

## Exam Questions MCI A-Level-1

MuleSoft Certified Integration Architect - Level 1

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



#### NEW QUESTION 1

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

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 2

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

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

**Answer:** B

#### Explanation:

- \* Clustering is a group of servers or mule runtime which acts as a single unit.
  - \* Mulesoft Enterprise Edition supports scalable clustering to provide high availability for the Mulesoft application.
  - \* In simple terms, virtual servers composed of multiple nodes and they communicate and share information through a distributed shared memory grid.
  - \* By default, Mulesoft ensures the High availability of applications if clustering implemented.
  - \* Let's consider the scenario one of the nodes in cluster crashed or goes down and under maintenance. In such cases, Mulesoft will ensure that requests are processed by other nodes in the cluster. Mulesoft clustering also ensures that the request is load balanced between all the nodes in a cluster.
  - \* Clustering is only supported by on-premise Mule runtime and it is not supported in Cloudhub.
- Correct answer is External monitoring tools or log aggregators must be configured to recognize the new nodes
- \* Rest of the options are automatically taken care of when a new node is added in cluster.

#### NEW QUESTION 3

An organization has implemented a continuous integration (CI) lifecycle that promotes Mule applications through code, build, and test stages. To standardize the organization's CI journey, a new dependency control approach is being designed to store artifacts that include information such as dependencies, versioning, and build promotions.

To implement these process improvements, the organization will now require developers to maintain all dependencies related to Mule application code in a shared location.

What is the most idiomatic (used for its intended purpose) type of system the organization should use in a shared location to standardize all dependencies related to Mule application code?

- A. A MuleSoft-managed repository at repository.mulesoft.org
- B. A binary artifact repository
- C. API Community Manager
- D. The Anypoint Object Store service at cloudhub.io

**Answer:** C

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

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

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

**Answer: B**

#### Explanation:

As the stakeholders are semitechnical users, preferred option is Create Anypoint Exchange entries with pages elaborating the integration design, including API notebooks (where applicable) to help the stakeholders understand and interact with the Mule applications and APIs at various levels of technical depth

#### NEW QUESTION 6

A REST API is being designed to implement a Mule application.

What standard interface definition language can be used to define REST APIs?

- A. Web Service Definition Language(WSDL)
- B. OpenAPI Specification (OAS)
- C. YAML
- D. AsyncAPI Specification

**Answer: B**

#### NEW QUESTION 7

An organization has implemented the cluster with two customer hosted Mule runtimes is hosting an application.

This application has a flow with a JMS listener configured to consume messages from a queue destination. As an integration architect can you advise which JMS listener configuration must be used to receive messages in all the nodes of the cluster?

- A. Use the parameter primaryNodeOnly= "false" on the JMS listener
- B. Use the parameter primaryNodeOnly= "false" on the JMS listener with a shared subscription
- C. Use the parameter primaryNodeOnly= "true" on the JMS listener with a non-shared subscription
- D. Use the parameter primaryNodeOnly= "true" on the JMS listener

**Answer: A**

#### NEW QUESTION 8

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

The organization does not currently use AWS in any way.

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

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

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

**Answer: D**

#### Explanation:

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

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

#### NEW QUESTION 9

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

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 10**

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

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

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

**Answer: B**

**NEW QUESTION 10**

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

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

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

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

**Answer: A**

**Explanation:**

When an operation in a Mule application fails to connect to an external server, the default behavior is for the operation to fail immediately and return a connectivity error. You can modify this default behavior by configuring a reconnection strategy for the operation. You can configure a reconnection strategy for an operation either by modifying the operation properties or by modifying the configuration of the global element for the operation. The following are the available reconnection strategies and their behaviors: None Is the default behavior, which immediately returns a connectivity error if the attempt to connect is unsuccessful Standard (reconnect) Sets the number of reconnection attempts and the interval at which to execute them before returning a connectivity error Forever (reconnect-forever) Attempts to reconnect continually at a given interval

**NEW QUESTION 12**

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

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

**Answer: C**

**Explanation:**

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

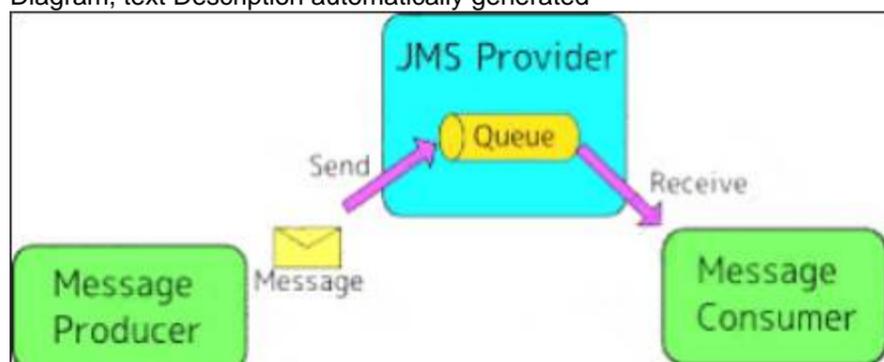
\* JMS Connector enables sending and receiving messages to queues and topics for any message service that implements the JMS specification.

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

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

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

Diagram, text Description automatically generated



**NEW QUESTION 17**

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

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

**Answer: C**

**Explanation:**

Correct answer is Use a Batch job scope to bulk insert records into the database

\* Batch Job is most efficient way to manage millions of records. A few points to note here are as follows :

Reliability: If you want reliabilty while processing the records, i.e should the processing survive a runtime crash or other unhappy scenarios, and when restarted process all the remaining records, if yes then go for batch as it uses persistent queues.

Error Handling: In Parallel for each an error in a particular route will stop processing the remaining records in that route and in such case you'd need to handle it using on error continue, batch process does not stop during such error instead you can have a step for failures and have a dedicated handling in it.

Memory footprint: Since question said that there are millions of records to process, parallel for each will aggregate all the processed records at the end and can possibly cause Out Of Memory.

Batch job instead provides a BatchResult in the on complete phase where you can get the count of failures and success. For huge file processing if order is not a concern definitely go ahead with Batch Job

**NEW QUESTION 20**

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

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

**Answer: B**

**Explanation:**

Types of logging:

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

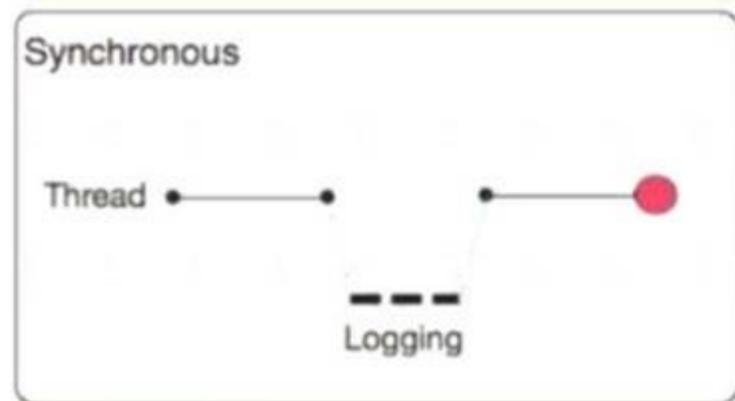
The execution of the thread that is processing your message is interrupted to wait for the log message to be fully output before it can continue

Performance degrades because of synchronous logging

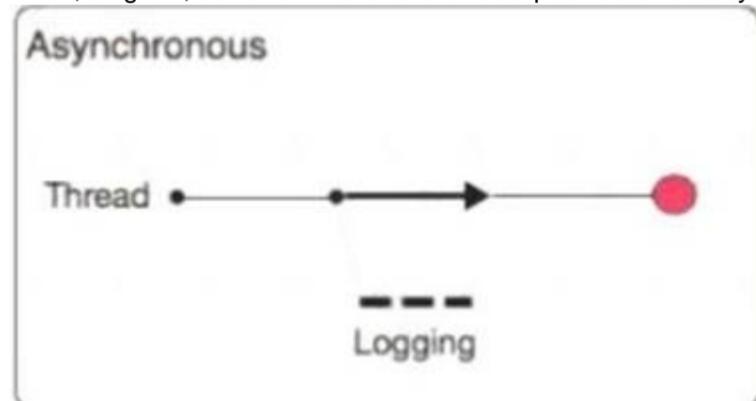
Used when the log is used as an audit trail or when logging ERROR/CRITICAL messages

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

Chart, diagram Description automatically generated



Chart, diagram, box and whisker chart Description automatically generated



B) Asynchronous:

