
B. Each request is routed to ONE ARBiTRARY CloudHub worker in the PRIMARY Availability Zone (AZ)
C. Each request Is routed to ONE ARBiTRARY CloudHub worker out of ALL three CloudHub workers
D. Each request is routed (scattered) to ALL three CloudHub workers at the same time

Answer: C

Explanation:
Correct behavior is Each request is routed to ONE ARBITRARY CloudHub worker out of ALL three CloudHub workers

NEW QUESTION 41
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 44
An organization uses one specific CloudHub (AWS) region for all CloudHub deployments. How are CloudHub workers assigned to availability zones (AZs) when the organization's Mule applications are deployed to CloudHub in that region?

A. Workers belonging to a given environment are assigned to the same AZ within that region.
B. AZs are selected as part of the Mule application's deployment configuration.
C. Workers are randomly distributed across available AZs within that region.
D. An AZ is randomly selected for a Mule application, and all the Mule application's CloudHub workers are assigned to that one AZ

Answer: C

Explanation:
Correct answer is Workers are randomly distributed across available AZs within that region. This ensure high availability for deployed mule applications Mulesoft documentation reference :
https://docs.mulesoft.com/runtime-manager/cloudhub-hadr

NEW QUESTION 48
A mule application is required to periodically process large data set from a back-end database to Salesforce CRM using batch job scope configured properly process the higher rate of records.
The application is deployed to two cloudhub workers with no persistence queues enabled. What is the consequence if the worker crashes during records processing?

A. Remaining records will be processed by a new replacement worker
B. Remaining records be processed by second worker
C. Remaining records will be left and processed
D. All the records will be processed from scratch by the second worker leading to duplicate processing

Answer: C

NEW QUESTION 53
A company is implementing a new Mule application that supports a set of critical functions driven by a rest API enabled, claims payment rules engine hosted on oracle ERP. As designed the mule application requires many data transformation operations as it performs its batch processing logic.
The company wants to leverage and reuse as many of its existing java-based capabilities (classes, objects, data model etc.) as possible
What approach should be considered when implementing required data mappings and transformations between Mule application and Oracle ERP in the new Mule application?

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

Answer: C

NEW QUESTION 54
An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?

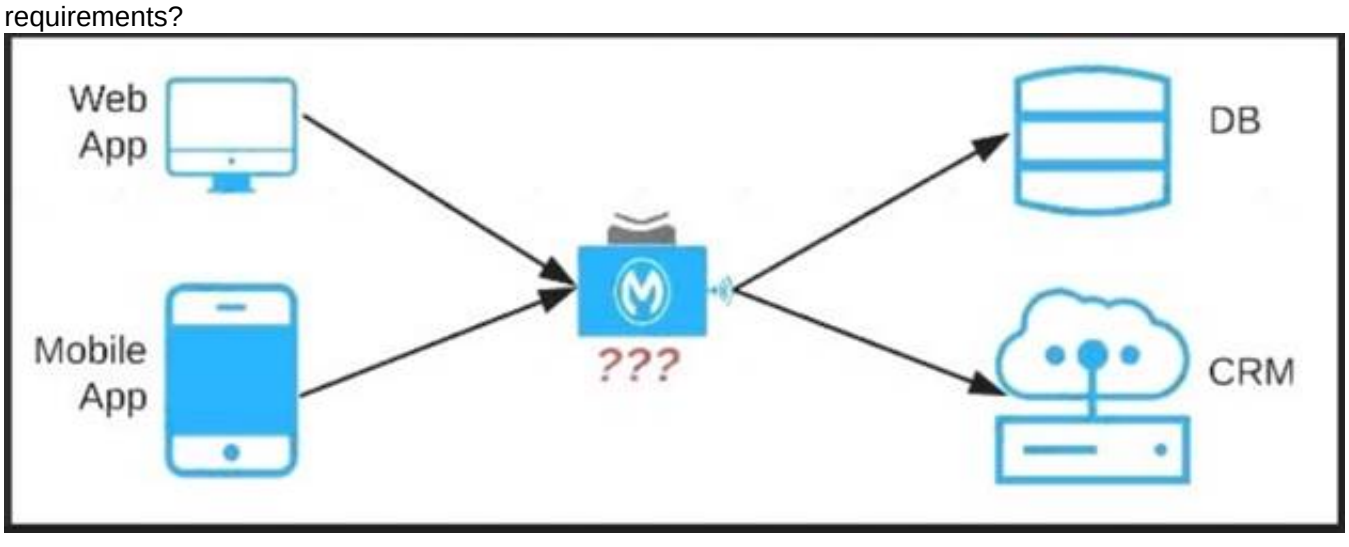

A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
B. One set of APIs (Experience API, Process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

Answer: C

Explanation:
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API
Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.
Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.
Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system
A screenshot of a computer Description automatically generated with low confidence

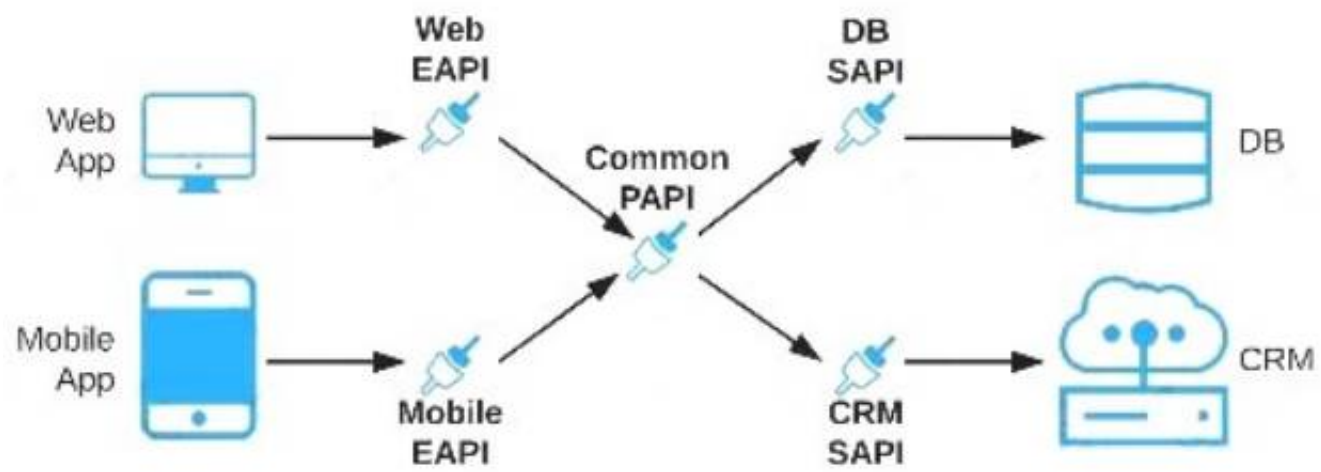

Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API
Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.
Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.
Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system
Diagram Description automatically generated

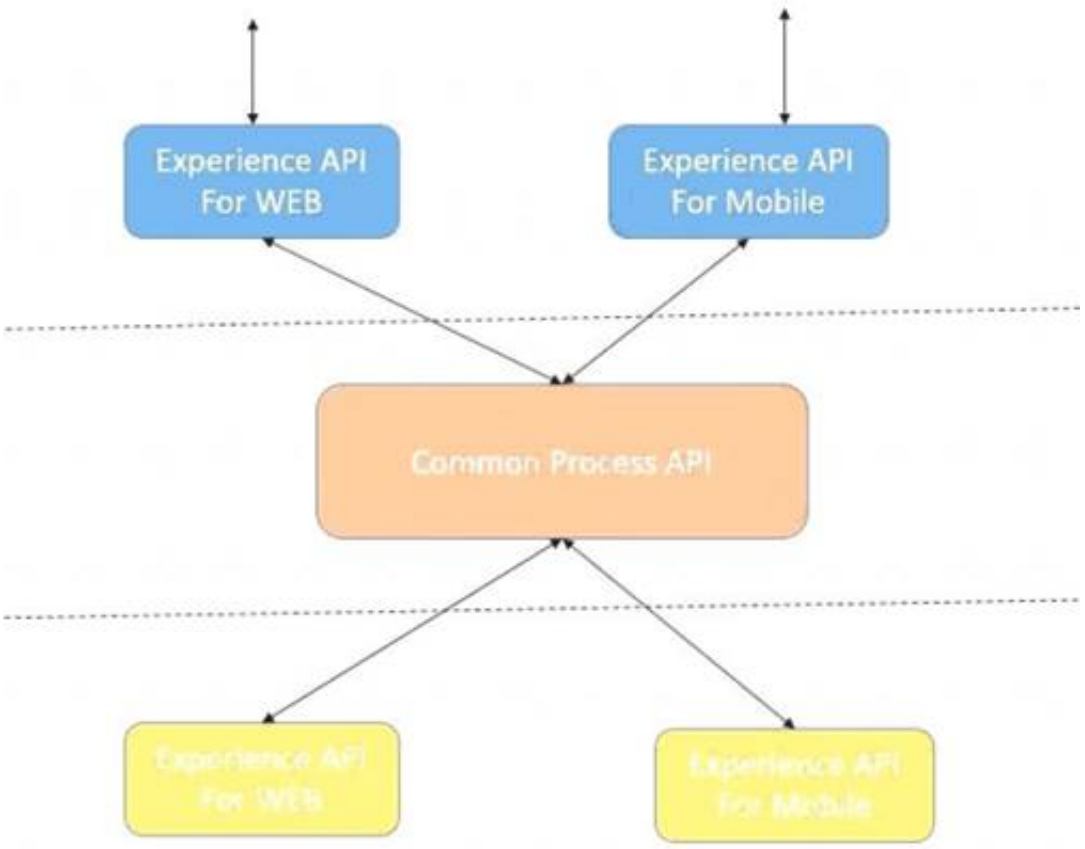

NEW QUESTION 58
A project team is working on an API implementation using the RAML definition as a starting point. The team has updated the definition to include new operations and has published a new version to exchange. Meanwhile another team is working on a mule application consuming the same API implementation.
During the development what has to be performed by the mule application team to take advantage of the newly added operations?

A. Scaffold the client application with the new definition
B. Scaffold API implementation application with the new definition
C. Update the REST connector from exchange in the client application
D. Update the API connector in the API implementation and publish to exchange

Answer: C

NEW QUESTION 62
An organization has several APIs that accept JSON data over HTTP POST. The APIs are all publicly available and are associated with several mobile applications and web applications. The organization does NOT want to use any authentication or compliance policies for these APIs, but at the same time, is worried that some bad actor could send payloads that could somehow compromise the applications or servers running the API implementations. What out-of-the-box Anypoint Platform policy can address exposure to this threat?

