

Exam Questions MCI A-Level-1

MuleSoft Certified Integration Architect - Level 1

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



NEW QUESTION 1

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

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

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

Answer: A

NEW QUESTION 2

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

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

Answer: A

NEW QUESTION 3

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 4

A Mule application is being designed for deployment to a single CloudHub worker. The Mule application will have a flow that connects to a SaaS system to perform some operations each time the flow is invoked.

The SaaS system connector has operations that can be configured to request a short-lived token (fifteen minutes) that can be reused for subsequent connections within the fifteen minute time window. After the token expires, a new token must be requested and stored.

What is the most performant and idiomatic (used for its intended purpose) Anypoint Platform component or service to use to support persisting and reusing tokens in the Mule application to help speed up reconnecting the Mule application to the SaaS application?

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

Answer: D

NEW QUESTION 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

A finance giant is planning to migrate all its Mule applications to Runtime fabric (RTF). Currently all Mule applications are deployed cloud hub using automated CI/CD scripts.
 As an integration architect, which of the below step would you suggest to ensure that the applications from cloudhub are migrated properly to Runtime Fabric (RTF) with an assumption that organization is keen on keeping the same deployment strategy.

- A. No changes need to be made to POM.xml file and CI/CD script should be modified as per the RTF configurations
- B. runtimeFabric dependency should be added as a mule plug-in to POM.xml file and CI/CD script should be modified as per the RTF configurations
- C. runtimeFabric deployment should be added to POM.xml file in all the mule applications and CI/CD script should be modified as per the RTF configurations
- D. runtimeFabric profile should be added mule configuration files in the mule applications and CI/CD script should be modified as per the RTF configurations

Answer: C

NEW QUESTION 8


A company is using Mulesoft to develop API's and deploy them to Cloudhub and on premises targets. Recently it has decided to enable Runtime Fabric deployment option as well and infrastructure is set up for this option.
 What can be used to deploy Runtime Fabric?

- A. AnypointCLI
- B. Anypoint platform REST API's
- C. Directly uploading ajar file from the Runtime manager
- D. Mule maven plug-in

Answer: D

NEW QUESTION 9

Refer to the exhibit.



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

What is the most effective way to design for these requirements in order to minimize the impact of future change?


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

Answer: C

Explanation:

* Scatter-Gather is used for parallel processing, to improve performance. In this scenario, input files are coming from different vendors so mostly at different times. Goal here is to minimize the impact of future change. So scatter Gather is not the correct choice.
 * If we use 1 API to receive all files from different Vendors, any new vendor addition will need changes to that 1 API to accommodate new requirements. So Option A and C are also ruled out.
 * Correct answer is Create an API that receives the file and invokes a Process API with the data contained in the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed. Answer to this question lies in the API led connectivity approach.
 * API-led connectivity is a methodical way to connect data to applications through a series of reusable and purposeful modern APIs that are each developed to play a specific role – unlock data from systems, compose data into processes, or deliver an experience. System API : System API tier, which provides consistent, managed, and secure access to backend systems. Process APIs : Process APIs take core assets and combines them with some business logic to create a higher level of value. Experience APIs : These are designed specifically for consumption by a specific end-user app or device.
 So in case of any future plans , organization can only add experience API on addition of new Vendors, which reuse the already existing process API. It will keep impact minimal.

Diagram Description automatically generated



NEW QUESTION 10

An organization is designing an integration solution to replicate financial transaction data from a legacy system into a data warehouse (DWH). The DWH must contain a daily snapshot of financial transactions, to be delivered as a CSV file. Daily transaction volume exceeds tens of millions of records, with significant spikes in volume during popular shopping periods. What is the most appropriate integration style for an integration solution that meets the organization's current requirements?

- A. Event-driven architecture
- B. Microservice architecture
- C. API-led connectivity
- D. Batch-triggered ETL

Answer: D

Explanation:

Correct answer is Batch-triggered ETL Within a Mule application, batch processing provides a construct for asynchronously processing larger-than-memory data sets that are split into individual records. Batch jobs allow for the description of a reliable process that automatically splits up source data and stores it into persistent queues, which makes it possible to process large data sets while providing reliability. In the event that the application is redeployed or Mule crashes, the job execution is able to resume at the point it stopped.