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

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

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

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

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

**NEW QUESTION 22**

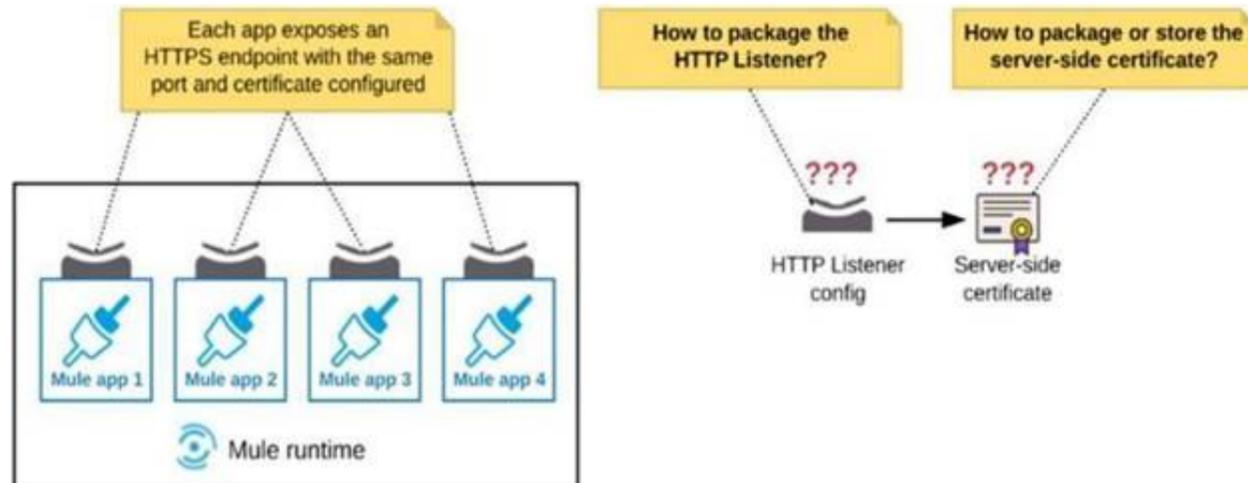
An organization is designing a Mule application to periodically poll an SFTP location for new files containing sales order records and then process those sales orders. Each sales order must be processed exactly once. To support this requirement, the Mule application must identify and filter duplicate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

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

**Answer: C**

**NEW QUESTION 27**

Refer to the exhibit.



An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a server-side certificate that rotates often. What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

- A. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint Package the server-side certificate in ALL Mule APPLICATIONS that need to expose an HTTPS endpoint
- B. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint
- C. Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION
- D. Package an HTTPS Listener configuration In all Mule APPLICATIONS that need to expose an HTTPS endpoint Package the server-side certificate in a NEW Mule DOMAIN project
- E. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing It from all Mule applications that need to expose an HTTPS endpoint
- F. Package the server-side certificate in the SAME Mule DOMAIN project Go to Set

**Answer: B**

**Explanation:**

In this scenario, both A & C will work, but A is better as it does not require repackage to the domain project at all. Correct answer is Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint. Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION.

What is Mule Domain Project?

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

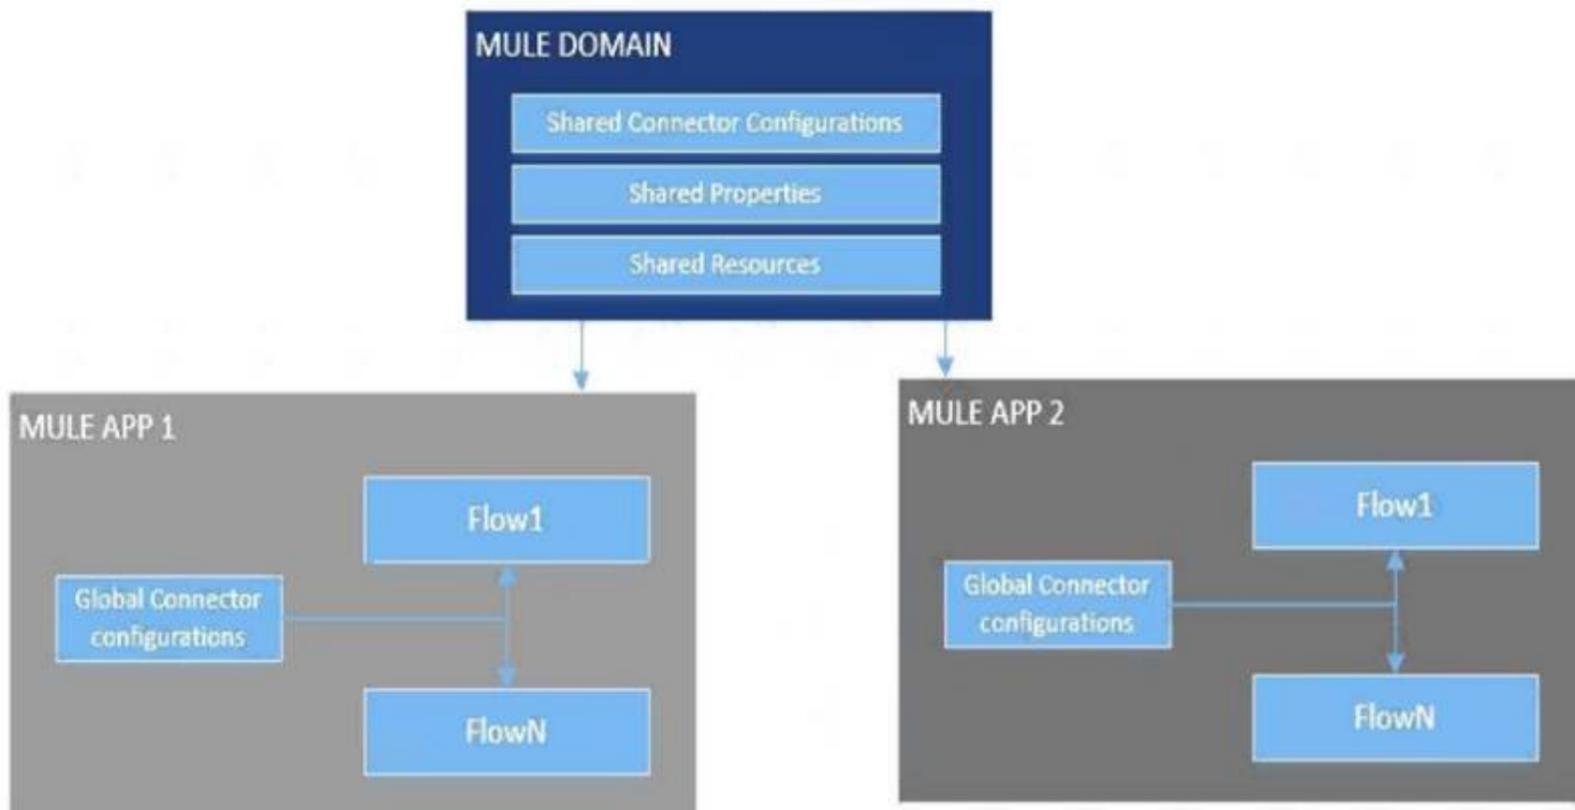
- Expose multiple services within the domain through the same port.
- Share the connection to persistent storage.
- Share services between apps through a well-defined interface.
- Ensure consistency between apps upon any changes because the configuration is only set in one place.

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

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

You can start using domains in only three steps:

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



Use a certificate defined in already deployed Mule domain Configure the certificate in the domain so that the API proxy HTTPS Listener references it, and then deploy the secure API proxy to the target Runtime Fabric, or on-premises target. (CloudHub is not supported with this approach because it does not support Mule domains.)

**NEW QUESTION 29**

An Organization has previously provisioned its own AWS VPC hosting various servers. The organization now needs to use Cloudhub to host a Mule application that will implement a REST API once deployed to Cloudhub, this Mule application must be able to communicate securely with the customer-provisioned AWS VPC resources within the same region, without being interceptable on the public internet.

What Anypoint Platform features should be used to meet these network communication requirements between Cloudhub and the existing customer-provisioned AWS VPC?

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

**Answer: A**

**Explanation:**

Correct answer is: Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC

\* Connecting to your Anypoint VPC extends your corporate network and allows CloudHub workers to access resources behind your corporate firewall.

\* You can connect on-premises data centers through a secured VPN tunnel, or a private AWS VPC through VPC peering, or by using AWS Direct Connect.

MuleSoft Doc Reference : <https://docs.mulesoft.com/runtime-manager/virtual-private-cloud>

**NEW QUESTION 30**

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

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

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

**Answer: A**

**NEW QUESTION 34**

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

- A. Compile, package, unit test, deploy, create associated API instances in API Manager
- B. Import from API designer, compile, package, unit test, deploy, publish to Am/point Exchange
- C. Compile, package, unit test, validate unit test coverage, deploy
- D. Compile, package, unit test, deploy, integration test

**Answer: C**

**NEW QUESTION 39**

An organization is using Mulesoft cloudhub and develops API's in the latest version. As a part of requirements for one of the API's, third party API needs to be called. The security team has made it clear that calling any external API needs to have include listing

As an integration architect please suggest the best way to accomplish the design plan to support these requirements?

- A. Implement include list IP on the cloudhub VPC firewall to allow the traffic
- B. Implement the validation of included IP operation
- C. Implement the Any point filter processor to implement the include list IP
- D. Implement a proxy for the third party API and enforce the IP include list policy and call this proxy from the flow of the API

**Answer:** D

#### NEW QUESTION 41

What is true about automating interactions with Anypoint Platform using tools such as Anypoint Platform REST API's, Anypoint CLI or the Mule Maven plugin?

- A. By default, the Anypoint CLI and Mule Maven plugin are not included in the Mule runtime
- B. Access to Anypoint Platform API's and Anypoint CLI can be controlled separately through the roles and permissions in Anypoint platform, so that specific users can get access to Anypoint CLI while others get access to the platform API's
- C. Anypoint Platform API's can only automate interactions with CloudHub while the Mule maven plugin is required for deployment to customer hosted Mule runtimes
- D. API policies can be applied to the Anypoint platform API's so that only certain LOS's has access to specific functions

**Answer:** A

#### Explanation:

Correct answer is By default, the Anypoint CLI and Mule Maven plugin are not included in the Mule runtime Maven is not part of runtime though it is part of studio. You do not need it to deploy in order to deploy your app. Same is the case with CLI.