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

Answer: D

Explanation:
We need to note few things about the scenario which will help us in reaching the correct solution.
Point 1 : The APIs are all publicly available and are associated with several mobile applications and web applications. This means Apply an IP blacklist policy is not viable option. as blacklisting IPs is limited to partial web traffic. It can't be useful for traffic from mobile application
Point 2 : The organization does NOT want to use any authentication or compliance policies for these APIs. This means we can not apply HTTPS mutual authentication scheme.
Header injection or removal will not help the purpose.
By its nature, JSON is vulnerable to JavaScript injection. When you parse the JSON object, the malicious code inflicts its damages. An inordinate increase in the size and depth of the JSON payload can indicate injection. Applying the JSON threat protection policy can limit the size of your JSON payload and thwart recursive additions to the JSON hierarchy.
Hence correct answer is Apply a JSON threat protection policy to all APIs to detect potential threat vectors

NEW QUESTION 63
A leading bank implementing new mule API.
The purpose of API to fetch the customer account balances from the backend application and display them on the online platform the online banking platform. The online banking platform will send an array of accounts to Mule API get the account balances.
As a part of the processing the Mule API needs to insert the data into the database for auditing purposes and this process should not have any performance related implications on the account balance retrieval flow
How should this requirement be implemented to achieve better throughput?

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

Answer: D

NEW QUESTION 64
Refer to the exhibit.

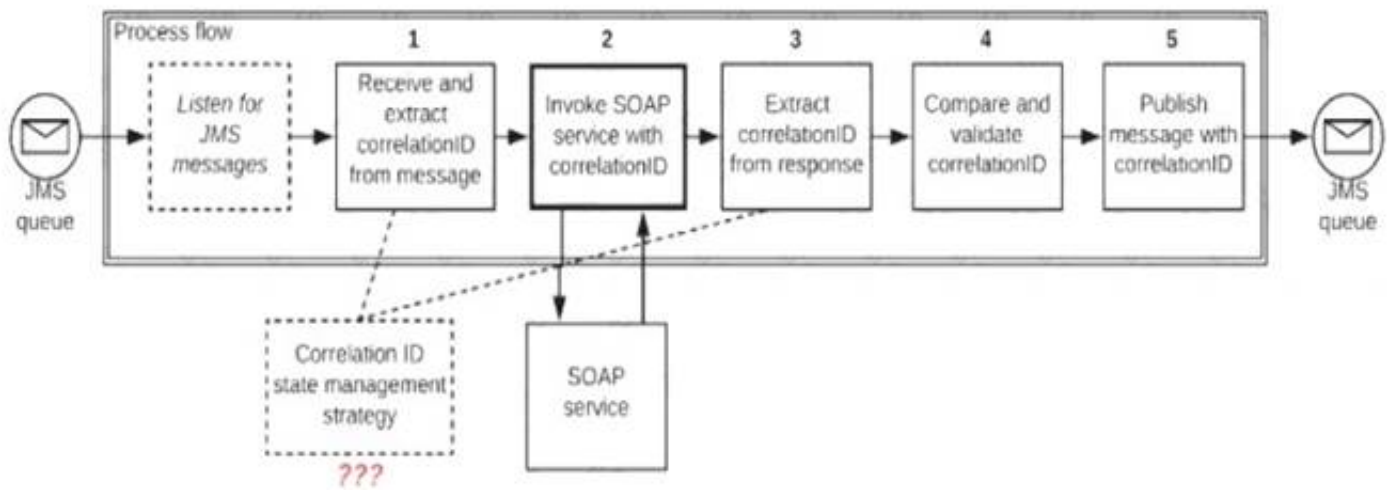

A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to receive JMS messages from a JMS queue. To process each received JMS message, the following steps are performed in a flow:
Step l: The JMS Correlation ID header is read from the received JMS message.
Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.
Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.
Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.
Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.
Where should the Mule application store the JMS Correlation ID values received in Step 1 and Step 3 so that the validation in Step 4 can be performed, while also making the overall Mule application highly available, fault-tolerant, performant, and maintainable?

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

Answer: C

Explanation:

* If we store Correlation id value in step 1 as Mule event variables/attributes, the values will be cleared after server restart and we want system to be fault tolerant.
* The Correlation ID value in Step 1 should be stored in a persistent object store.
* We don't need to store Correlation ID value in Step 3 to persistent object store. We can store it but as we also need to make application performant. We can avoid this step of accessing persistent object store.
* Accessing persistent object stores slow down the performance as persistent object stores are by default stored in shared file systems.
* As the SOAP service is idempotent in nature. In case of any failures , using this Correlation ID saved in first step we can make call to SOAP service and validate the Correlation ID.

Top of Form

Additional Information:

* Competing Consumers

are multiple consumers that are all created to receive messages from a single
Point-to-Point Channel. When the channel delivers a message, any of the consumers could potentially receive it. The messaging system's implementation determines which consumer actually receives the message, but in effect the consumers compete with each other to be the receiver. Once a consumer receives a message, it can delegate to the rest of its application to help process the message.

Diagram Description automatically generated

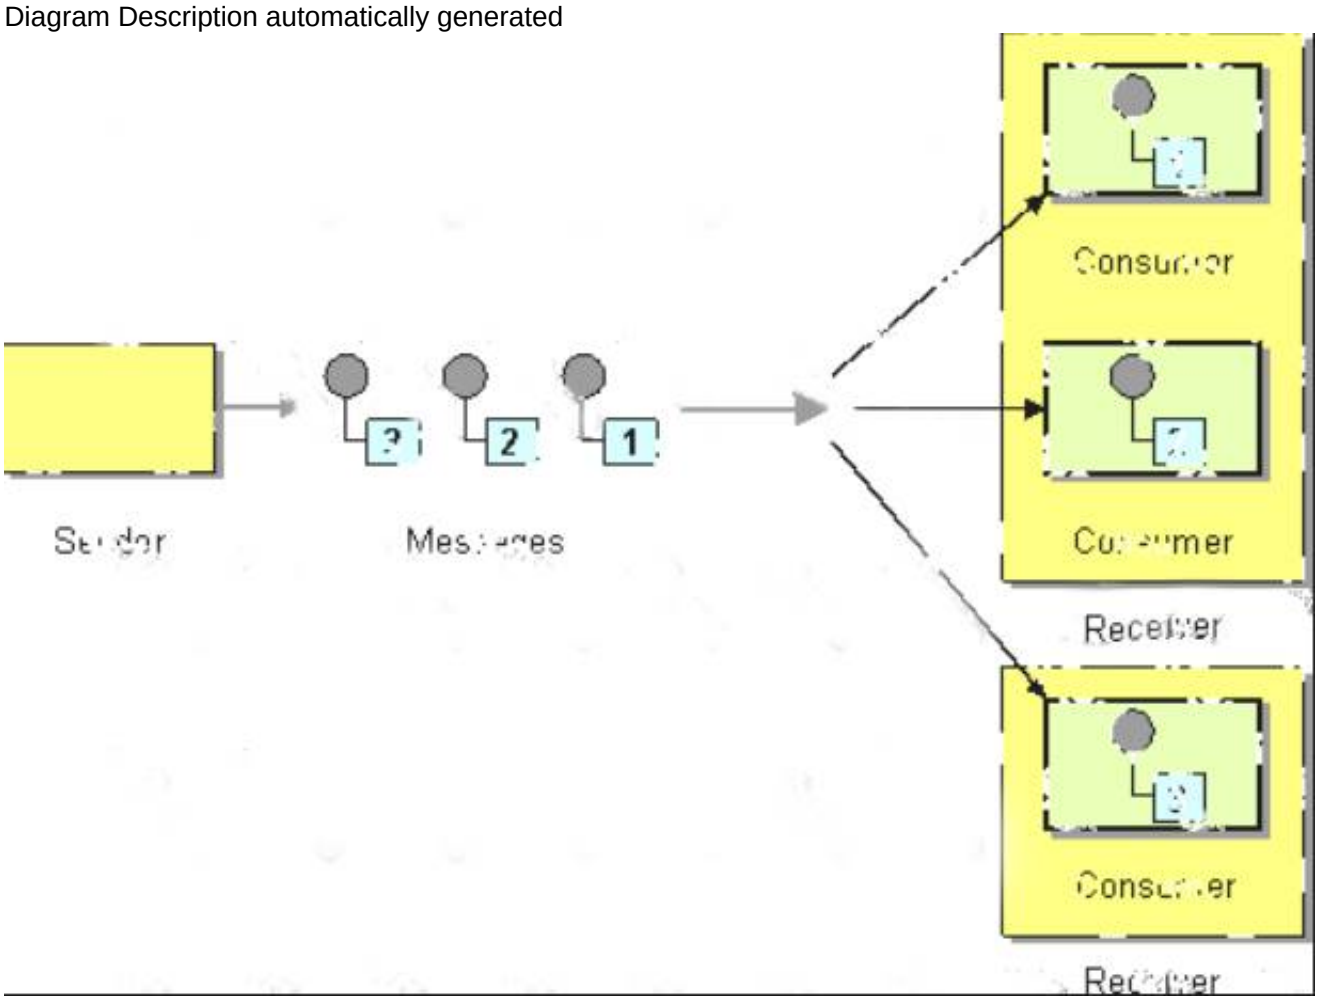

* In case you are unaware about term idempotent re is more info:

Idempotent operations means their result will always same no matter how many times these operations are invoked.