NEW QUESTION 10

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

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


Answer: C

Explanation:

* To send message or receive JMS (Java Message Service) message no separate network connection need to be established. So option A, C and D are ruled out. Correct Answer The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider.
 * JMS Connector enables sending and receiving messages to queues and topics for any message service that implements the JMS specification.
 * JMS is a widely used API for message-oriented middleware.
 * It enables the communication between different components of a distributed application to be loosely coupled, reliable, and asynchronous.

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

Diagram, text Description automatically generated



NEW QUESTION 11

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

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

Answer: A

NEW QUESTION 12

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

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

Answer: D

NEW QUESTION 13

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

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

Answer: B

Explanation:

High availability is about up-time of your application.

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

NEW QUESTION 18

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

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

Answer: A

Explanation:

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

NEW QUESTION 23

A Mule application is being designed to do the following:

Step 1: Read a SalesOrder message from a JMS queue, where each SalesOrder consists of a header and a list of SalesOrderLineItems.

Step 2: Insert the SalesOrder header and each SalesOrderLineItem into different tables in an RDBMS.

Step 3: Insert the SalesOrder header and the sum of the prices of all its SalesOrderLineItems into a table in a different RDBMS.

No SalesOrder message can be lost and the consistency of all SalesOrder-related information in both RDBMSs must be ensured at all times.

What design choice (including choice of transactions) and order of steps addresses these requirements?

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

Answer: A

Explanation:

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

Option D says Perform BOTH DB inserts in ONE DB transaction.

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

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

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

Additional Information about transactions:

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

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

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

Use JMS ack if

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

NEW QUESTION 24

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

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

Answer: C

NEW QUESTION 26

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

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

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

Answer: D

NEW QUESTION 28

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 33

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

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

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

Answer: D

Explanation:

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

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

Hence the correct option is file based repeatable stream

NEW QUESTION 35

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 40

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 Until Successful scope can be implemented while calling backend API's. The Until Successful scope repeatedly triggers the scope's components (including flow references) until they all succeed or until a maximum number of retries is exceeded. The scope provides 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 41

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

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

Answer: AB

NEW QUESTION 43

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 47

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 52

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 54

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 56

A new Mule application under development must implement extensive data transformation logic. Some of the data transformation functionality is already available as external transformation services that are mature and widely used across the organization; the rest is highly specific to the new Mule application.

The organization follows a rigorous testing approach, where every service and application must be extensively acceptance tested before it is allowed to go into production.

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

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


Answer: D

Explanation:

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

NEW QUESTION 59

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 64

An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders. The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS).

At a minimum, what must be stored in each truststore and keystore of these two Mule applications to properly support two-way TLS between the two Mule applications while properly protecting each Mule application's keys?

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

Answer: C

NEW QUESTION 68

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

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

Answer: C

Explanation:

Correct answer is Workers are randomly distributed across available AZs within that region. This ensure high availability for deployed mule applications Mulesoft documentation reference :

<https://docs.mulesoft.com/runtime-manager/cloudhub-hadr>

NEW QUESTION 70

A company is implementing a new Mule application that supports a set of critical functions driven by a rest API enabled, claims payment rules engine hosted on oracle ERP. As designed the mule application requires many data transformation operations as it performs its batch processing logic.

The company wants to leverage and reuse as many of its existing java-based capabilities (classes, objects, data model etc.) as possible

What approach should be considered when implementing required data mappings and transformations between Mule application and Oracle ERP in the new Mule application?

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

Answer: C

NEW QUESTION 71

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

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

Answer: B

Explanation:

An SLA for the upstream API CANNOT be provided.

NEW QUESTION 72

What best describes the Fully Qualified Domain Names (FQDNs), also known as DNS entries, created when a Mule application is deployed to the CloudHub Shared Worker Cloud?

- A. A fixed number of FQDNs are created, IRRESPECTIVE of the environment and VPC design
- B. The FQDNs are determined by the application name chosen, IRRESPECTIVE of the region
- C. The FQDNs are determined by the application name, but can be modified by an administrator after deployment
- D. The FQDNs are determined by both the application name and the region

Answer: D

Explanation:

Every Mule application deployed to CloudHub receives a DNS entry pointing to the CloudHub. The DNS entry is a CNAME for the CloudHub Shared Load Balancer in the region to which the Mule application is deployed. When we deploy the application on CloudHub, we get a generic url to access the endpoints. Generic URL looks as below:

<application-name>.<region>.cloudhub.io <application-name> is the deployed application name which is unique across all the MuleSoft clients. <region> is the region name in which an application is deployed.