#### NEW QUESTION 42

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 44

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 46

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 51

A mule application designed to fulfil two requirements

a) Processing files are synchronously from an FTPS server to a back-end database using VM intermediary queues for load balancing VM events  
 b) Processing a medium rate of records from a source to a target system using batch job scope  
 Considering the processing reliability requirements for FTPS files, how should VM queues be configured for processing files as well as for the batch job scope if the application is deployed to Cloudhub workers?

- A. Use Cloud hub persistent queues for FTPS files processing There is no need to configure VM queues for the batch jobs scope as it uses by default the worker's disc for VM queueing
- B. Use Cloud hub persistent VM queue for FTPS file processing There is no need to configure VM queues for the batch jobs scope as it uses by default the worker's JVM memory for VM queueing
- C. Use Cloud hub persistent VM queues for FTPS file processing Disable VM queue for the batch job scope
- D. Use VM connector persistent queues for FTPS file processing Disable VM queue for the batch job scope

**Answer: C**

#### NEW QUESTION 53

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 54

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

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

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

**Answer: D**

#### Explanation:

Anypoint Platform provides a scripting and command-line tool for both Anypoint Platform and Anypoint Platform Private Cloud Edition (Anypoint Platform PCE). The command-line interface (CLI) supports both the interactive shell and standard CLI modes and works with: Anypoint Exchange Access management Anypoint Runtime Manager

#### NEW QUESTION 58

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 61

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 retriee
- E. A dead-letter object store configured in the Mule application

**Answer: B**

#### NEW QUESTION 63

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

What is the best way to implement the data transformation logic for this new Mule application while minimizing the overall testing effort?

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 66**

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**

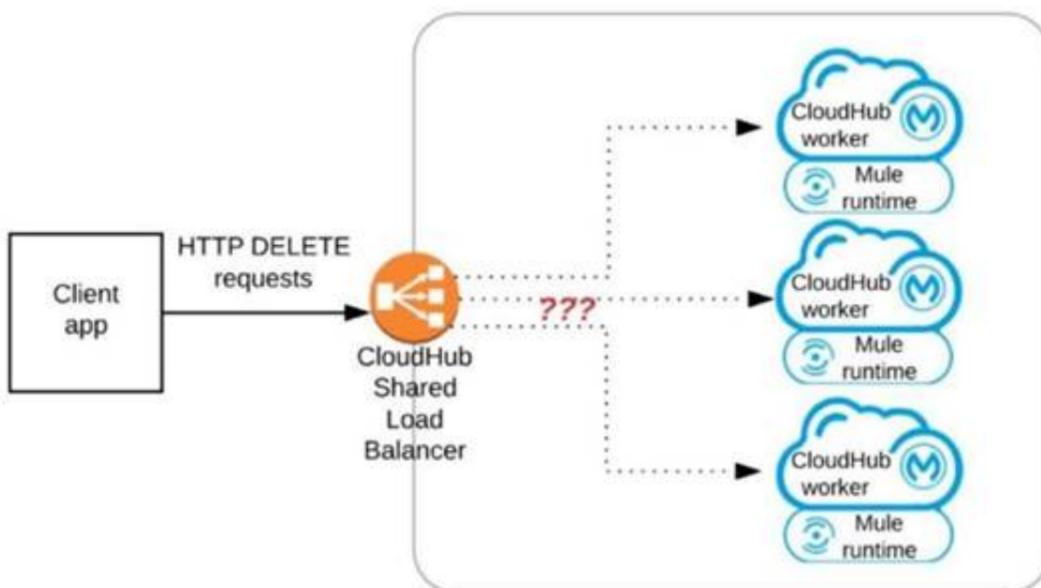
**Explanation:**

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

**NEW QUESTION 69**

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

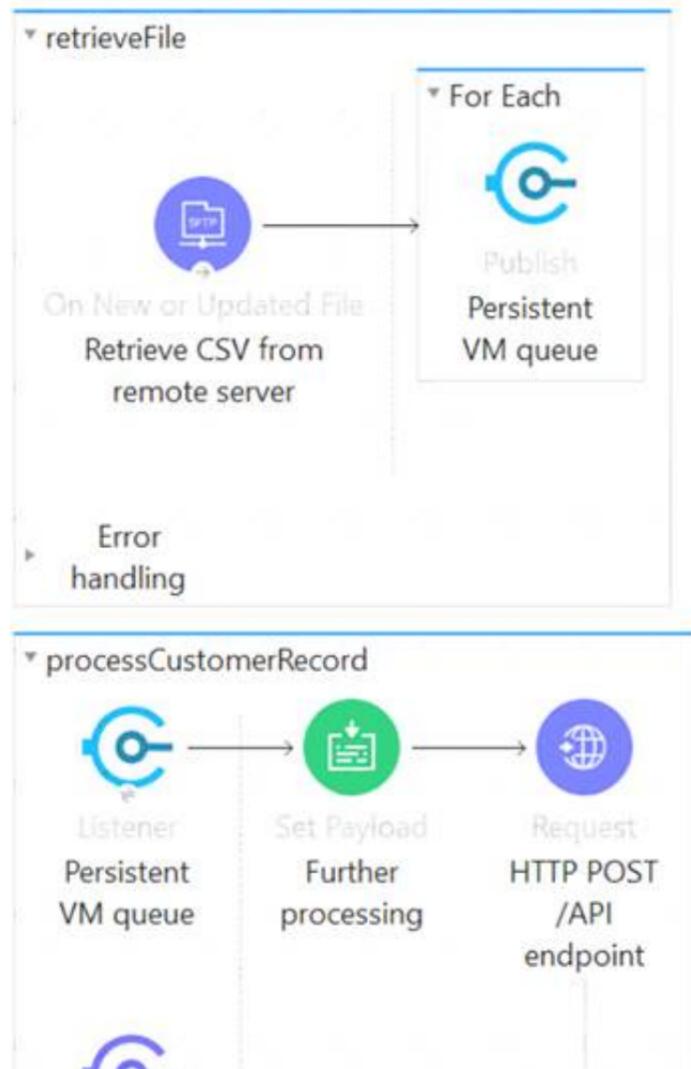
**Answer: C**

**Explanation:**

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

**NEW QUESTION 73**

Refer to the exhibit.



This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow's VM Listener receives messages from the same VM queue and then processes each message separately.

How are messages routed to the cloudhub workers as messages are received by the VM Listener?

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

**Answer: B**

**NEW QUESTION 75**

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 76**

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

(Choose two answers)

- A. In the Java Invoke static connector configuration, configure a path and name of the dependency jar file
- B. Add the dependency jar file to the java classpath by setting the JVM parameters
- C. Use Maven command to include the dependency jar file when packaging the application
- D. Configure the dependency as a shared library in the project POM
- E. Update mule-artefact.json to export the Java package

**Answer: BD**

**NEW QUESTION 79**

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 84**

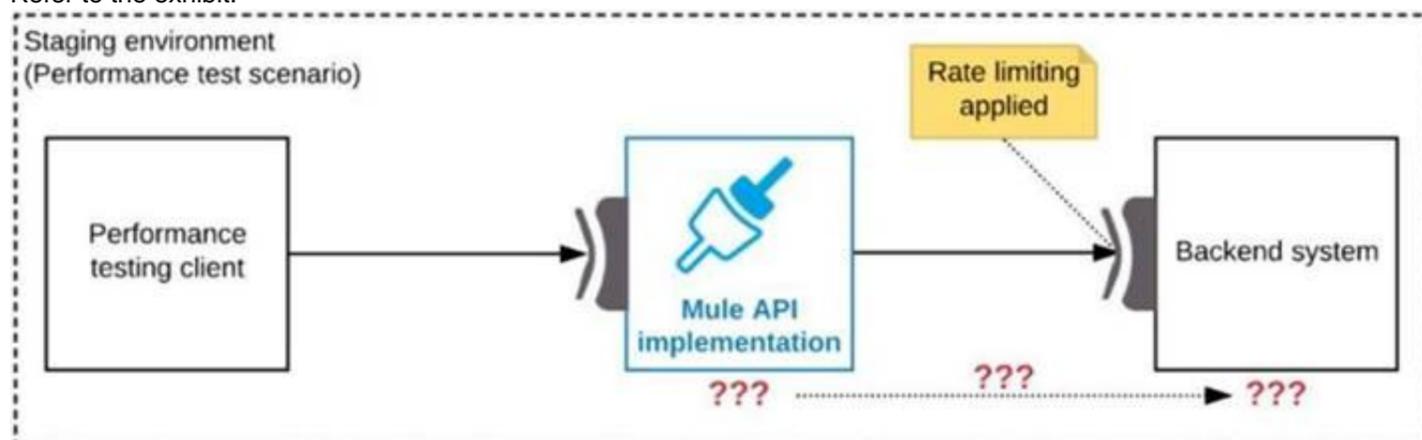
An organization is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted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

- A. An Until Successful scope to call the back-end system One or more ActiveMQ long-retry queues One or more ActiveMQ dead-letter queues for manual processing
- B. One or more On Error scopes to assist calling the back-end system An Until Successful scope containing VM components for long retries A persistent dead-letter VM queue configured in CloudHub
- C. One or more On Error scopes to assist calling the back-end system One or more ActiveMQ long-retry queues A persistent dead-letter object store configured in the CloudHub Object Store service
- D. A Batch Job scope to call the back-end system An Until Successful scope containing Object Store components for long retries A dead-letter object store configured in the Mule application

Answer: A

**NEW QUESTION 87**

Refer to the exhibit.