Table Description automatically generated

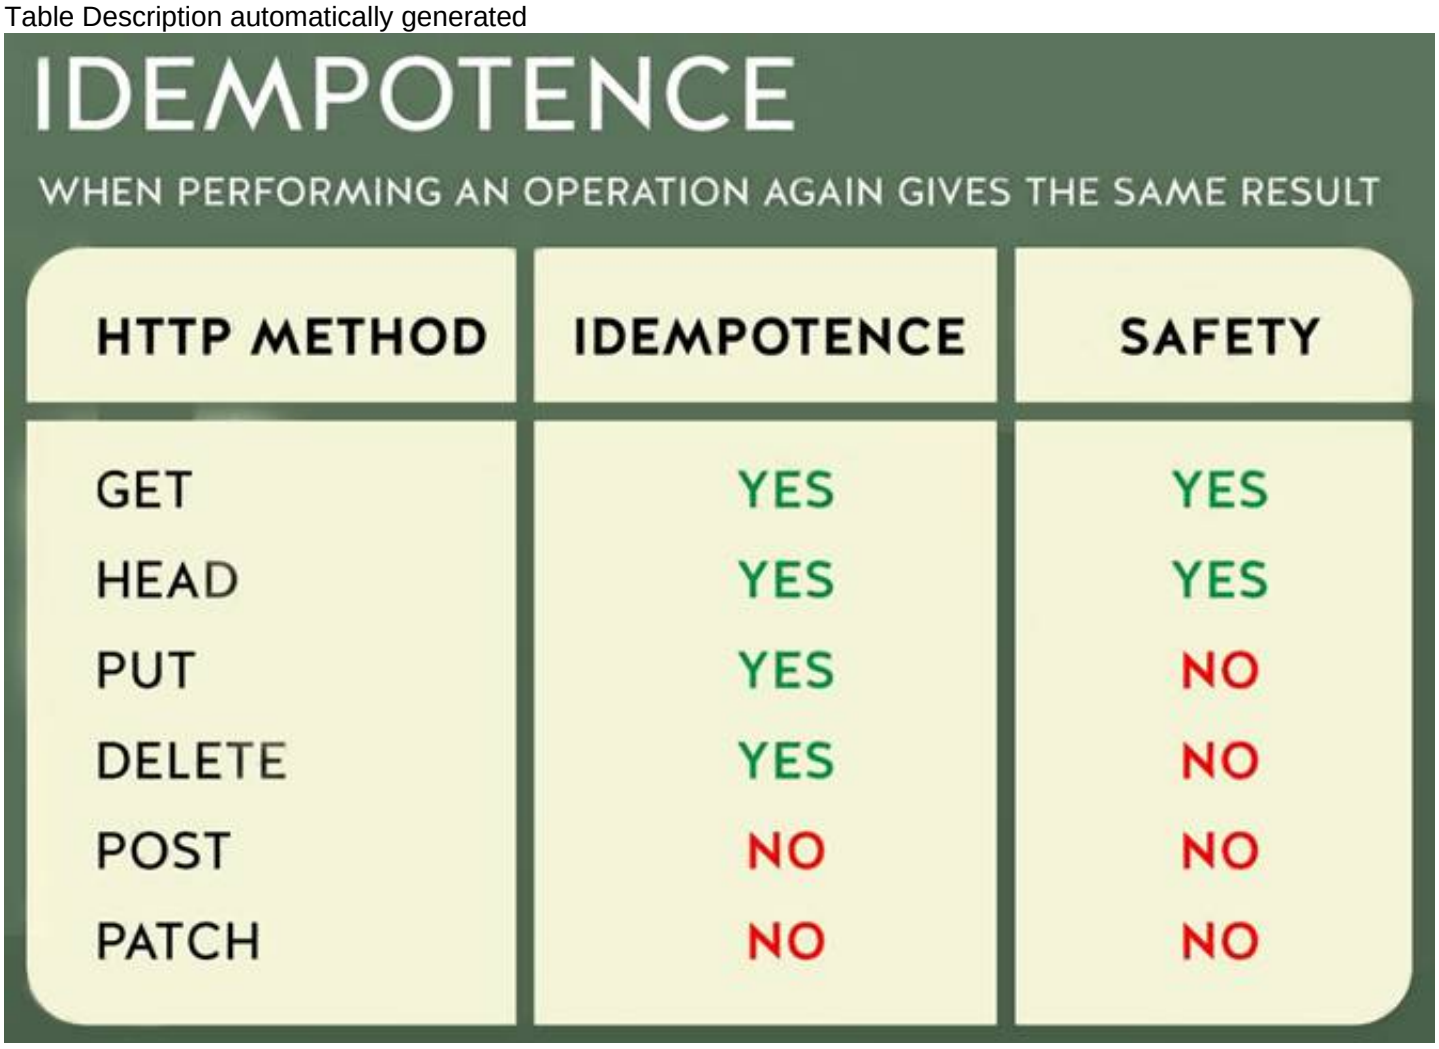

Bottom of Form

NEW QUESTION 65

A Mule application name Pub uses a persistence object store. The Pub Mule application is deployed to Cloudhub and it configured to use Object Store v2. Another Mule application name sub is being developed to retrieve values from the Pub Mule application persistence object Store and will also be deployed to cloudhub.

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

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

Answer: D

Explanation:
* The Object Store V2 API enables API access to Anypoint Platform Object Store v2.
* You can configure a Mule app to use the Object Store REST API to store and retrieve values from an object store in another Mule app. However, Object Store v2 is not designed for app-to-app communication. To share data between two Mule4 apps, use a queue in Anypoint MQ.
* The Object Store v2 APIs enable you to use REST to perform the following:
- Retrieve a list of object stores and keys associated with an application.
- Store and retrieve key-value pairs in an object store.
- Delete key-value pairs from an object store.
- Retrieve Object Store usage statistics for your organization.
- Object Store provides these APIs: Object Store API
Object Store Stats API

NEW QUESTION 69
Refer to the exhibit.

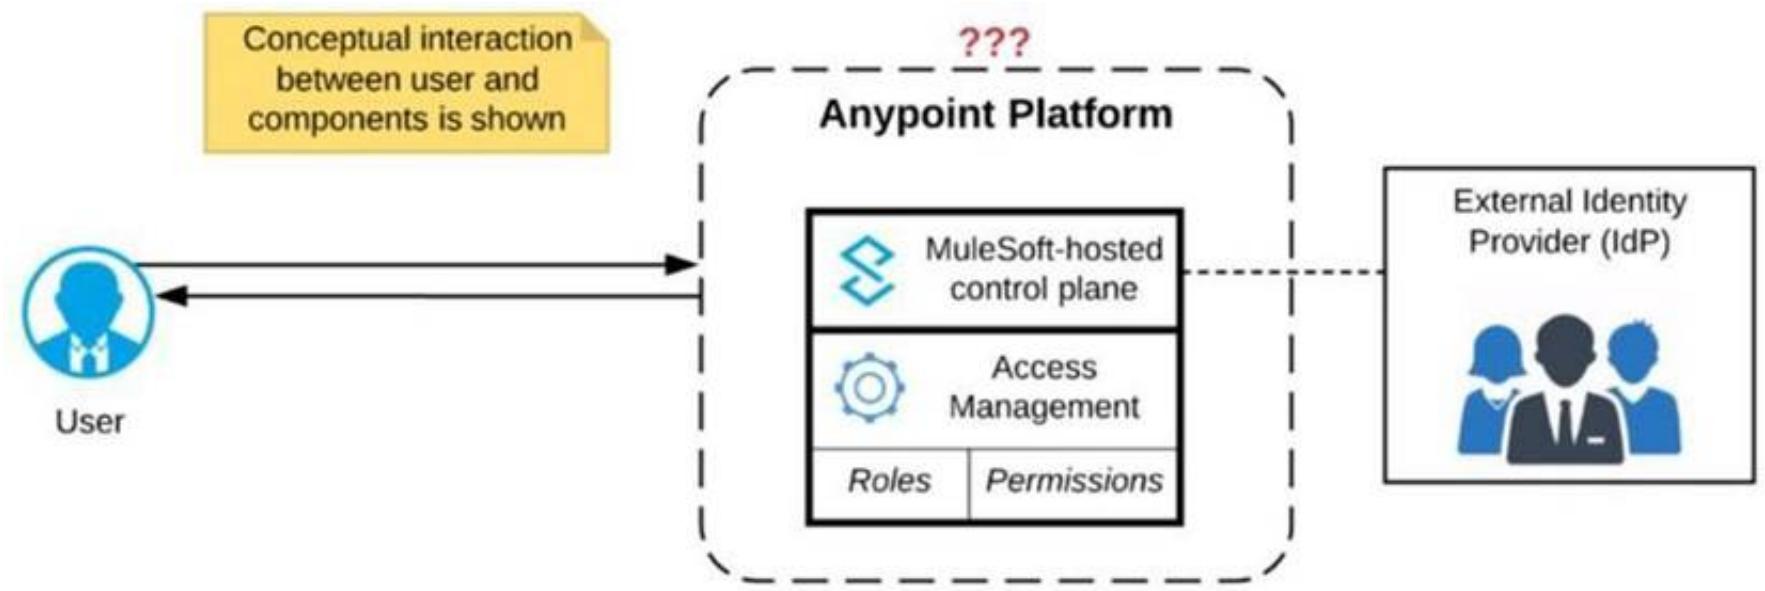

Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.
What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled via the external Identity Provider?

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

Answer: B

Explanation:
* By default, Anypoint Platform performs its own user management
– For user management, one external IdP can be integrated with the Anypoint Platform organization (note: not at business group level)
– Permissions and access control are still enforced inside Anypoint Platform and CANNOT be controlled via the external Identity Provider * As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). * You can map users in a federated organization’s group to a role which also gives the flexibility of controlling the business group within Anypoint Platform to which the user belongs to.
Also user can nbe removed from external identity
management system when they no longer work for the organization. So they wont be able to authenticate using SSO to login to Anypoint Platform. * Using external identity we can no change permissions of a particular role in Mulesoft Anypoint platform.
* So Correct answer is Assigning Anypoint Platform permissions to a role

