

Exam Questions MCI A-Level-1

MuleSoft Certified Integration Architect - Level 1

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



NEW QUESTION 1

An Integration Mule application is being designed to synchronize customer data between two systems. One system is an IBM Mainframe and the other system is a Salesforce Marketing Cloud (CRM) instance. Both systems have been deployed in their typical configurations, and are to be invoked using the native protocols provided by Salesforce and IBM.

What interface technologies are the most straightforward and appropriate to use in this Mule application to interact with these systems, assuming that Anypoint Connectors exist that implement these interface technologies?

- A. IBM: DB access CRM: gRPC
- B. IBM: REST CRM: REST
- C. IBM: Active MQ CRM: REST
- D. IBM: CICS CRM: SOAP

Answer: D

Explanation:

Correct answer is IBM: CICS CRM: SOAP

* Within Anypoint Exchange, MuleSoft offers the IBM CICS connector. Anypoint Connector for IBM CICS Transaction Gateway (IBM CTG Connector) provides integration with back-end CICS apps using the CICS Transaction Gateway.

* Anypoint Connector for Salesforce Marketing Cloud (Marketing Cloud Connector) enables you to connect to the Marketing Cloud API web services (now known as the Marketing Cloud API), which is also known as the Salesforce Marketing Cloud. This connector exposes convenient operations via SOAP for exploiting the capabilities of Salesforce Marketing Cloud.

NEW QUESTION 2

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

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

What must the architect do to achieve the caching objective?

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

Answer: A

NEW QUESTION 3

A company is modernizing its legal systems to accelerate access to applications and data while supporting the adoption of new technologies. The key to achieving this business goal is unlocking the companies' key systems and data including microservices running under Docker and Kubernetes containers using APIs.

Considering the current aggressive backlog and project delivery requirements the company wants to take a strategic approach in the first phase of its transformation projects by quickly deploying APIs in mule runtime that are able to scale, connect to on-premises systems and migrate as needed.

Which runtime deployment option supports company's goals?

- A. Customer hosted self-provisioned runtimes
- B. Cloudhub runtimes
- C. Runtime fabric on self-managed Kubernetes
- D. Runtime fabric on VMware metal

Answer: C

NEW QUESTION 4

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 sequentially. Any 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 parallel. Any 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 sequentially. Any 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 parallel. Any error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations.

Answer: A

NEW QUESTION 5

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

- A. User management
- B. Roles and permissions
- C. Dedicated load balancers
- D. Client Management

Answer: D

Explanation:

Correct answer is Client Management

- * Anypoint Platform acts as a client provider by default, but you can also configure external client providers to authorize client applications.
- * As an API owner, you can apply an OAuth 2.0 policy to authorize client applications that try to access your API. You need an OAuth 2.0 provider to use an OAuth 2.0 policy.
- * You can configure more than one client provider and associate the client providers with different environments. If you configure multiple client providers after you have already created environments, you can associate the new client providers with the environment.
- * You should review the existing client configuration before reassigning client providers to avoid any downtime with existing assets or APIs.
- * When you delete a client provider from your master organization, the client provider is no longer available in environments that used it.
- * Also, assets or APIs that used the client provider can no longer authorize users who want to access them.

-----MuleSoft
Reference: <https://docs.mulesoft.com/access-management/managing-api-clients>
<https://www.folkstalk.com/2019/11/mulesoft-integration-and-platform.html>

NEW QUESTION 6

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 7

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 8

A finance giant is planning to migrate all its Mule applications to Runtime fabric (RTF). Currently all Mule applications are deployed cloud hub using automated CI/CD scripts.


As an integration architect, which of the below step would you suggest to ensure that the applications from cloudhub are migrated properly to Runtime Fabric (RTF) with an assumption that organization is keen on keeping the same deployment strategy.

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

Answer: C

NEW QUESTION 9

Refer to the exhibit.



A Mule application is being designed to be deployed to several CloudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.

A watermark will be used to only retrieve those Satesforce Accounts that have been modified since the last time the integration logic ran.

What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?

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

Answer: B

Explanation:

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

* Mule provides two types of object stores:


- 1) In-memory store – stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime.
- 2) Persistent store – Mule persists data when an object store is explicitly configured to be persistent.

In a standalone Mule runtime, Mule creates a default persistent store in the file system. If you do not specify an object store, the default persistent object store is used.