One of the backend systems invoked by an API implementation enforces rate limits on the number of requests a particular client can make. Both the backend system and the API implementation are deployed to several non-production environments in addition to production. Rate limiting of the backend system applies to all non-production environments. The production environment, however, does NOT have any rate limiting. What is the most effective approach to conduct performance tests of the API implementation in a staging (non-production) environment?

- A. Create a mocking service that replicates the backend system's production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests
- B. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
- C. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situation
- D. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
- E. Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale

Answer: A

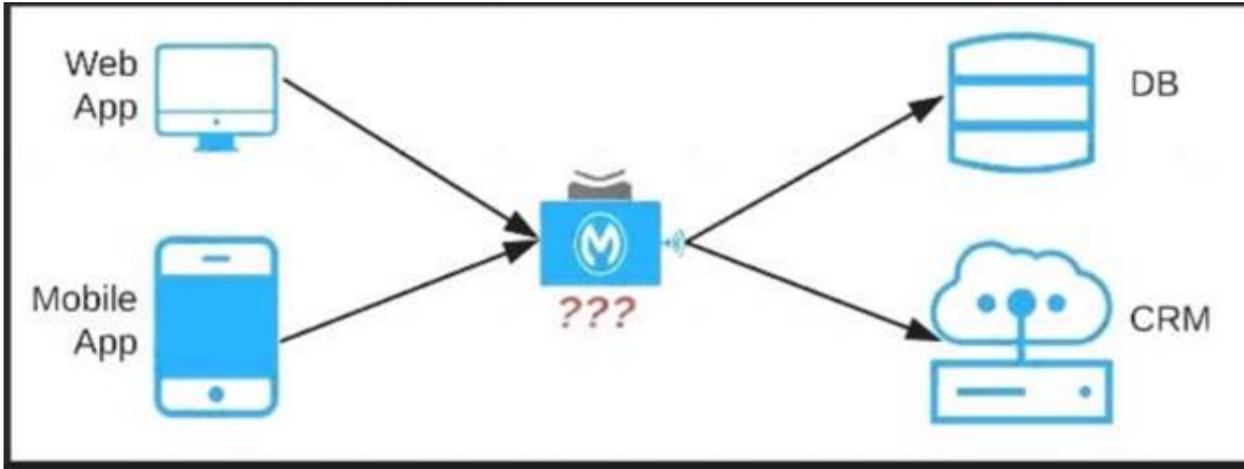
**Explanation:**

Correct answer is Create a mocking service that replicates the backend system’s production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests

- \* MUnit is for only Unit and integration testing for APIs and Mule apps. Not for performance Testing, even if it has the ability to Mock the backend.
- \* Bypassing the backend invocation defeats the whole purpose of performance testing. Hence it is not a valid answer.
- \* Scaled down performance tests cant be relied upon as performance of API's is not linear against load.

**NEW QUESTION 90**

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



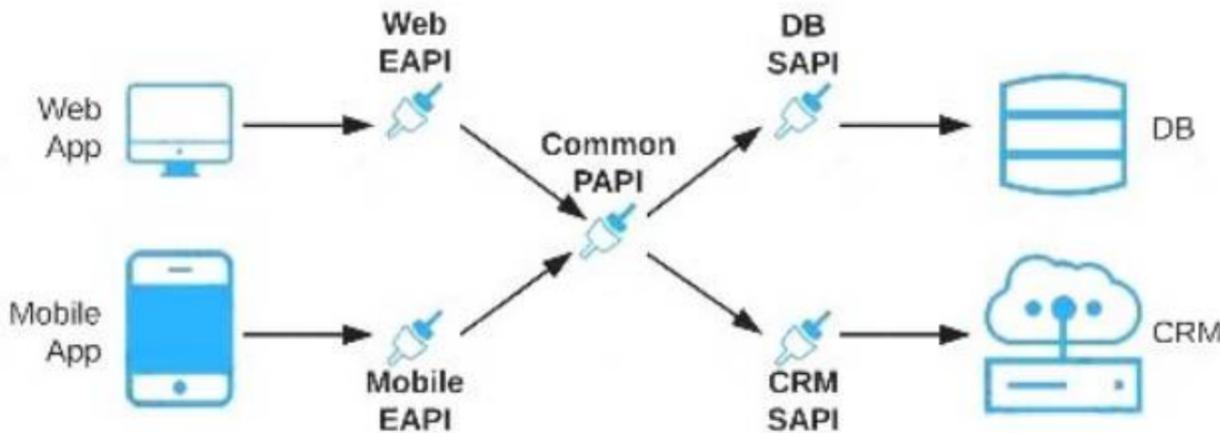
- A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
- B. One set of APIs (Experience API, Process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

**Answer: C**

**Explanation:**

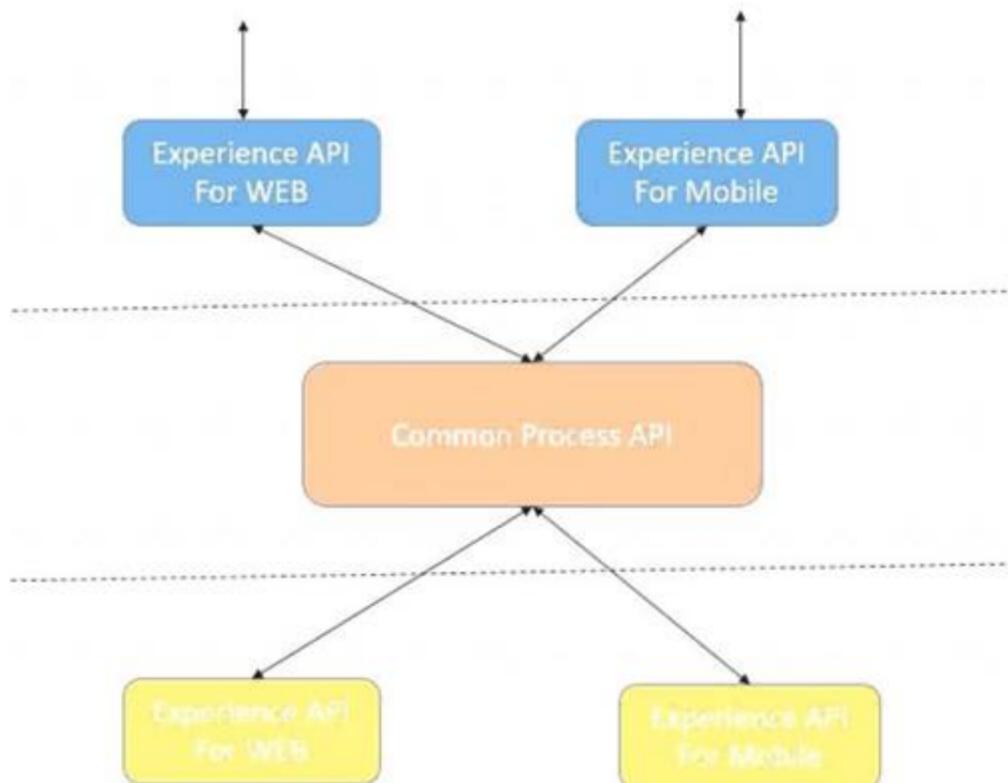
Lets analyze the situation in regards to the different options available  
 Option : A common Experience API but separate Process APIs  
 Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements 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 91**

An organization is creating a Mule application that will be deployed to CloudHub. The Mule application has a property named dbPassword that stores a database user's password.

The organization's security standards indicate that the dbPassword property must be hidden from every Anypoint Platform user after the value is set in the Runtime Manager Properties tab.

What configuration in the Mule application helps hide the dbPassword property value in Runtime Manager?

- A. Use secure::dbPassword as the property placeholder name and store the cleartext (unencrypted) value in a secure properties placeholder file
- B. Use secure::dbPassword as the property placeholder name and store the property encrypted value in a secure properties placeholder file
- C. Add the dbPassword property to the secureProperties section of the pom.xml file
- D. Add the dbPassword property to the secureProperties section of the mule-artifact.json file

**Answer: B**

**NEW QUESTION 92**

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 97**

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.