NEW QUESTION 70
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

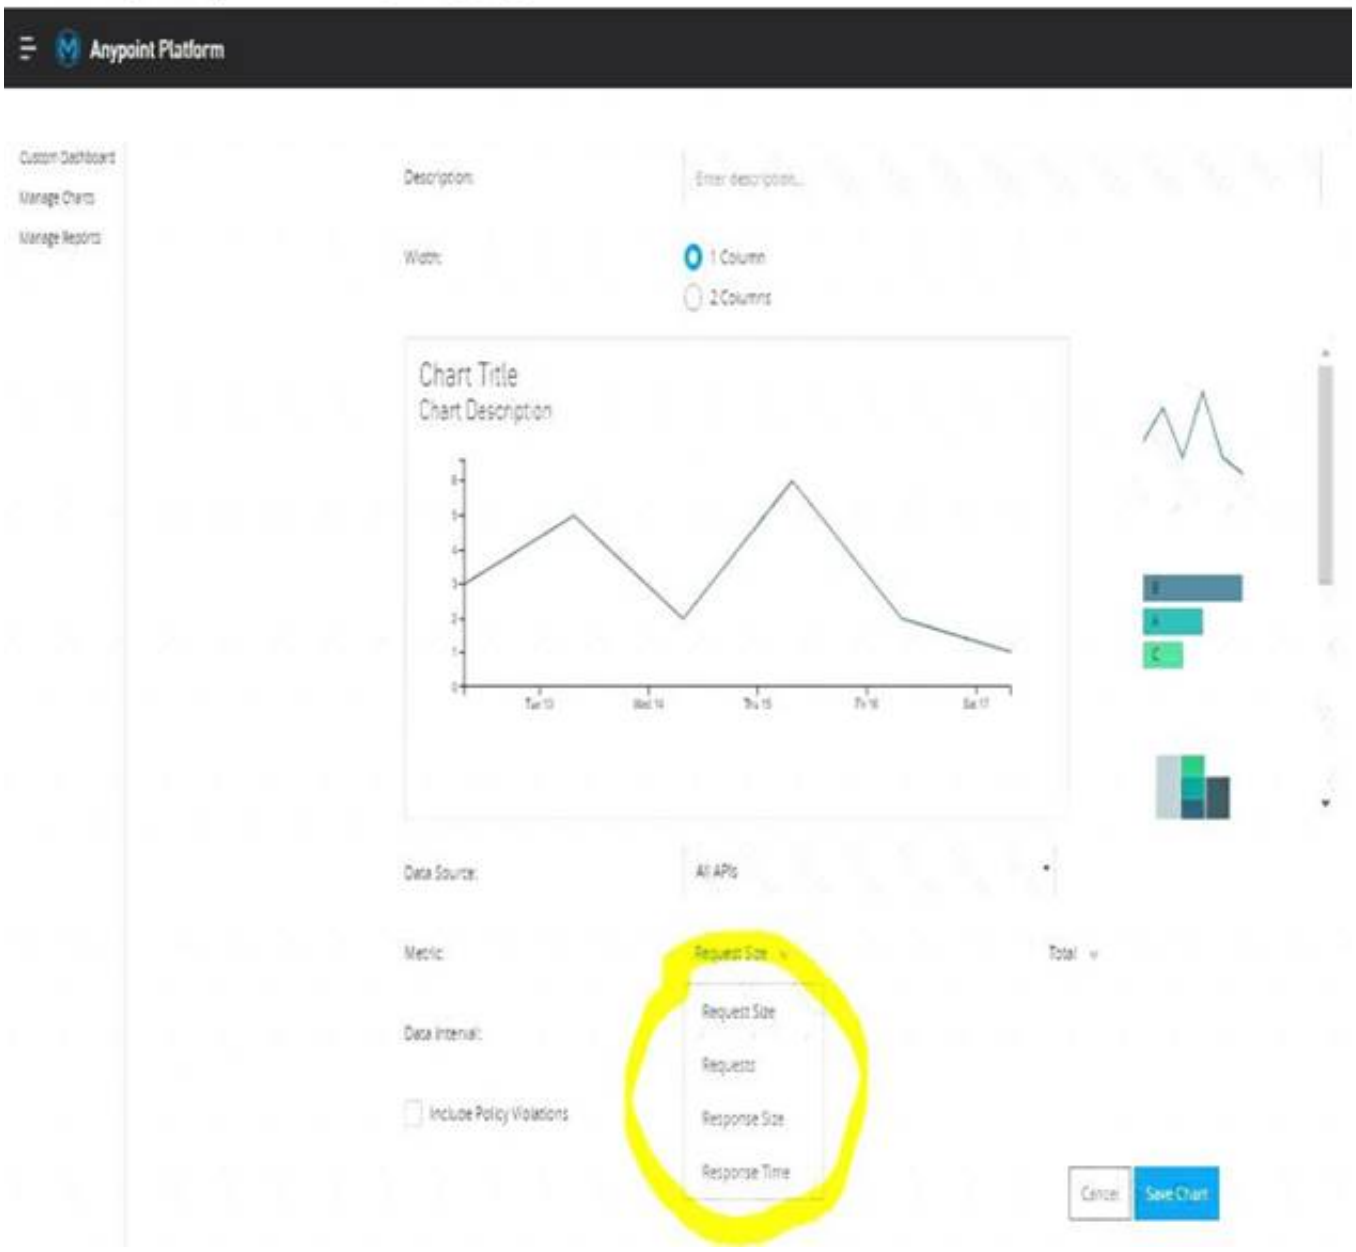

NEW QUESTION 71
When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

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 73
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 76
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 80
An airline is architecting an API connectivity project to integrate its flight data into an online aggregation website. The interface must allow for secure communication high-performance and asynchronous message exchange.
What are suitable interface technologies for this integration assuming that Mulesoft fully supports these technologies and that Anypoint connectors exist for these interfaces?

A. AsyncAPI over HTTPSAMQP with RabbitMQ JSON/REST over HTTPS
B. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
C. CSV over FTP YAM L over TLS JSON over HTTPS
D. SOAP over HTTPS HOP over TLS gRPC over HTTPS

Answer: A

NEW QUESTION 82
An auto mobile company want to share inventory updates with dealers Dl and D2 asynchronously and concurrently via queues Q1 and Q2. Dealer Dl must consume the message from the queue Q1 and dealer D2 to must consume a message from the queue Q2.
Dealer D1 has implemented a retry mechanism to reprocess the transaction in case of any errors while processing the inventers updates. Dealer D2 has not implemented any retry mechanism.

How should the dealers acknowledge the message to avoid message loss and minimize impact on the current implementation?

A. Dealer D1 must use auto acknowledgement and dealer D2 can use manual acknowledgement and acknowledge the message after successful processing
B. Dealer D1 can use auto acknowledgement and dealer D2 can use IMMEDIATE acknowledgement and acknowledge the message of successful processing
C. Dealer D1 and dealer D2 must use AUTO acknowledgement and acknowledge the message after successful processing
D. Dealer D1 can use AUTO acknowledgement and dealer D2 must use manual acknowledgement and acknowledge the message after successful processing

Answer: D

NEW QUESTION 87
In one of the critical payment related mule application, transaction is being used . As an enhancement to implementation , scatter gather route is introduced which is also the part of transaction group. Scatter gather route has 4 routes.
What will be the behavior of the Mule application in case of error occurs in 4th route of the scatter-gather router and transaction needs to be rolled back?

A. Only errored route will be rolled back
B. All routes will be rolled back
C. Scatter Gather router cannot be part of transaction

Answer: B

Explanation:
•Scatter Gather: When running within a transaction, Scatter Gather does not execute in parallel. This means that the second route is executed after the first one is processed, the third after the second one, etc. In case of error, all routes will be rolled back

NEW QUESTION 88
An organization is evaluating using the CloudHub shared Load Balancer (SLB) vs creating a CloudHub dedicated load balancer (DLB). They are evaluating how this choice affects the various types of certificates used by CloudHub deplpoyed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates.
What type of restrictions exist on the types of certificates that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

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

Answer: A

Explanation:
https://docs.mulesoft.com/runtime-manager/dedicated-load-balancer-tutorial

NEW QUESTION 93
An API client is implemented as a Mule application that includes an HTTP Request operation using a default configuration. The HTTP Request operation invokes an external API that follows standard HTTP status code conventions, which causes the HTTP Request operation to return a 4xx status code.
What is a possible cause of this status code response?

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

Answer: D

Explanation:
Correct choice is: "The external API reported an error with the HTTP request that was received from the outbound HTTP Request operation of the Mule application"
Understanding HTTP 4XX Client Error Response Codes : A 4XX Error is an error that arises in cases where there is a problem with the user's request, and not with the server.
Such cases usually arise when a user's access to a webpage is restricted, the user misspells the URL, or when a webpage is nonexistent or removed from the public's view.
In short, it is an error that occurs because of a mismatch between what a user is trying to access, and its availability to the user — either because the user does not have the right to access it, or because what the user is trying to access simply does not exist. Some of the examples of 4XX errors are
400 Bad Request The server could not understand the request due to invalid syntax. 401 Unauthorized Although the HTTP standard specifies "unauthorized", semantically this response means "unauthenticated". That is, the client must authenticate itself to get the requested response. 403 Forbidden The client does not have access rights to the content; that is, it is unauthorized, so the server is refusing to give the requested resource. Unlike 401, the client's identity is known to the server. 404 Not Found The server can not find the requested resource. In the browser, this means the URL is not recognized. In an API, this can also mean that the endpoint is valid but the resource itself does not exist. Servers may also send this response instead of 403 to hide the existence of a resource from an unauthorized client. This response code is probably the most famous one due to its frequent occurrence on the web. 405 Method Not Allowed The request method is known by the server but has been disabled and cannot be used. For example, an API may forbid DELETE-ing a resource. The two mandatory methods, GET and HEAD, must never be disabled and should not return this error code. 406 Not Acceptable This response is sent when the web server, after performing server-driven content negotiation, doesn't find any content that conforms to the criteria given by the user agent. The external API reported that the API implementation has moved to a different external endpoint cannot be the correct answer as in this situation 301 Moved Permanently The URL of the requested resource has been changed permanently. The new URL is given in the response.
----------In Lay man's term the scenario would be: API CLIENT —> MuleSoft API - HTTP request "Hey, API.. process this" —> External API API CLIENT <– MuleSoft API - http response "I'm sorry Client.. something is wrong with that request" <– (4XX) External API

NEW QUESTION 97

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

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

Answer: B

Explanation:
Correct answer is Only MuleSoft provided certificates can be used for server side certificate
* The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
* You would need to use dedicated load balancer which can enable you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
* To use a dedicated load balancer in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.
Additional Info on SLB Vs DLB:
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/ip-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 98
A corporation has deployed multiple mule applications implementing various public and private API's to different cloudhub workers. These API's arc Critical applications that must be highly available and in line with the reliability SLA as defined by stakeholders.
How can API availability (liveliness or readiness) be monitored so that Ops team receives outage notifications?

A. Enable monitoring of individual applications from Anypoint monitoring
B. Configure alerts with failure conditions in runtime manager
C. Configure alerts failure conditions in API manager
D. Use any point functional monitoring test API's functional behavior

Answer: A

NEW QUESTION 101
What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

A. To send Mule application log entries to Splunk
B. To change tog4j2 tog levels in Anypoint Runtime Manager without having to restart the Mule application
C. To log certain messages to a custom log category
D. To directly reference one shared and customized log4j2.xml file from multiple Mule applications

Answer: D

Explanation:
* Correct answer is To directly reference one shared and customized log4j2.xml file from multiple Mule applications. Key word to note in the answer is directly.
* By default, CloudHub replaces a Mule application's log4j2.xml file with a CloudHub log4j2.xml file. This specifies the CloudHub appender to write logs to the CloudHub logging service.
* You cannot modify CloudHub log4j2.xml file to add any custom appender. But there is a process in order to achieve this. You need to raise a request on support portal to disable CloudHub provided Mule application log4j2 file.
Graphical user interface, application, Word Description automatically generated

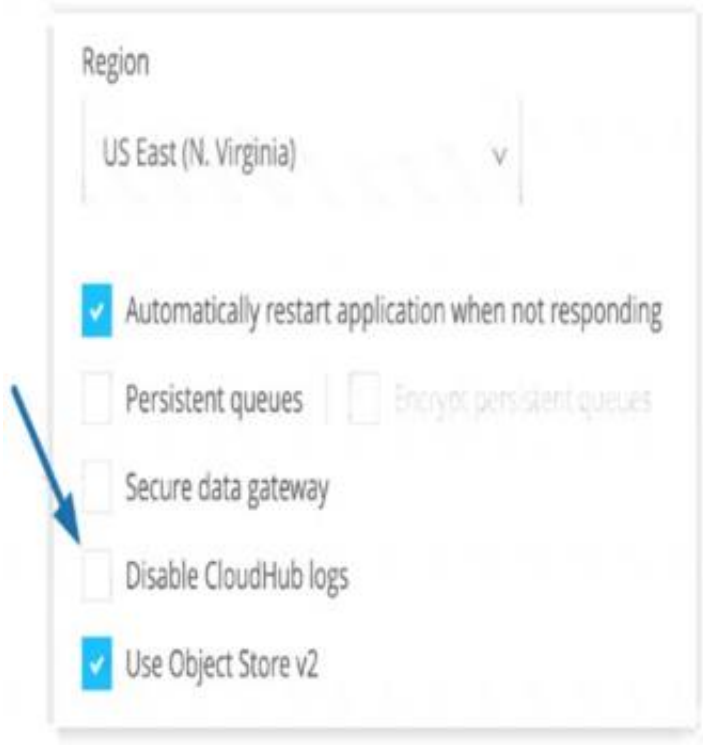

* Once this is done , Mule application's log4j2.xml file is used which you can use to send/export application logs to other log4j2 appenders, such as a custom logging system MuleSoft does not own any responsibility for lost logging data due to misconfiguration of your own log4j appender if it happens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