The public CloudHub (shared) load balancer already redirects these requests, where myApp is the name of the Mule application deployment to CloudHub: HTTP requests to http://myApp.


<region>.cloudhub.io redirects to

http://mule-worker-myApp.<region>.cloudhub.io:8081

HTTPS traffic to https://myApp.<region>.cloudhub.io redirects to https://mule-worker-myApp.<region>.cloudhub.io:8082

NEW QUESTION 76

Refer to the exhibit.



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

- Step 1: The JMS Correlation ID header is read from the received JMS message.
- Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.
- Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.
- Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.
- Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.

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

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

Answer: C

Explanation:

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

* As the SOAP service is idempotent in nature. In case of any failures , using this Correlation ID saved in first step we can make call to SOAP service and validate the Correlation ID.

Top of Form


Additional Information:

* Competing Consumers

are multiple consumers that are all created to receive messages from a single

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

Diagram Description automatically generated



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

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

Table Description automatically generated

IDEMPOTENCE


WHEN PERFORMING AN OPERATION AGAIN GIVES THE SAME RESULT

HTTP METHOD	IDEMPOTENCE	SAFETY
GET	YES	YES
HEAD	YES	YES
PUT	YES	NO
DELETE	YES	NO
POST	NO	NO
PATCH	NO	NO

Bottom of Form

NEW QUESTION 81

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 84

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

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

Answer: B

Explanation:

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

NEW QUESTION 86

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

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

Answer: D

NEW QUESTION 87

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

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

Answer: A

NEW QUESTION 91

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

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

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 93

An auto mobile company want to share inventory updates with dealers D1 and D2 asynchronously and concurrently via queues Q1 and Q2. Dealer D1 must consume the message from the queue Q1 and dealer D2 to must consume a message from the queue Q2.

Dealer D1 has implemented a retry mechanism to reprocess the transaction in case of any errors while processing the inventers updates. Dealer D2 has not implemented any retry mechanism.

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

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

Answer: D

NEW QUESTION 97

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

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

Answer: CD

NEW QUESTION 99

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 102

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 103

A rate limiting policy has been applied to a soap V1.2 API published in Clondhub. 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 H1TP 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 106

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

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

Answer: B

NEW QUESTION 111

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 112

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>https://docs.mulesoft.com/runtime-manager/ib-whitelists</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 113

An organization uses Mule runtimes which are managed by Anypoint Platform - Private Cloud Edition. What MuleSoft component is responsible for feeding analytics data to non-MuleSoft analytics platforms?

- A. Anypoint Exchange
- B. The Mule runtimes
- C. Anypoint API Manager
- D. Anypoint Runtime Manager

Answer: D

Explanation:

Correct answer is Anypoint Runtime Manager

MuleSoft Anypoint Runtime Manager (ARM) provides connectivity to Mule Runtime engines deployed across your organization to provide centralized management, monitoring and analytics reporting. However, most enterprise customers find it necessary for these on-premises runtimes to integrate with their existing non MuleSoft analytics / monitoring systems such as Splunk and ELK to support a single pane of glass view across the infrastructure.

* You can configure the Runtime Manager agent to export data to external analytics tools.

Using either the Runtime Manager cloud console or Anypoint Platform Private Cloud Edition, you can:


--> Send Mule event notifications, including flow executions and exceptions, to Splunk or ELK.

--> Send API Analytics to Splunk or ELK. Sending data to third-party tools is not supported for applications deployed on CloudHub.

You can use the CloudHub custom log appender to integrate with your logging system. Reference: <https://docs.mulesoft.com/runtime-manager/>

<https://docs.mulesoft.com/release-notes/runtime-manager-agent/runtime-manager-agent-release-notes>

Diagram Description automatically generated




Additional Info:

It can be achieved in 3 steps:

- 1) register an agent to a runtime manager,
- 2) configure a gateway to enable API analytics to be sent to non MuleSoft analytics platform (Splunk for ex.) – as highlighted in the following diagram and
- 3) setup dashboards.