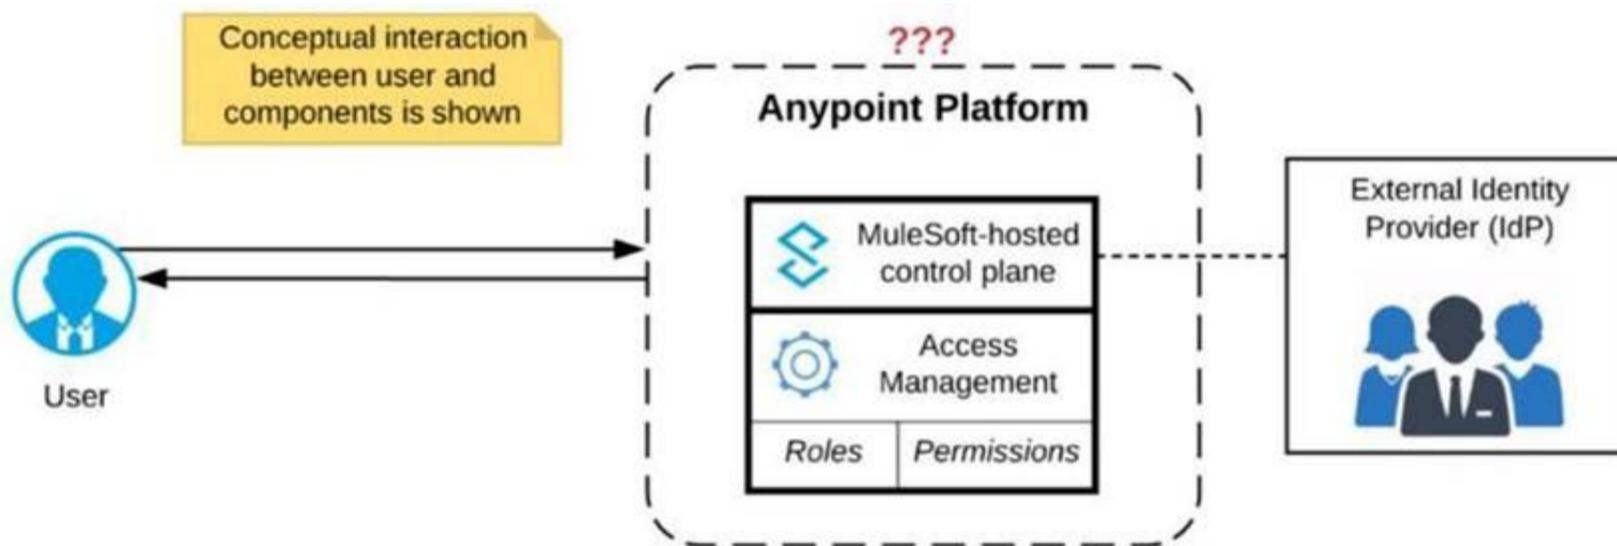
**Answer: D**

**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 98**

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 100**

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.

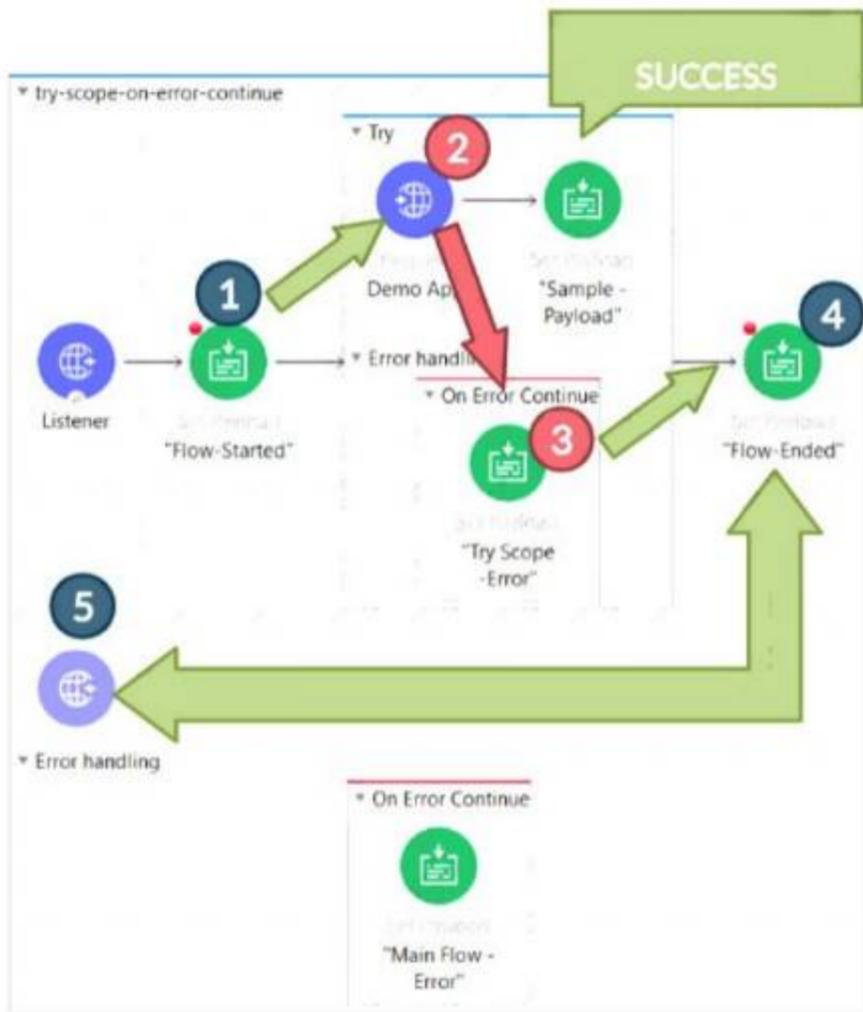
\* Scatter gather is used to club multiple responses together before processing. In this scenario, we need sequential processing. So option A is out of choice.

\* 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 103**

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 107**

A Mule application is synchronizing customer data between two different database systems.

What is the main benefit of using eXtended Architecture (XA) transactions over local transactions to synchronize these two different database systems?

- A. An XA transaction synchronizes the database systems with the least amount of Mule configuration or coding
- B. An XA transaction handles the largest number of requests in the shortest time
- C. An XA transaction automatically rolls back operations against both database systems if any operation fails
- D. An XA transaction writes to both database systems as fast as possible

**Answer: B**

**NEW QUESTION 108**

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

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

**Answer: D**

**NEW QUESTION 113**

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

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

- A. EDJFACT XML over SFTP JSON/REST over HTTPS
- B. SOAP over HTTPS HOP over TLS gRPC over HTTPS

- C. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
- D. CSV over FTP YAML over TLS JSON over HTTPS

**Answer:** C

**Explanation:**

As per definition of API by Mulesoft, it is Application Programming Interface using HTTP-based protocols. Non-HTTP-based programmatic interfaces are not APIs. \* HTTP-based programmatic interfaces are APIs even if they don't use REST or JSON. Hence implementation based on Java RMI, CORBA/IIOP, raw TCP/IP interfaces are not API's as they are not using HTTP.

\* One more thing to note is FTP was not built to be secure. It is generally considered to be an insecure protocol because it relies on clear-text usernames and passwords for authentication and does not use encryption.

\* Data sent via FTP is vulnerable to sniffing, spoofing, and brute force attacks, among other basic attack methods.

Considering the above points only correct option is

- XML over ActiveMQ
- XML over SFTP
- XML/REST over HTTPS

**NEW QUESTION 115**

A rate limiting policy has been applied to a soap V1.2 API published in Cloudfoundry. The API implementation catches errors in a global error handler on error propagate in the main flow for HTTP: RETRY\_EXHAUSTED with HTTP status set to 429 and any with the HTTP status set to 500.

What is the expected HTTP status when the client exceeds the quota of the API calls?

- A. HTTP status 429 as defined in the HTTP:RETRY\_EXHAUSTED error handler in the API
- B. HTTP status 500 as defined in the ANY error handler in the API since an API:RETRY\_EXHAUSTED will be generated
- C. HTTP status 401 unauthorized for policy violation
- D. HTTP status 400 from the rate-limiting policy violation since the call does not reach the back-end

**Answer:** A

**NEW QUESTION 119**

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 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 123**