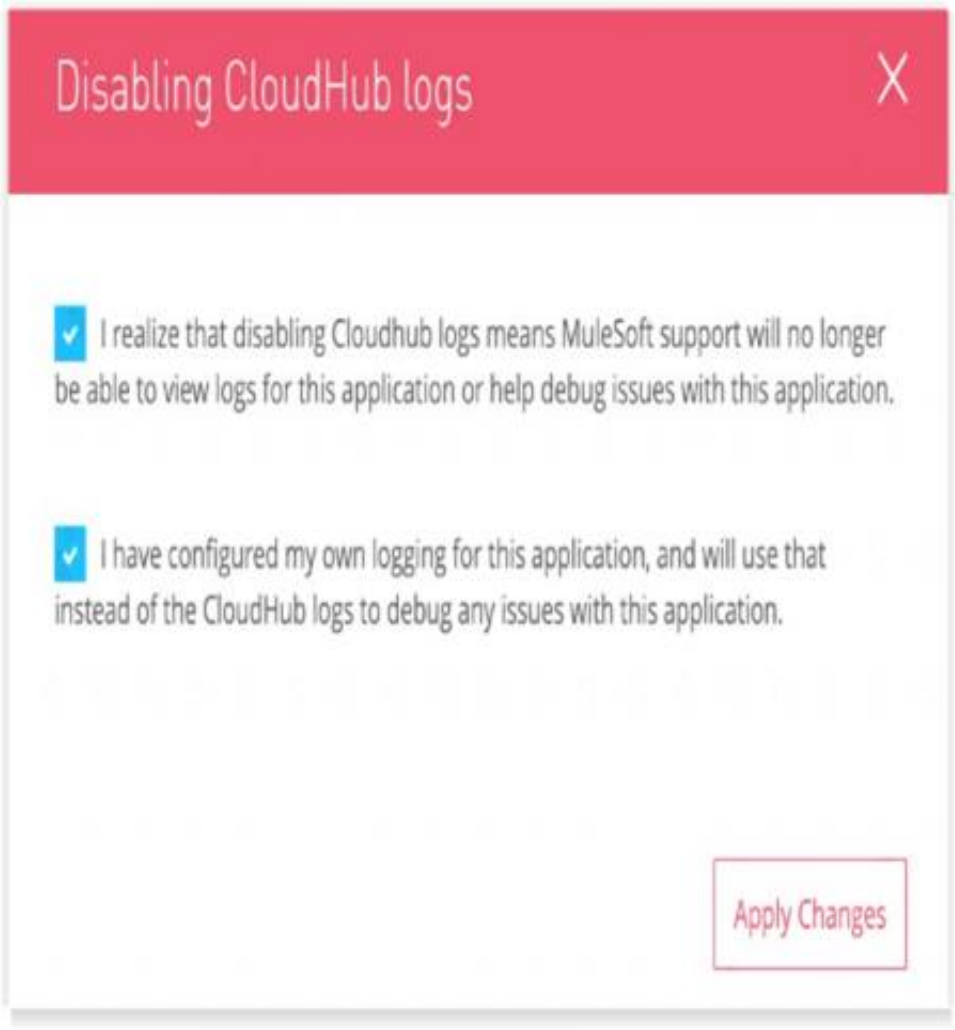

* One more difference between customer-hosted Mule runtimes and CloudHub deployed mule instance is that
- CloudHub system log messages cannot be sent to external log management system without installing custom CH logging configuration through support
- where as Customer-hosted runtime can send system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

NEW QUESTION 103
A Mule application contains a Batch Job with two Batch Steps (Batch_Step_l and Batch_Step_2). A payload with 1000 records is received by the Batch Job.
How many threads are used by the Batch Job to process records, and how does each Batch Step process records within the Batch Job?

A. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and RECORDS are processed IN PARALLEL within and between the two Batch Steps
B. Each Batch Job uses a SINGLE THREAD for all Batch steps Each Batch step instance receives ONE record at a time as the payload, and RECORDS are processed IN ORDER, first through Batch_Step_l and then through Batch_Step_2

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

Answer: A

Explanation:
* Each Batch Job uses SEVERAL THREADS for the Batch Steps
* Each Batch Step instance receives ONE record at a time as the payload. It's not received in a block, as it does not wait for multiple records to be completed before moving to next batch step. (So Option D is out of choice)
* RECORDS are processed IN PARALLEL within and between the two Batch Steps.
* RECORDS are not processed in order. Let's say if second record completes batch_step_1 before record 1, then it moves to batch_step_2 before record 1. (So option C and D are out of choice)
* A batch job is the scope element in an application in which Mule processes a message payload as a batch of records. The term batch job is inclusive of all three phases of processing: Load and Dispatch, Process, and On Complete.
* A batch job instance is an occurrence in a Mule application whenever a Mule flow executes a batch job. Mule creates the batch job instance in the Load and Dispatch phase. Every batch job instance is identified internally using a unique String known as batch job instance id.

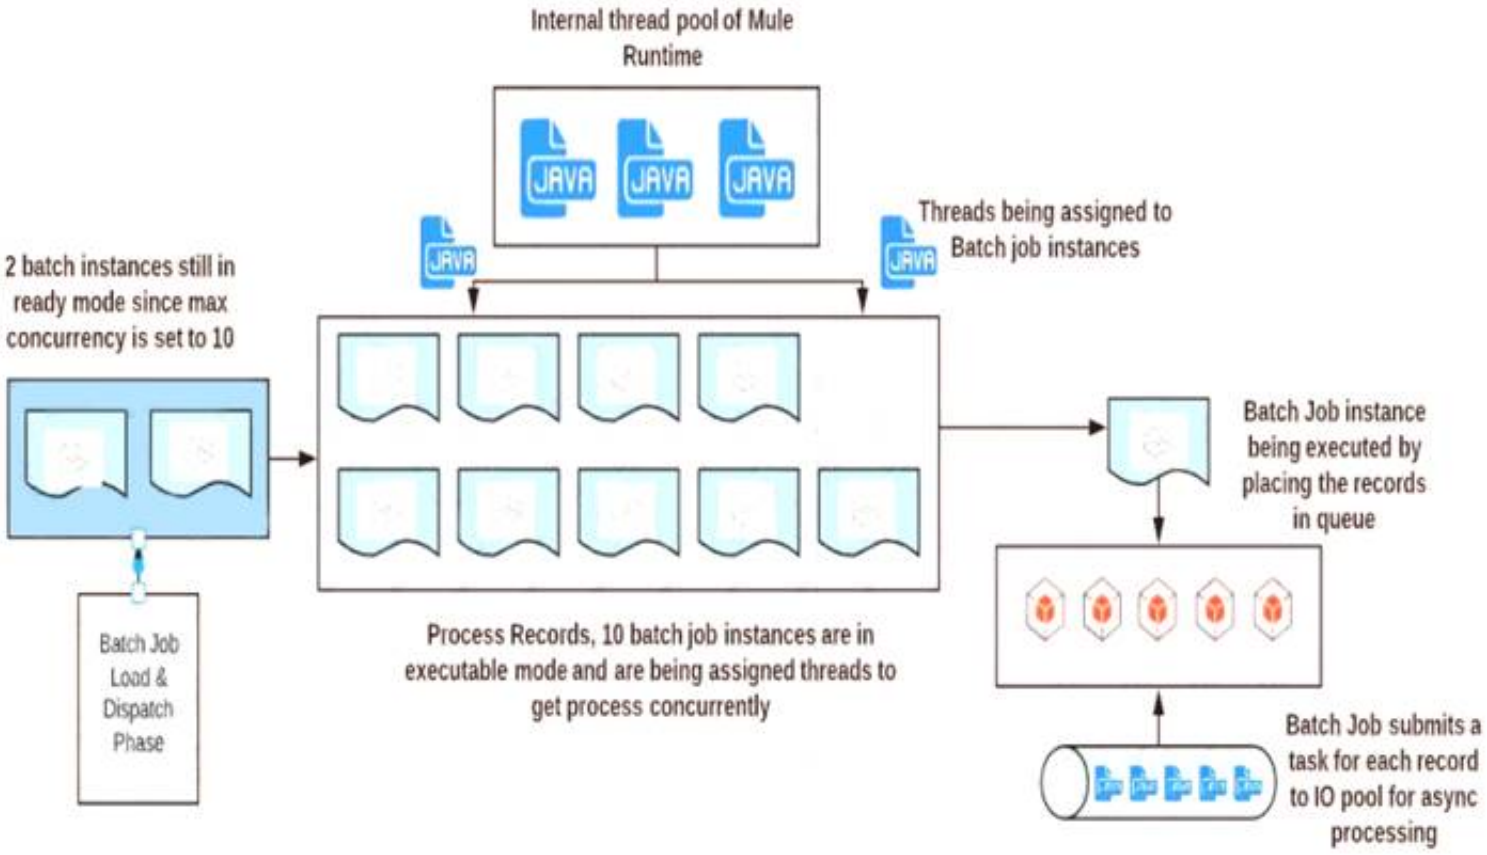

NEW QUESTION 107
A manufacturing company is planning to deploy Mule applications to its own Azure Kubernetes Service infrastructure.
The organization wants to make the Mule applications more available and robust by deploying each Mule application to an isolated Mule runtime in a Docker container while 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 108
A project uses Jenkins to implement CI/CD process. It was observed that each Mule package contains some of the Jenkins files and folders for configurations of CI/CD jobs.
As these files and folders are not part of the actual package, expectation is that these should not be part of deployed archive.
Which file can be used to exclude these files and folders from the deployed archive?