Diagram Description automatically generated



NEW QUESTION 117

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

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

Answer: C

Explanation:

Connected Apps

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

Mule Ref Doc : <https://docs.mulesoft.com/access-management/connected-apps-overview>

NEW QUESTION 121

As a part of business requirement , old CRM system needs to be integrated using Mule application. CRM system is capable of exchanging data only via SOAP/HTTP protocol. As an integration architect who follows API led approach , what is the the below step you will perform so that you can share document with CRM team?

- A. Create RAML specification using Design Center
- B. Create SOAP API specification using Design Center
- C. Create WSDL specification using text editor
- D. Create WSDL specification using Design Center

Answer: C

Explanation:

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

NEW QUESTION 125

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

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

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

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

Answer: C

Explanation:

* "100% increase in the throughput of the API" might look correct, as the number of requests processed per second might increase, but is it guaranteed to increase by 100%? Using 4 nodes will definitely increase throughput of system. But it is cant be precisely said if there would be 100% increase in throughput as it depends on many other factors. Also it is nowhere mentioned in the description that all nodes have same CPU/memory assigned. The question is about the guaranteed behavior * Increasing number of nodes will have no impact on response time as we are scaling application horizontally and not vertically. Similarly there is no change in JVM heap memory usage. * So Correct answer is 50% reduction in the number of requests being received by each node This is because of the two reasons. 1) API is mentioned as stateless 2) Load Balancer is used

NEW QUESTION 126

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 130

A corporation has deployed multiple mule applications implementing various public and private API's to different cloudhub workers. These API's arc Critical applications that must be highly available and in line with the reliability SLA as defined by stakeholders.

How can API availability (liveliness or readiness) be monitored so that Ops team receives outage notifications?

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

Answer: A

NEW QUESTION 135

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

- A. When there is one deployment of the API implementation to CloudHub and another one to customer hosted mule runtime that must share the cache state.
- B. When there are two CloudHub deployments of the API implementation by two Anypoint Platform business groups to the same CloudHub region that must share the cache state.
- C. When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.
- D. When there are three CloudHub deployments of the API implementation to three separate CloudHub regions that must share the cache state.

Answer: C

Explanation:

Object Store Connector is a Mule component that allows for simple key-value storage. Although it can serve a wide variety of use cases, it is mainly design for: - Storing synchronization information, such as watermarks. - Storing temporal information such as access tokens. - Storing user information. Additionally, Mule Runtime uses Object Stores to support some of its own components, for example: - The Cache module uses an Object Store to maintain all of the cached data. - The OAuth module (and every OAuth enabled connector) uses Object Stores to store the access and refresh tokens. Object Store data is in the same region as


the worker where the app is initially deployed. For example, if you deploy to the Singapore region, the object store persists in the Singapore region. MuleSoft Reference : <https://docs.mulesoft.com/object-store-connector/1.1/> Data can be shared between different instances of the Mule application. This is not recommended for Inter Mule app communication. Coming to the question, object store cannot be used to share cached data if it is deployed as separate Mule applications or deployed under separate Business Groups. Hence correct answer is When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.

NEW QUESTION 137

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 138

What condition requires using a CloudHub Dedicated Load Balancer?

- A. When cross-region load balancing is required between separate deployments of the same Mule application
- B. When custom DNS names are required for API implementations deployed to customer-hosted Mule runtimes
- C. When API invocations across multiple CloudHub workers must be load balanced
- D. When server-side load-balanced TLS mutual authentication is required between API implementations and API clients

Answer: D

Explanation:

Correct answer is When server-side load-balanced TLS mutual authentication is required between API implementations and API clients CloudHub dedicated load balancers (DLBs) are an optional component of Anypoint Platform that enable you to route external HTTP and HTTPS traffic to multiple Mule applications deployed to CloudHub workers in a Virtual Private Cloud (VPC). Dedicated load balancers enable you to:

- * Handle load balancing among the different CloudHub workers that run your application.
- * Define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
- * Configure proxy rules that map your applications to custom domains. This enables you to host your applications under a single domain

NEW QUESTION 139

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 140

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 145

A Mule application uses the Database connector. What condition can the Mule application automatically adjust to or recover from without needing to restart or redeploy the Mule application?