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

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

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

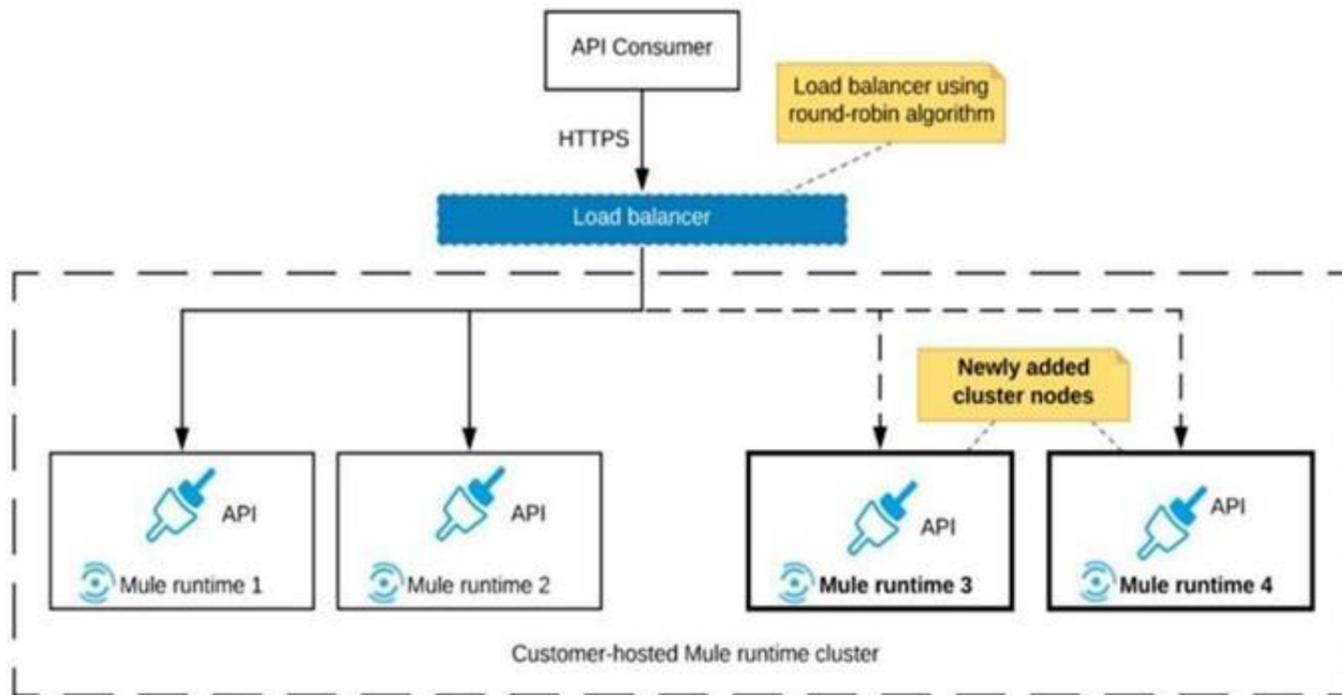
**Answer:** A

**Explanation:**

SLAs and non-functional requirements to access the backend systems. Only this option actually speaks to design parameters and reqs. \* Below two are technical implementations and not the part of design: - Snapsho of the Mule application's flows, including their error handling - The credentials to access the backend systems and contact details for the administrator of each system \* List of consumers is not relevant to the design

**NEW QUESTION 128**

Refer to the exhibit.



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

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

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

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

Answer: D

**NEW QUESTION 129**

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

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

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

Answer: D

**NEW QUESTION 133**

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

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

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

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

Answer: D

**NEW QUESTION 137**

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

**NEW QUESTION 141**

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

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

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

**Answer: D**

**Explanation:**

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

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

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

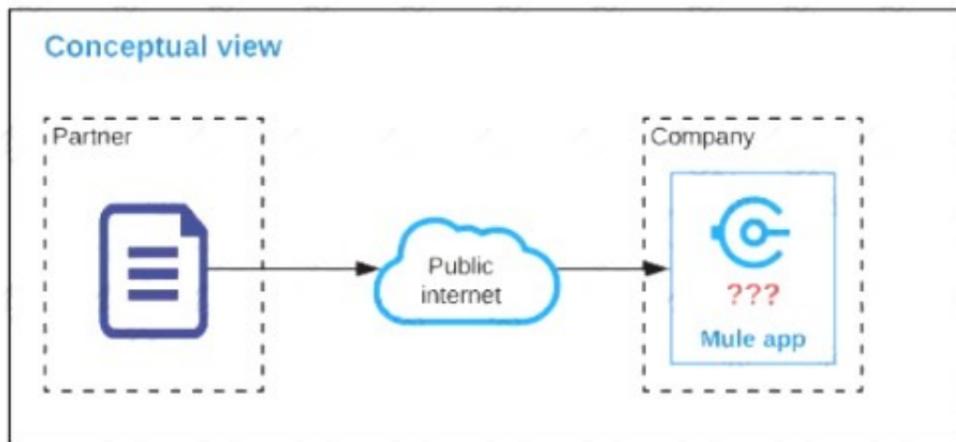
Reference: <https://www.oasis-open.org/committees/download.php/16768/wss-v1.1-spec-os-SAMLTokenProfile.pdf>

**NEW QUESTION 145**

Refer to the exhibit.

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

What Anypoint Connector can be used in the organization's Mule application to securely receive data from this external business partner?



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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 150**

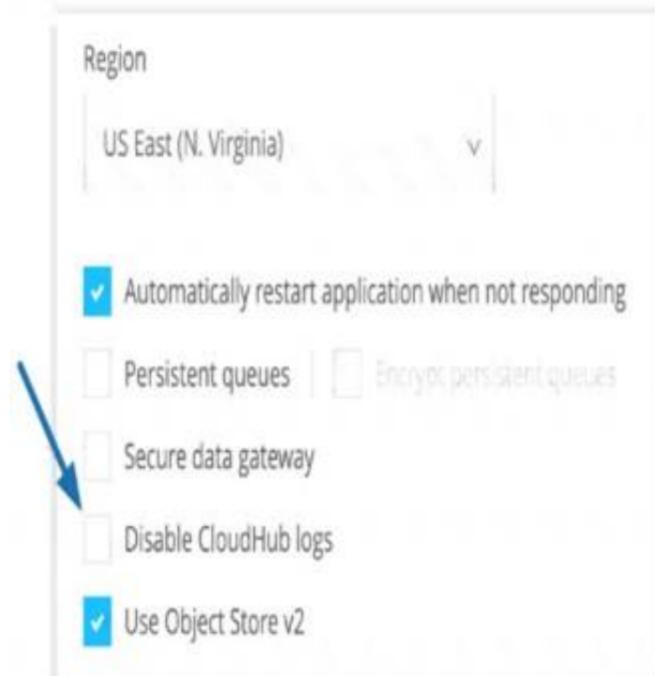
What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

- 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

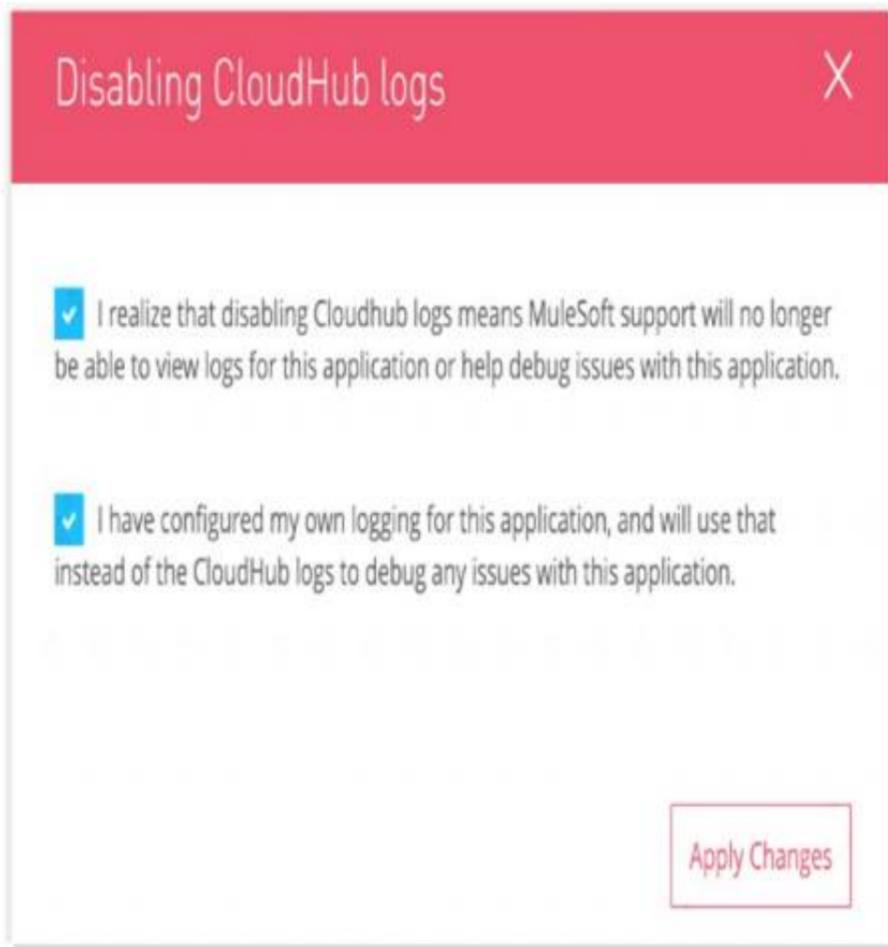
**Answer: D**

**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 152**

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 153**

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. muleInclude
- D. \_muleExclude

**Answer: D**

**NEW QUESTION 156**

To implement predictive maintenance on its machinery equipment, ACME Tractors has installed thousands of IoT sensors that will send data for each machinery asset as sequences of JMS messages, in near real-time, to a JMS queue named SENSOR\_DATA on a JMS server. The Mule application contains a JMS Listener operation configured to receive incoming messages from the JMS servers SENSOR\_DATA JMS queue. The Mule application persists each received JMS message, then sends a transformed version of the corresponding Mule event to the machinery equipment back-end systems. The Mule application will be deployed to a multi-node, customer-hosted Mule runtime cluster. Under normal conditions, each JMS message should be processed exactly once. How should the JMS Listener be configured to maximize performance and concurrent message processing of the JMS queue?

- A. Set numberOfConsumers = 1 Set primaryNodeOnly = false
- B. Set numberOfConsumers = 1 Set primaryNodeOnly = true
- C. Set numberOfConsumers to a value greater than one Set primaryNodeOnly = true
- D. Set numberOfConsumers to a value greater than one Set primaryNodeOnly = false

**Answer: D**

**NEW QUESTION 159**

What API policy would LEAST likely be applied to a Process API?