A. muleignore
B. _unTrackMule
C. mulelnclude
D. _muleExclude

Answer: D

NEW QUESTION 111
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

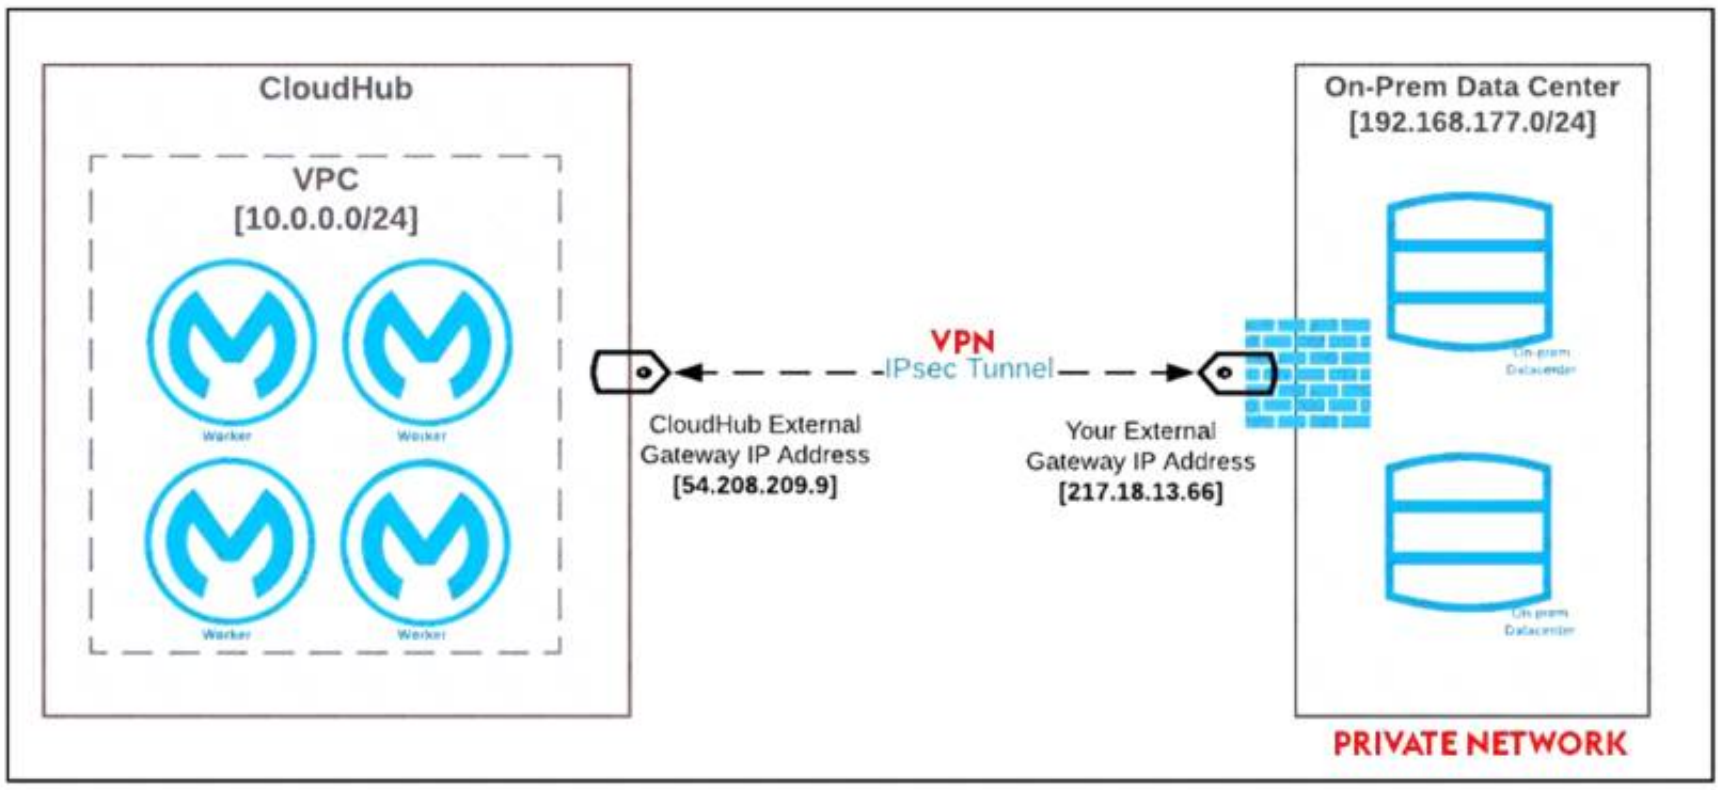

NEW QUESTION 116
Refer to the exhibit.

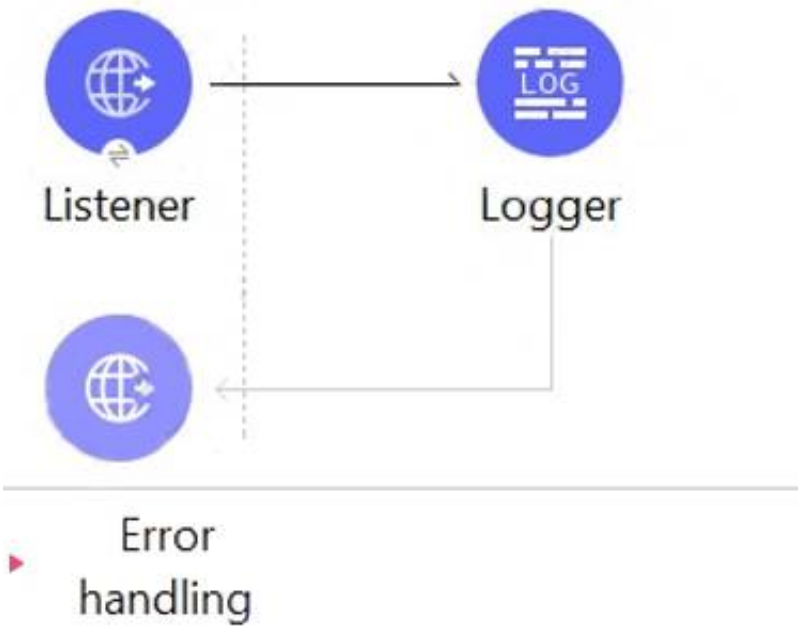

The HTTP Listener and the Logger are being handled from which thread pools respectively?

A. CPU_INTENSIVE and Dedicated Selector pool
B. UBER and NONBLOCKING
C. Shared Selector Pool and CPU LITE
D. BLOCKING _IO and UBER

Answer: C

NEW QUESTION 117
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 118
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 119
An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/'CD pipeline. Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?

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

Answer: D

NEW QUESTION 120
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 125
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 127
An organization has just developed a Mule application that implements a REST API. The mule application will be deployed to a cluster of customer hosted Mule runtimes.
What additional infrastructure component must the customer provide in order to distribute inbound API requests across the Mule runtimes of the cluster?

A. A message broker
B. An HTTP Load Balancer
C. A database
D. An Object Store

Answer: B

Explanation:
Correct answer is An HTTP Load Balancer.
Key thing to note here is that we are deploying application to customer hosted Mule runtime. This means we will need load balancer to route the requests to different instances of the cluster.

NEW QUESTION 130
A marketing organization is designing a Mule application to process campaign data. The Mule application will periodically check for a file in a SFTP location and process the records in the file. The size of the file can vary from 10MB to 5GB. Due to the limited availabiltty of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.
The application must transform and send different formats of this file to three different downstream SFTP locations.
What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

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

Answer: A

NEW QUESTION 132
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

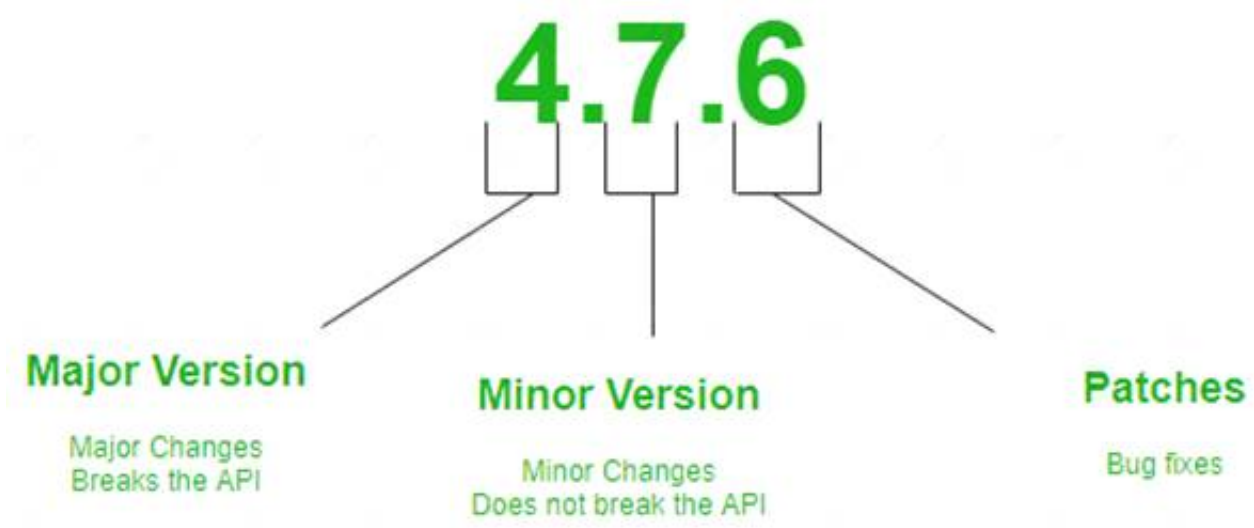

NEW QUESTION 136
A company is designing a mule application to consume batch data from a partner's ftps server The data files have been compressed and then digitally signed using PGP.
What inputs are required for the application to securely consumed these files?

A. ATLS context Key Store requiring the private key and certificate for the company PGP public key of partnerPGP private key for the company
B. ATLS context first store containing a public certificate for partner ftps server and the PGP public key of the partnerTLS contact Key Store containing the FTP credentials
C. TLS context trust or containing a public certificate for the ftps server The FTP username and password The PGP public key of the partner
D. The PGP public key of the partner The PGP private key for the company The FTP username and password

Answer: D

NEW QUESTION 139
The implementation of a Process API must change. What is a valid approach that minimizes the impact of this change on API clients?

A. Implement required changes to the Process API implementation so that whenever possible, the Process API's RAML definition remains unchanged
B. Update the RAML definition of the current Process API and notify API client developers by sending them links to the updated RAML definition
C. Postpone changes until API consumers acknowledge they are ready to migrate to a new Process API or API version
D. Implement the Process API changes in a new API implementation, and have the old API implementation return an HTTP status code 301 - Moved Permanently to inform API clients they should be calling the new API implementation

Answer: A

Explanation:
* Option B shouldn't be used unless extremely needed, if RAML is changed, client needs to accommodate changes. Question is about minimizing impact on Client. So this is not a valid choice.
* Option C isn't valid as Business can't stop for consumers acknowledgment.
* Option D again needs Client to accommodate changes and isn't viable option.
* Best choice is A where RAML definition isn't changed and underlined functionality is changed without any dependency on client and without impacting client.

NEW QUESTION 141
An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL https://api.acme.com
As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub
What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

A. Deploy the Mule applications to CloudhubUpdate the CNAME record for an api.acme.com in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map URLs to their corresponding Mule applications
B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB)Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications
C. Deploy the Mule applications to CloudhubCreate CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancerApply mapping rules in the SLB to map URLs to their corresponding Mule applications
D. Deploy the Mule applications to CloudhubUpdate the CNAME record for api.acme.com in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB)Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

Answer: A

Explanation:
https://help.mulesoft.com/s/feed/0D52T000055pzgsSAA.