- A. One of the stored procedures being called by the Mule application has been renamed
- B. The database server was unavailable for four hours due to a major outage but is now fully operational again
- C. The credentials for accessing the database have been updated and the previous credentials are no longer valid
- D. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade

Answer: B

Explanation:

* Any change in the application will require a restart except when the issue outside the app. For below situations, you would need to redeploy the code after doing necessary changes

- One of the stored procedures being called by the Mule application has been renamed. In this case, in the Mule application you will have to do changes to accommodate the new stored procedure name.
- Required redesign of Mule applications to follow microservice architecture principles. As code is changed, deployment is must
- If the credentials changed and you need to update the connector or the properties.
- The credentials for accessing the database have been updated and the previous credentials are no longer valid. In this situation you need to restart or redeploy depending on how credentials are configured in Mule application.

* So Correct answer is The database server was unavailable for four hours due to a major outage but is now fully operational again as this is the only external issue to application.

NEW QUESTION 146

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 147

An organization is migrating all its Mule applications to Runtime Fabric (RTF). None of the Mule applications use Mule domain projects. Currently, all the Mule applications have been manually deployed to a server group among several customer hosted Mule runtimes. Port conflicts between these Mule application deployments are currently managed by the DevOps team who carefully manage Mule application properties files. When the Mule applications are migrated from the current customer-hosted server group to Runtime Fabric (RTF), for the Mule applications need to be rewritten and what DevOps port configuration responsibilities change or stay the same?

- A. Yes, the Mule applications Must be rewritten DevOps No Longer needs to manage port conflicts between the Mule applications
- B. Yes, the Mule applications Must be rewritten DevOps Must Still Manage port conflicts.
- C. NO, The Mule applications do NOT need to be rewritten DevOps MUST STILL manage port conflicts
- D. NO, the Mule applications do NO need to be rewritten DevOps NO LONGER needs to manage port conflicts between the Mule applications.

Answer: C

Explanation:

* Anypoint Runtime Fabric is a container service that automates the deployment and orchestration of your Mule applications and gateways.

* Runtime Fabric runs on customer-managed infrastructure on AWS, Azure, virtual machines (VMs) or bare-metal servers.

* As none of the Mule applications use Mule domain projects. applications are not required to be rewritten. Also when applications are deployed on RTF, by default ingress is allowed only on 8081.

* Hence port conflicts are not required to be managed by DevOps team

NEW QUESTION 152

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

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

Answer: D

Explanation:

Correct answer is using below set of activities Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing We will see why this is correct answer but before that lets understand few of the concepts which we need to know. Until Successful Scope The Until Successful scope processes messages through its processors until the entire operation succeeds. Until Successful repeatedly retries to process a message that is attempting to complete an activity such as: - Dispatching to outbound endpoints, for example, when calling a remote web service that may have availability issues. - Executing a component method, for example, when executing on a Spring bean that may depend on unreliable resources. - A sub-flow execution, to keep re-executing several actions until they all succeed, - Any other message processor execution, to allow more complex scenarios. How this will help requirement : Using Until Successful Scope we can retry sending the order to backend systems in case of error to avoid manual processing later. Retry values can be configured in Until Successful Scope Apache ActiveMQ It is an open source message broker written in Java together with a full Java Message Service client ActiveMQ has the ability to deliver

messages with delays thanks to its scheduler. This functionality is the base for the broker redelivery plug-in. The redelivery plug-in can intercept dead letter processing and reschedule the failing messages for redelivery. Rather than being delivered to a DLQ, a failing message is scheduled to go to the tail of the original queue and redelivered to a message consumer. How this will help requirement : If backend application is down for a longer duration where Until Successful Scope wont work, then we can make use of ActiveMQ long retry Queue. The redelivery plug-in can intercept dead letter processing and reschedule the failing messages for redelivery. Mule Reference:

<https://docs.mulesoft.com/mule-runtime/4.3/migration-core-until-successful>

NEW QUESTION 156

49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58.

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

The retailer has obtained a public X.509 certificate for the name api.acme.com, signed by a reputable CA, to be used as the server certificate.

Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

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

Answer: C

Explanation:

* An X.509 certificate is a vital safeguard against malicious network impersonators. Without x.509 server authentication, man-in-the-middle attacks can be initiated by malicious access points, compromised routers, etc.