MuleSoft Reference: <https://docs.mulesoft.com/mule-runtime/3.9/mule-object-stores>

NEW QUESTION 10

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

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

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

Answer: B

NEW QUESTION 11

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 16

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 18

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

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

Answer: A

NEW QUESTION 19

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

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

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

Answer: B

NEW QUESTION 20

A Mule application is being designed to receive nightly a CSV file containing millions of records from an external vendor over SFTP. The records from the file need to be validated, transformed, and then written to a database. Records can be inserted into the database in any order.

In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

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

Answer: C

Explanation:

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

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

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

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

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

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

NEW QUESTION 24

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

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

Answer: C

Explanation:

JMX Management Java Management Extensions (JMX) is a simple and standard way to manage applications, devices, services, and other resources. JMX is dynamic, so you can use it to monitor and manage resources as they are created, installed, and implemented. You can also use JMX to monitor and manage the

Java Virtual Machine (JVM). Each resource is instrumented by one or more Managed Beans, or MBeans. All MBeans are registered in an MBean Server. The JMX server agent consists of an MBean Server and a set of services for handling Mbeans. There are several agents provided with Mule for JMX support. The easiest way to configure JMX is to use the default JMX support agent. Log4J Agent The log4j agent exposes the configuration of the Log4J instance used by Mule for JMX management. You enable the Log4J agent using the <jmx-log4j> element. It does not take any additional properties MuleSoft Reference: <https://docs.mulesoft.com/mule-runtime/3.9/jmx-management>

NEW QUESTION 26

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

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

Answer: D

Explanation:

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

NEW QUESTION 31

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

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

Answer: D

Explanation:

* Mule High Availability Clustering provides basic failover capability for Mule. * When the primary Mule Runtime becomes unavailable, for example, because of a fatal JVM or hardware failure or it's taken offline for maintenance, a backup Mule Runtime immediately becomes the primary node and resumes processing where the failed instance left off. * After a system administrator recovers a failed Mule Runtime server and puts it back online, that server automatically becomes the backup node. In this case, Alice, once up, will become backup

-----Reference: <https://docs.mulesoft.com/mule-runtime/4.3/hadr-guide> So correct choice is : Files are polled from Bob node Bob is now the primary node

NEW QUESTION 33

An Organization has previously provisioned its own AWS VPC hosting various servers. The organization now needs to use Cloudhub to host a Mule application that will implement a REST API once deployed to Cloudhub, this Mule application must be able to communicate securely with the customer-provisioned AWS VPC resources within the same region, without being interceptable on the public internet. What Anypoint Platform features should be used to meet these network communication requirements between Cloudhub and the existing customer-provisioned AWS VPC?

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

Answer: A

Explanation:

Correct answer is: Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC
 * Connecting to your Anypoint VPC extends your corporate network and allows CloudHub workers to access resources behind your corporate firewall.
 * You can connect on-premises data centers through a secured VPN tunnel, or a private AWS VPC through VPC peering, or by using AWS Direct Connect.
 MuleSoft Doc Reference : <https://docs.mulesoft.com/runtime-manager/virtual-private-cloud>

NEW QUESTION 34

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 37

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 Am/point Exchange
- B. Compile, package, unit test, validate unit test coverage, deploy
- C. Compile, package, unit test, deploy, integration test

Answer: C

NEW QUESTION 38

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 42

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 47

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 49

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 and systems and support mule applications from change

Answer: AB

NEW QUESTION 50

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 53

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

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

Answer: C

NEW QUESTION 55

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

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

Answer: C

NEW QUESTION 59

An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments.
What must the testing team do before they can start integration testing the API in the Staging environment?

- A. They must access the API portal and create an API notebook using the Client ID and Client Secret supplied by the API portal in the Staging environment
- B. They must request access to the API instance in the Staging environment and obtain a Client ID and Client Secret to be used for testing the API
- C. They must be assigned as an API version owner of the API in the Staging environment
- D. They must request access to the Staging environment and obtain the Client ID and Client Secret for that environment to be used for testing the API

Answer: B

Explanation:

- * It's mentioned that the API is governed by a Client ID Enforcement policy in all environments.
- * Client ID Enforcement policy allows only authorized applications to access the deployed API implementation.
- * Each authorized application is configured with credentials: client_id and client_secret.
- * At runtime, authorized applications provide the credentials with each request to the API implementation. MuleSoft Reference: <https://docs.mulesoft.com/api-manager/2.x/policy-mule3-client-id-based-policies>

NEW QUESTION 63

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

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

Answer: BD

NEW QUESTION 66

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 67


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

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

Answer: A

NEW QUESTION 72

Refer to the exhibit.