NEW QUESTION 144
What Is a recommended practice when designing an integration Mule 4 application that reads a large XML payload as a stream?

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

Answer: C

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

NEW QUESTION 149
As a part of project , existing java implementation is being migrated to Mulesoft. Business is very tight on the budget and wish to complete the project in most economical way possible.
Canonical object model using java is already a part of existing implementation. Same object model is required by mule application for a business use case. What is the best way to achieve this?

A. Make use of Java module
B. Create similar model for Mule applications
C. Create a custom application to read Java code and make it available for Mule application
D. Use Anypoint exchange

Answer: A

Explanation:
Mule 4 is built to:
•Minimize the need for custom code.
•Avoid the need for you to know or understand Java.
However, some advanced uses cases require integration with custom Java code, such as:
•Reuse of a library, such as a tax calculation library.
•Reuse of a canonical object model that is standard in the organization.
•Execution of custom logic using Java.
Mule ref doc : https://docs.mulesoft.com/java-module/1.2/

NEW QUESTION 152
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 154
What is not true about Mule Domain Project?

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

Answer: C

Explanation:
* Mule Domain Project is ONLY available for customer-hosted Mule runtimes, but not for Anypoint Runtime Fabric
* Mule domain project is available for Hybrid and Private Cloud (PCE). Rest all provide application isolation and can't support domain project.
What is Mule Domain Project?
* A Mule Domain Project is implemented to configure the resources that are shared among different projects. These resources can be used by all the projects associated with this domain. Mule applications can be associated with only one domain, but a domain can be associated with multiple projects. Shared resources allow multiple development teams to work in parallel using the same set of reusable connectors. Defining these connectors as shared resources at the domain level allows the team to: - Expose multiple services within the domain through the same port. - Share the connection to persistent storage. - Share services between apps through a well-defined interface. - Ensure consistency between apps upon any changes because the configuration is only set in one place.
* Use domains Project to share the same host and port among multiple projects. You can declare the http connector within a domain project and associate the domain project with other projects. Doing this also allows to control thread settings, keystore configurations, time outs for all the requests made within multiple applications. You may think that one can also achieve this by duplicating the http connector configuration across all the applications. But, doing this may pose a nightmare if you have to make a change and redeploy all the applications.
* If you use connector configuration in the domain and let all the applications use the new domain instead of a default domain, you will maintain only one copy of the http connector configuration. Any changes will require only the domain to the redeployed instead of all the applications.
You can start using domains in only three steps:
1) Create a Mule Domain project
2) Create the global connector configurations which needs to be shared across the applications inside the Mule Domain project
3) Modify the value of domain in mule-deploy.properties file of the applications Graphical user interface Description automatically generated

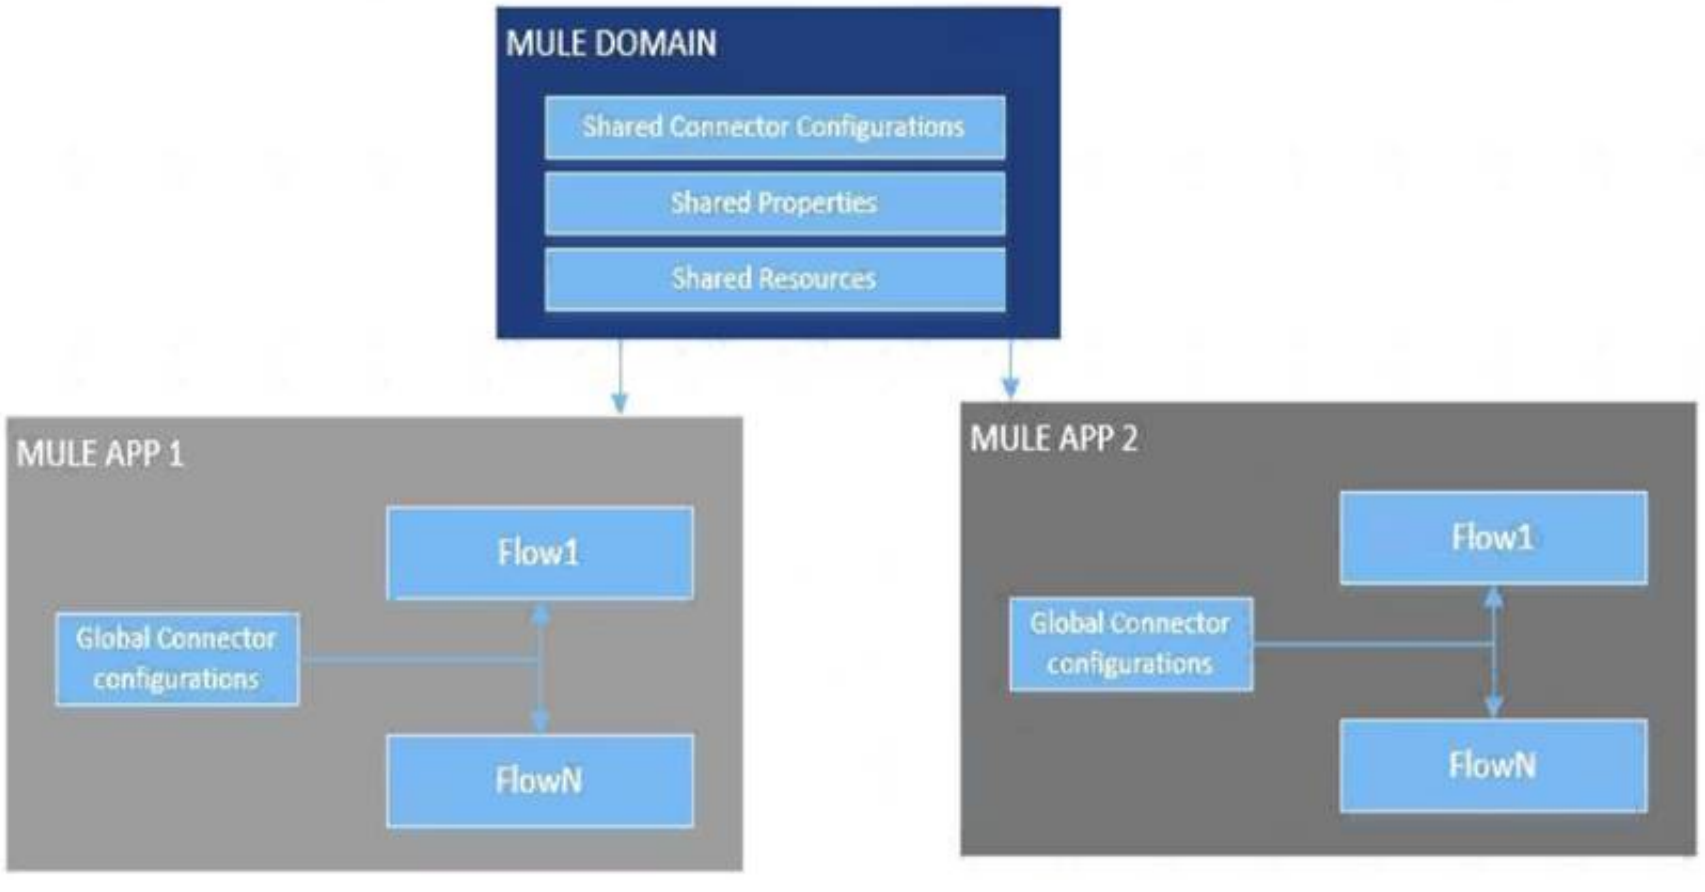

NEW QUESTION 155
A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is 58. proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application.
What non-functional 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 156
A global, high-volume shopping Mule application is being built and will be deployed to CloudHub. To improve performance, the Mule application uses a Cache scope that maintains cache state in a CloudHub object store. Web clients will access the Mule application over HTTP from all around the world, with peak volume coinciding with business hours in the web client's geographic location. To achieve optimal performance, what Anypoint Platform region should be chosen for the CloudHub object store?

A. Choose the same region as to where the Mule application is deployed
B. Choose the US-West region, the only supported region for CloudHub object stores
C. Choose the geographically closest available region for each web client
D. Choose a region that is the traffic-weighted geographic center of all web clients

Answer: A

Explanation:
CloudHub object store should be in same region where the Mule application is deployed. This will give optimal performance.
Before learning about Cache scope and object store in Mule 4 we understand what is in general Caching is and other related things.
WHAT DOES “CACHING” MEAN?
Caching is the process of storing frequently used data in memory, file system or database which saves processing time and load if it would have to be accessed from original source location every time.
In computing, a cache is a high-speed data storage layer which stores a subset of data, so that future requests for that data are served up faster than is possible by accessing the data’s primary storage location. Caching allows you to efficiently reuse previously retrieved or computed data.
How does Caching work?
The data in a cache is generally stored in fast access hardware such as RAM (Random-access memory) and may also be used in correlation with a software component. A cache’s primary purpose is to increase data retrieval performance by reducing the need to access the underlying slower storage layer.
Caching in MULE 4
In Mule 4 caching can be achieved in mule using cache scope and/or object-store. Cache scope internally uses Object Store to store the data.
What is Object Store
Object Store lets applications store data and states across batch processes, Mule components, and applications, from within an application. If used on cloud hub,

the object store is shared between applications deployed on Cluster.
Cache Scope is used in below-mentioned cases:
Need to store the whole response from the outbound processor
Data returned from the outbound processor does not change very frequently
As Cache scope internally handle the cache hit and cache miss scenarios it is more readable Object Store is used in below-mentioned cases:
Need to store custom/intermediary data To store watermarks
Sharing the data/stage across applications, schedulers, batch.
If CloudHub object store is in same region where the Mule application is deployed it will aid in fast access of data and give optimal performance.

NEW QUESTION 158
Refer to the exhibit.