- A. Custom circuit breaker
- B. Client ID enforcement
- C. Rate limiting
- D. JSON threat protection

**Answer:** D

**Explanation:**

Key to this question lies in the fact that Process API are not meant to be accessed directly by clients. Lets analyze options one by one. Client ID enforcement : This is applied at process API level generally to ensure that identity of API clients is always known and available for API-based analytics Rate Limiting : This policy is applied on Process Level API to secure API's against degradation of service that can happen in case load received is more than it can handle Custom circuit breaker : This is also quite useful feature on process level API's as it saves the API client the wasted time and effort of invoking a failing API. JSON threat protection : This policy is not required at Process API and rather implemented as Experience API's. This policy is used to safeguard application from malicious attacks by injecting malicious code in JSON object. As ideally Process API's are never called from external world , this policy is never used on Process API's Hence correct answer is

JSON threat protection MuleSoft Documentation Reference : <https://docs.mulesoft.com/api-manager/2.x/policy-mule3-json-threat>

**NEW QUESTION 164**

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

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

**Answer:** D

**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 166**

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 170**

An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub.

For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub's default zero-downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers.This is expected to result in several Mule application deployments per hour.

- A. 10.0.0.0/21(2048 IPs)
- B. 10.0.0.0/22(1024IPs)
- C. 10.0.0.0/23(512 IPs)
- D. 10.0.0.0/24(256 IPs)

**Answer:** A

**Explanation:**

\* When you create an Anypoint VPC, the range of IP addresses for the network must be specified in the form of a Classless Inter-Domain Routing (CIDR) block, using CIDR notation.

\* This address space is reserved for Mule workers, so it cannot overlap with any address space used in your data center if you want to peer it with your VPC.

\* To calculate the proper sizing for your Anypoint VPC, you first need to understand that the number of dedicated IP addresses is not the same as the number of workers you have deployed.

\* For each worker deployed to CloudHub, the following IP assignment takes place: For better fault tolerance, the VPC subnet may be divided into up to four Availability Zones.

\* A few IP addresses are reserved for infrastructure. At least two IP addresses per worker to perform at zero-downtime.

\* Hence in this scenario 2048 IP's are required to support the requirement.

#### NEW QUESTION 171

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 responses Then 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 characteristics Then 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 system
- E. Then upscale performance results to full production scale

**Answer: C**

#### NEW QUESTION 174

An external web UI application currently accepts occasional HTTP requests from client web browsers to change (insert, update, or delete) inventory pricing information in an inventory system's database. Each inventory pricing change must be transformed and then synchronized with multiple customer experience systems in near real-time (in under 10 seconds). New customer experience systems are expected to be added in the future.

The database is used heavily and limits the number of SELECT queries that can be made to the database to 10 requests per hour per user.

What is the most scalable, idiomatic (used for its intended purpose), decoupled, reusable, and maintainable integration mechanism available to synchronize each inventory pricing change with the various customer experience systems in near real-time?

- A. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the watermark attribute set to an appropriate database column In the same flow, use a Scatter-Gather to call each customer experience system's REST API with transformed inventory-pricing records
- B. Add a trigger to the inventory-pricing database table so that for each change to the inventory pricing database, a stored procedure is called that makes a REST call to a Mule application Write the Mule application to publish each Mule event as a message to an Anypoint MQ exchange Write other Mule applications to subscribe to the Anypoint MQ exchange, transform each received message, and then update the Mule application's corresponding customer experience system(s)
- C. Replace the external web UI application with a Mule application to accept HTTP requests from client web browsers In the same Mule application, use a Batch Job scope to test if the database request will succeed, aggregate pricing changes within a short time window, and then update both the inventory pricing database and each customer experience system using a Parallel For Each scope
- D. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the ID attribute set to an appropriate database column In the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queue Write other Mule applications to subscribe to the Anypoint MQ queue, transform each received message, and then update the Mule application's corresponding customer experience system(s)

**Answer: B**

#### NEW QUESTION 179

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

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

**Answer: A**

#### Explanation:

Correct answer is "Compile, package, unit test, validate unit test coverage, deploy"

Anypoint Platform supports continuous integration and continuous delivery using industry standard tools Mule Maven Plugin The Mule Maven plugin can automate building, packaging and deployment of Mule applications from source projects Using the Mule Maven plugin, you can automate your Mule application deployment to CloudHub, to Anypoint Runtime Fabric, or on-premises, using any of the following deployment strategies • CloudHub deployment • Runtime Fabric deployment • Runtime Manager REST API deployment • Runtime Manager agent deployment MUnit Maven Plugin The MUnit Maven plugin can automate test execution, and ties in with the Mule Maven plugin. It provides a full suite of integration and unit test capabilities, and is fully integrated with Maven and Surefire for integration with your continuous deployment environment. Since MUnit 2.x, the coverage report goal is integrated with the maven reporting section. Coverage Reports are generated during Maven's site lifecycle, during the coverage-report goal. One of the features of MUnit Coverage is to fail the build if a certain coverage level is not reached. MUnit is not used for integration testing Also publishing to Anypoint Exchange or to create associated API instances in API Manager is not a part of CI/CD pipeline which can be achieved using mulesoft provided maven plugin

Explanation

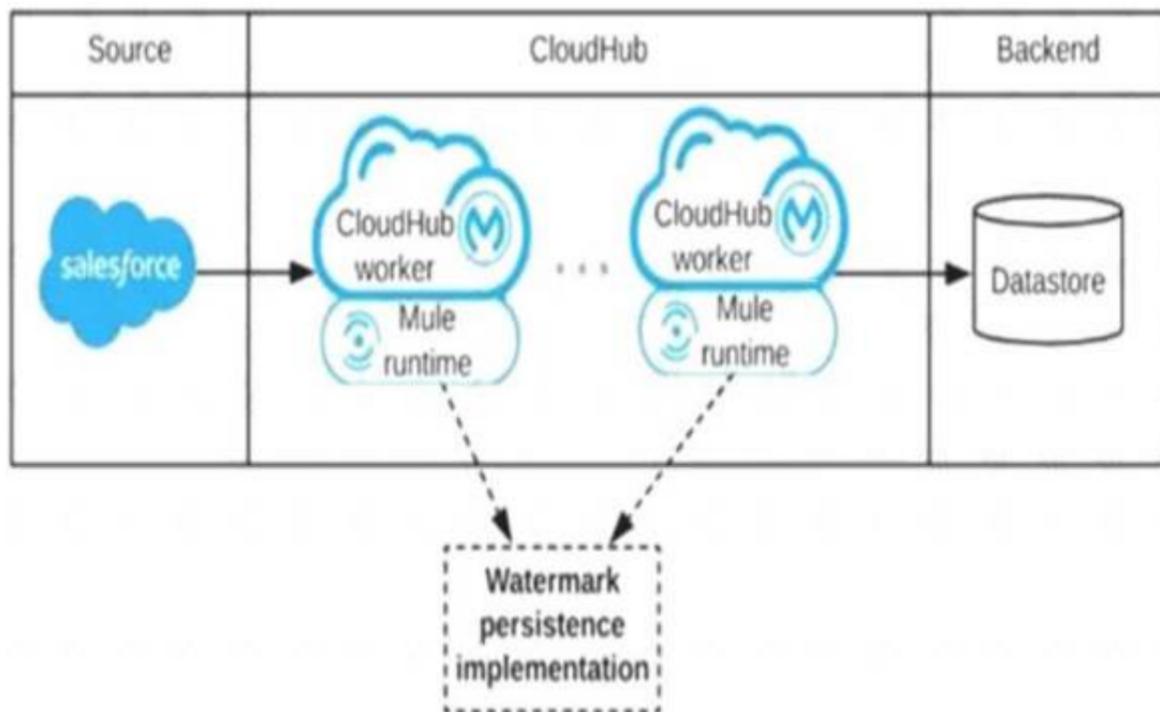
Architecture mentioned in the question can be diagrammatically put as below. Persistent Object Store is the correct answer .

\* Mule Object Stores: An object store is a facility for storing objects in or across Mule applications. Mule uses object stores to persist data for eventual retrieval.

Mule provides two types of object stores:

- 1) In-memory store – stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime. So we can't use in memory store in our scenario as we want to share watermark within all cloudhub workers
- 2) Persistent store – Mule persists data when an object store is explicitly configured to be persistent. Hence this watermark will be available even any of the worker goes down

Diagram Description automatically generated



**NEW QUESTION 182**

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

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

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

**Answer: D**

**Explanation:**

\* Sharing Mule applications as templates is a great way to share your work with other people who are in your organization in Anypoint Platform. When they need to build a similar application they can create the mule application using the template project from Anypoint studio.  
 \* Anypoint Templates are designed to make it easier and faster to go from a blank canvas to a production application. They're bit for bit Mule applications requiring only Anypoint Studio to build and design, and are deployable both on-premises and in the cloud.  
 \* Anypoint Templates are based on five common data Integration patterns and can be customized and extended to fit your integration needs. So even if your use case involves different endpoints or connectors than those included in the template, they still offer a great starting point.  
 Some of the best practices while creating the template project: - Define the common error handler as part of template project, either using pom dependency or mule config file - Define common logger/audit framework as part of the template project - Define the env specific properties and secure properties file as per the requirement - Define global.xml for global configuration - Define the config file for connector configuration like Http,Salesforce,File,FTP etc - Create separate folders to create DWL,Properties,SSL certificates etc - Add the dependency and configure the pom.xml as per the business need - Configure the mule-artifact.json as per the business need

**NEW QUESTION 184**

Anypoint Exchange is required to maintain the source code of some of the assets committed to it, such as Connectors, Templates, and API specifications.