* X.509 is most used for SSL/TLS connections to ensure that the client (e.g., a web browser) is not fooled by a malicious impersonator pretending to be a known, trustworthy website.

* Coming to the question , we can not use SLB here as SLB does not allow to define vanity domain names. * Hence we need to use DLB and add certificate in there

Hence correct answer is Add the X 509 certificate to the cloudhub Dedicated Load Balancer (DLB), not the Mule application. Create the CNAME for api.acme.com pointing to the DLB's record

NEW QUESTION 160

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 columnIn the same now, 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 applicationWrite 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 receivedmessage, 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 browsersIn 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 columnIn the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queueWrite 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 162

An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events

What are the two considerations that must be kept in mind while using test recorder (Choose two answers)

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

Answer: AD

NEW QUESTION 167

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 172

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 177

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

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

Answer: A

Explanation:

V2 Rest API is recommended for on premise applications to access Object Store. It also comes with overhead of encryption and security of using rest api. With Object Store v2, the API call is localized to the same data center as the Runtime Manager app.
 But in this case requirement is to access the OS of other mule application and not the same mule application. 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.

NEW QUESTION 179

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 180

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

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

Answer: B

NEW QUESTION 184

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

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

Answer: B

Explanation:

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

NEW QUESTION 187

What is not true about Mule Domain Project?

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

Answer: C

Explanation:

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

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

What is Mule Domain Project?

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


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

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

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

You can start using domains in only three steps:

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



NEW QUESTION 191

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

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

Answer: B

Explanation:

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

NEW QUESTION 192

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 194

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 198

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 203

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 206

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

How are the messages consumed by the Mule application?

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

Answer: C

Explanation:

Correct answer is Depending on the Listener operation configuration, either all messages are consumed by ONLY the primary cluster node or else EACH message is consumed by ANY ONE cluster node

For applications running in clusters, you have to keep in mind the concept of primary node and how the connector will behave. When running in a cluster, the JMS listener default behavior will be to receive messages only in the primary node, no matter what kind of destination you are consuming from. In case of consuming messages from a Queue, you'll want to change this configuration to receive messages in all the nodes of the cluster, not just the primary.

This can be done with the primaryNodeOnly parameter:

```
<jms:listener config-ref="config" destination="{inputQueue}" primaryNodeOnly="false"/>
```

NEW QUESTION 210

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 212

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 the functionality

NEW QUESTION 217

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

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

Answer: D

Explanation:

Before we answer this question, we need to understand what median (50th percentile) and 80th percentile means. If the 50th percentile (median) of a response time is 500ms that means that 50% of my transactions are either as fast or faster than 500ms.

If the 90th percentile of the same transaction is at 1000ms it means that 90% are as fast or faster and only 10% are slower. Now as per upstream SLA, 99th percentile is 800 ms which means 99% of the incoming requests should have response time less than or equal to 800 ms. But as per one of the backend API, their 95th percentile is 1000 ms which means that backend API will take 1000 ms or less than that for 95% of requests. As there are three API invocation from upstream API, we can not conclude a timeout that can be set to meet the desired SLA as backend SLA's do not support it.

Let see why other answers are not correct.

1) Do not set a timeout --> This can potentially violate SLA's of upstream API

2) Set a timeout of 100 ms; ---> This will not work as backend API has 100 ms as median meaning only 50% requests will be answered in this time and we will get

timeout for 50% of the requests. Important thing to note here is, All APIs need to be executed sequentially, so if you get timeout in first API, there is no use of going to second and third API. As a service provider you wouldn't want to keep 50% of your consumers dissatisfied. So not the best option to go with.

*To quote an example: Let's assume you have built an API to update customer contact details.

- First API is fetching customer number based on login credentials
- Second API is fetching Info in 1 table and returning unique key
- Third API, using unique key provided in second API as primary key, updating remaining details

* Now consider, if API times out in first API and can't fetch customer number, in this case, it's useless to call API 2 and 3 and that is why question mentions specifically that all APIs need to be executed sequentially.

3) Set a timeout of 50 ms --> Again not possible due to the same reason as above Hence correct answer is No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

NEW QUESTION 220

.....

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