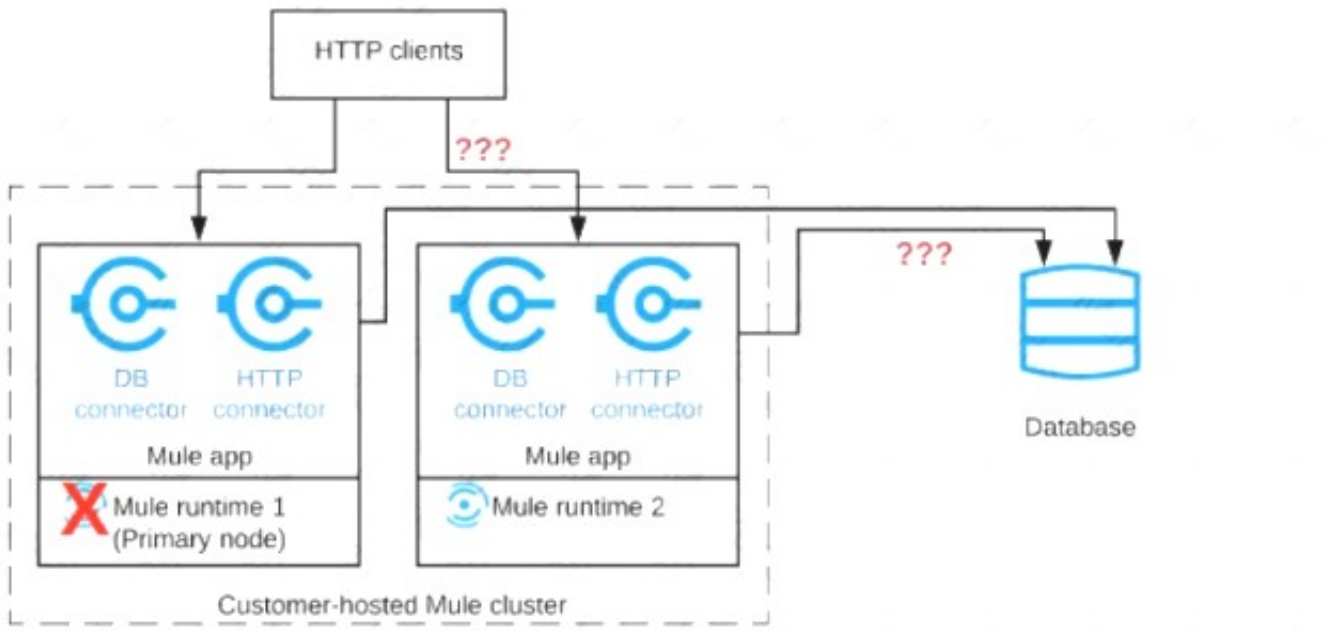

A Mule application is deployed to a cluster of two customer-hosted Mute runtimes. The Mute application has a flow that polls a database and another flow with an HTTP Listener. HTTP clients send HTTP requests directly to individual cluster nodes.
What happens to database polling and HTTP request handling 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 descripted 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

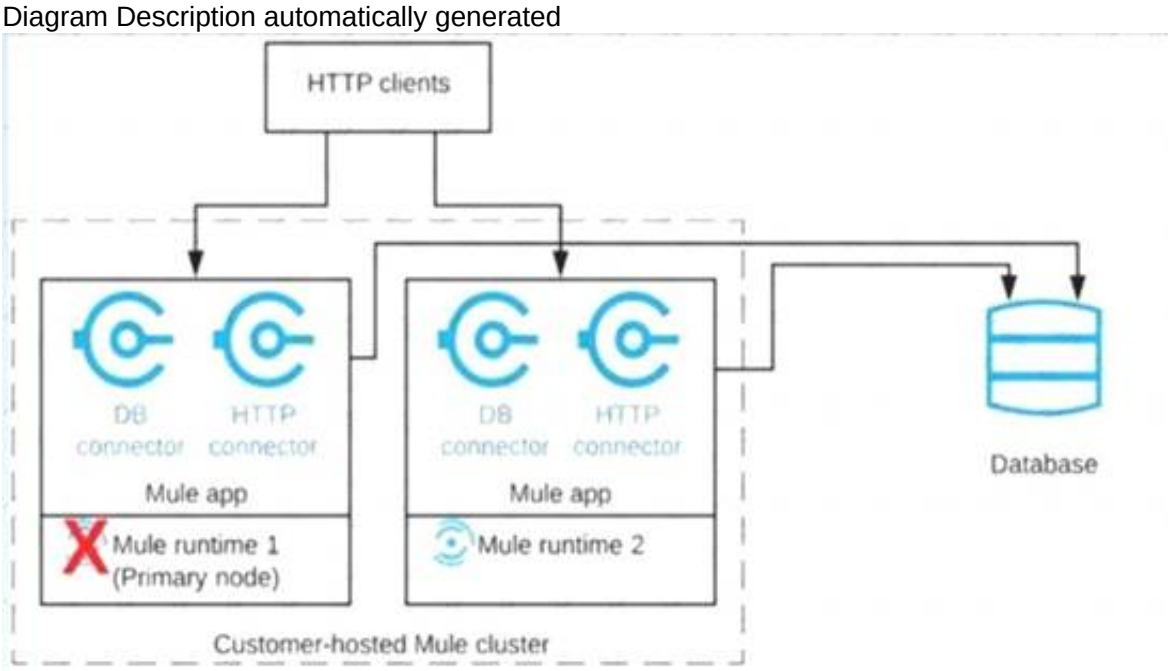

NEW QUESTION 162
A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.
The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.
What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

A. Retrieve the OpenAPI Specification file(s) from API Manager

B. Retrieve the WSDL file(s) from the deployed Mule application
C. Retrieve the RAML file(s) from the deployed Mule application
D. Retrieve the XML file(s) from Runtime Manager

Answer: D

NEW QUESTION 167
A banking company is developing a new set of APIs for its online business. One of the critical API's is a master lookup API which is a system API. This master lookup API uses persistent object store. This API will be used by all other APIs to provide master lookup data.

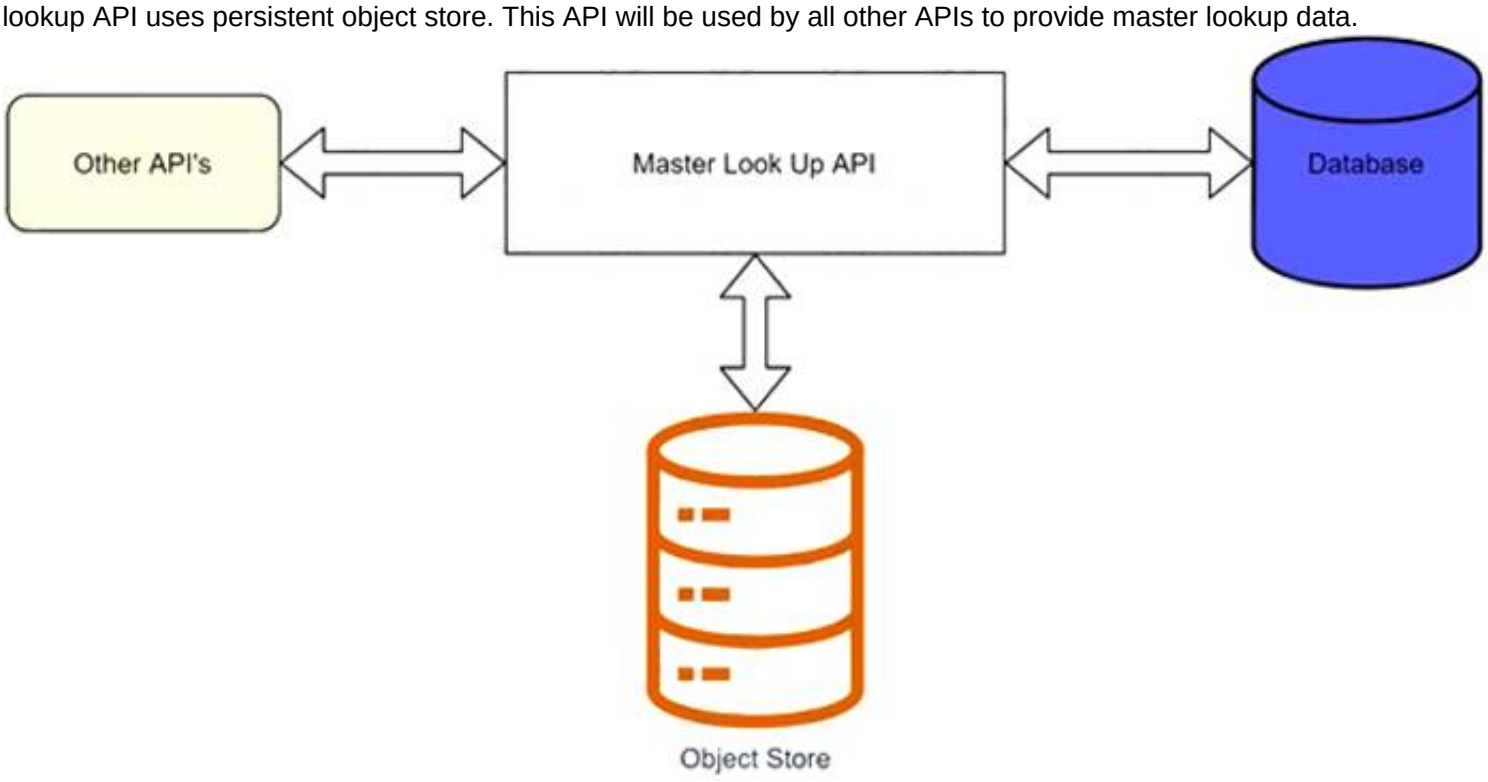

Master lookup API is deployed on two cloudhub workers of 0.1 vCore each because there is a lot of master data to be cached. Master lookup data is stored as a key value pair. The cache gets refreshed if they key is not found in the cache.
Doing performance testing it was observed that the Master lookup API has a higher response time due to database queries execution to fetch the master lookup data.
Due to this performance issue, go-live of the online business is on hold which could cause potential financial loss to Bank.
As an integration architect, which of the below option you would suggest to resolve performance issue?

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

Answer: A

NEW QUESTION 168
An integration Mule application is deployed to a customer-hosted multi-node Mule 4 runtime duster. The Mule application uses a Listener operation of a JMS connector to receive incoming messages from a JMS
queue.
How are the messages consumed by the Mule application?

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

Answer: C

Explanation:
Correct answer is Depending on the Listener operation configuration, either all messages are consumed by ONLY the primary cluster node or else EACH message is consumed by ANY ONE cluster node
For applications running in clusters, you have to keep in mind the concept of primary node and how the connector will behave. When running in a cluster, the JMS listener default behavior will be to receive messages only in the primary node, no matter what kind of destination you are consuming from. In case of consuming messages from a Queue, you'll want to change this configuration to receive messages in all the nodes of the cluster, not just the primary.
This can be done with the primaryNodeOnly parameter:
<jms:listener config-ref="config" destination="${inputQueue}" primaryNodeOnly="false"/>

NEW QUESTION 172
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 175
What Mule application can have API policies applied by Anypoint Platform to the endpoint exposed by that Mule application?

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

Answer: A

Explanation:
* HTTP/1.1 keeps all requests and responses in plain text format.
* HTTP/2 uses the binary framing layer to encapsulate all messages in binary format, while still maintaining HTTP semantics, such as verbs, methods, and headers. It came into use in 2015, and offers several methods to decrease latency, especially when dealing with mobile platforms and server-intensive graphics and videos
* Currently, Mule application can have API policies only for Mule application that accepts requests over HTTP/1x

NEW QUESTION 180
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your MCIA-Level-1 Exam with Our Prep Materials Via below:

https://www.certleader.com/MCIA-Level-1-dumps.html