What is the best way to use an organization's source-code management (SCM) system in this context?

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

**Answer: B**

**Explanation:**

\* Organization should continue to use SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching.  
 \* Reason is that Anypoint exchange is not full fledged version repositories like GitHub.  
 \* But at same time it is tightly coupled with Mule assets

**NEW QUESTION 186**

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 189**

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 193**

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 197**

A customer wants to use the mapped diagnostic context (MDC) and logging variables to enrich its logging and improve tracking by providing more context in the logs.

The customer also wants to improve the throughput and lower the latency of message processing.

As an Mulesoft integration architect can you advise, what should the customer implement to meet these requirements?

- A. Use synchronous logging and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
- B. Useasync logger at the level greater than INFO and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
- C. Useasync logger at the level equal to DEBUG orTRACEand use patternlayout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
- D. Use synchronous logging at the INFO DEBUG or Trace level and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables

**Answer:** B

**NEW QUESTION 202**

When the mule application using VM is deployed to a customer-hosted cluster or multiple cloudhub workers, how are messages consumed by the Mule engine?

- A. in non-deterministic way
- B. by starting an XA transaction for each new message
- C. in a deterministic way
- D. the primary only in order to avoid duplicate processing

**Answer:** C

**NEW QUESTION 204**

An insurance provider is implementing Anypoint platform to manage its application infrastructure and is using the customer hosted runtime for its business due to certain financial requirements it must meet. It has built a number of synchronous API's and is currently hosting these on a mule runtime on one server

These applications make use of a number of components including heavy use of object stores and VM queues. Business has grown rapidly in the last year and the insurance provider is starting to receive reports of reliability issues from its applications.

The DevOps team indicates that the API's are currently handling too many requests and this is over loading the server. The team has also mentioned that there is

a significant downtime when the server is down for maintenance.  
 As an integration architect, which option would you suggest to mitigate these issues?

- A. Add a load balancer and add additional servers in a server group configuration
- B. Add a load balancer and add additional servers in a cluster configuration
- C. Increase physical specifications of server CPU memory and network
- D. Change applications by use an event-driven model

**Answer: B**

**NEW QUESTION 207**

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

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

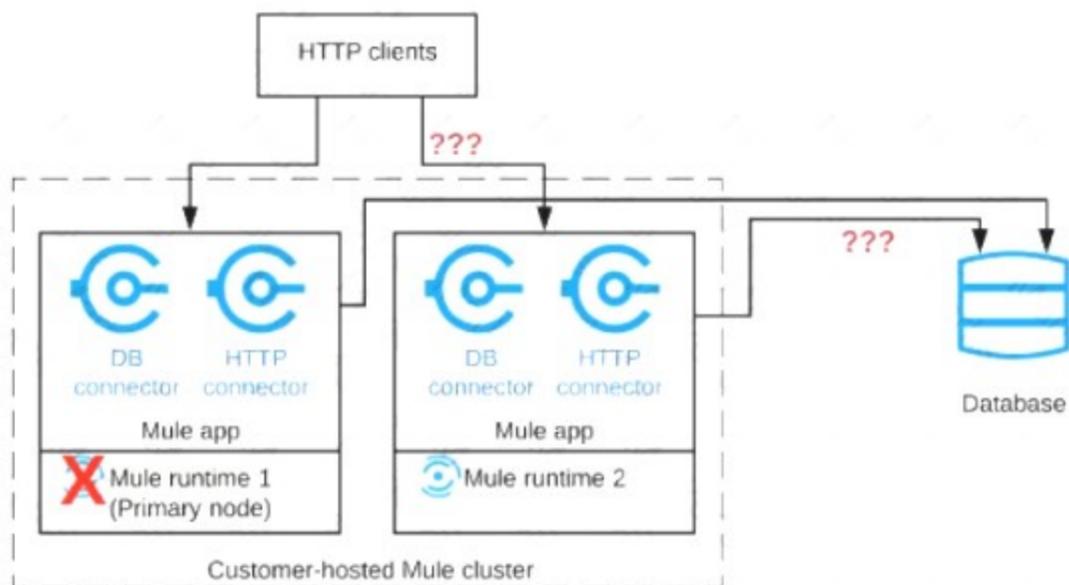
**Answer: A**

**Explanation:**

Context of the question is about managing and governing mule applications deployed on Anypoint platform. Anypoint API Manager (API Manager) is a component of Anypoint Platform that enables you to manage, govern, and secure APIs. It leverages the runtime capabilities of API Gateway and Anypoint Service Mesh, both of which enforce policies, collect and track analytics data, manage proxies, provide encryption and authentication, and manage applications.  
 Mule Ref Doc : <https://docs.mulesoft.com/api-manager/2.x/getting-started-proxy>

**NEW QUESTION 210**

Refer to the exhibit.



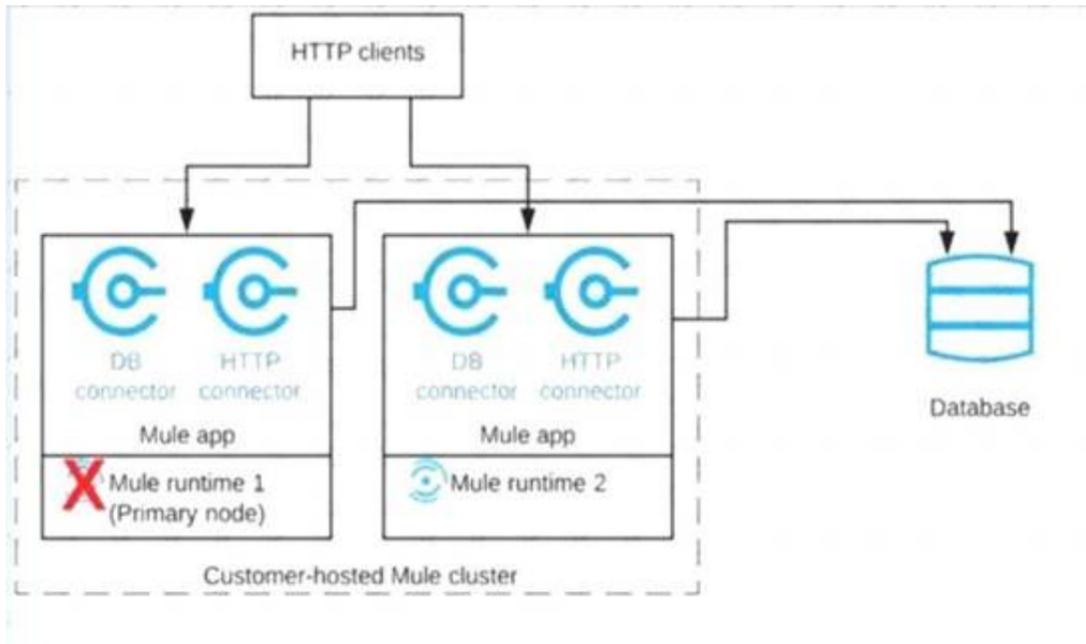
A Mule application is deployed to a cluster of two customer-hosted Mule runtimes. The Mule application has a flow that polls a database and another flow with an HTTP Listener. HTTP clients send HTTP requests directly to individual cluster nodes. What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has failed, 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 described 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 214**

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 218**

What limits if a particular Anypoint Platform user can discover an asset in Anypoint Exchange?

- A. Design Center and RAML were both used to create the asset
- B. The existence of a public Anypoint Exchange portal to which the asset has been published
- C. The type of the asset in Anypoint Exchange
- D. The business groups to which the user belongs

**Answer: D**

**Explanation:**

\* "The existence of a public Anypoint Exchange portal to which the asset has been published" - question does not mention anything about the public portal. Beside the public portal is open to the internet, to anyone. \* If you cannot find an asset in the current business group scopes, search in other scopes. In the left navigation bar click All assets (assets provided by MuleSoft and your own master organization), Provided by MuleSoft, or a business group scope. User belonging to one Business Group can see assets related to his group only Reference: <https://docs.mulesoft.com/exchange/to-find-info> <https://docs.mulesoft.com/exchange/asset-details> Correct answer is The business groups to which the user belongs

**NEW QUESTION 223**

How are the API implementation , API client, and API consumer combined to invoke and process an API ?

- A. The API consumer creates an API implementation , which receives API invocations from an API such that they are processed for an API client
- B. The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation
- C. An API client creates an API consumer, which receives API invocation from an API such that they are processed for an API implementation
- D. The API client creates an API consumer which sends API invocations to an API such that they are processed by API implementation

**Answer: C**

**Explanation:**

The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation This is based on below definitions API client • An application component • that accesses a service • by invoking an API of that service - by definition of the term API over HTTP API consumer • A business role, which is often assigned to an individual • that develops API clients, i.e., performs the activities necessary for enabling an API client to invoke APIs API implementation • An application component • that implements th functionality

**NEW QUESTION 224**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

### Money Back Guarantee

#### **MCIA-Level-1 Practice Exam Features:**

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