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

- A. Create a mocking service that replicates the backend system's production performance characteristics.Then configure the API implementation to use the mocking service and conduct the performance tests
- B. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system

- C. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situation
- D. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
- E. Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale

Answer: A


Explanation:

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

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

NEW QUESTION 77

Refer to the exhibit.



A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations are Mule applications executing in a customer -hosted Mule runtime.


End-to-end correlation of all HTTP requests and responses belonging to each individual checkout Instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the same correlation ID for all requests and responses belonging to the same checkout instance.

What is the most efficient way (using the least amount of custom coding or configuration) for the web store backend and the implementations of the Experience API and Process API to participate in end-to-end correlation of the API invocations for each checkout instance?

A)

The web store backend, being a Java EE application, automatically makes use of the thread-local correlation ID generated by the Java EE application server and automatically transmits that to the Experience API using HTTP-standard headers


No special code or configuration is included in the web store backend, Experience API, and Process API implementations to generate and manage the correlation ID



B)

The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-It HTTP request header In each API invocation belonging to that checkout


No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID




C)

The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

The Experience API implementation must be coded to also propagate the correlation ID to the Process API in a suitable HTTP request header



D)
 The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API
 The Experience API and Process API implementations must be coded to receive the custom correlation ID In the HTTP requests and propagate It in suitable HTTP request headers



- A. Option A
- B. Option B
- C. Option C
- D. Option D


Answer: B

Explanation:

: 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 propogate 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.

Mulesoft Reference:

<https://help.mulesoft.com/s/article/How-to-Set-Custom-Correlation-Id-for-Flows-with-HTTP-Endpoint-in-Mule>
 Graphical user interface, application, Word Description automatically generated



NEW QUESTION 79

In Anypoint Platform, a company wants to configure multiple identity providers (IdPs) for multiple lines of business (LOBs). Multiple business groups, teams, and environments have been defined for these LOBs.

What Anypoint Platform feature can use multiple IdPs across the company's business groups, teams, and environments?

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

Answer: A

Explanation:

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

NEW QUESTION 81

An XA transaction is being configured that involves a JMS connector listening for incoming JMS messages. What is the meaning of the timeout attribute of the XA transaction, and what happens after the timeout expires?

- A. The time that is allowed to pass between committing the transaction and the completion of the Mule flow After the timeout, flow processing triggers an error
- B. The time that is allowed to pass between receiving JMS messages on the same JMS connection After the timeout, a new JMS connection is established
- C. The time that is allowed to pass without the transaction being ended explicitly After the timeout, the transaction is forcefully rolled-back
- D. The time that is allowed to pass for state JMS consumer threads to be destroyed After the timeout, a new JMS consumer thread is created

Answer: C

Explanation:

- * Setting a transaction timeout for the Bitronix transaction manager Set the transaction timeout either
 - In wrapper.conf
 - In CloudHub in the Properties tab of the Mule application deployment The default is 60 secs. It is defined as mule.bitronix.transactiontimeout = 120
 - * This property defines the timeout for each transaction created for this manager.
- If the transaction has not terminated before the timeout expires it will be automatically rolled back.

Additional Info around Transaction Management:


Bitronix is available as the XA transaction manager for Mule applications
 To use Bitronix, declare it as a global configuration element in the Mule application
`<bti:transaction-manager />`

Each Mule runtime can have only one instance of a Bitronix transaction manager, which is shared by all Mule applications
 For customer-hosted deployments, define the XA transaction manager in a Mule domain
 – Then share this global element among all Mule applications in the Mule runtime Graphical user interface, table Description automatically generated with medium confidence

Transaction Management		
Characteristics	Local Transactions	Extended Architecture (XA) Transactions
Connector Requisite 1	All operations inside the transaction must belong to same Connector.	Operations inside the transaction may belong to different Connectors.
Connector Requisite 2	Connectors may not be XA enabled	Connectors must be XA enabled
Connector Requisite 3	Connectors should use single config reference	Connectors may use multiple config references
Available resources	JMS, VM, JDBC	JMS, VM, JDBC
Uses Two Phase Commit (2PC)	No	Yes
DB Operations	Perform database operation to only one database resource	Perform database operation to one or more transactional resource
Supports Nested Transactions	Does not support nested transactions.	Supports nested transactions.
Bitronix is available	No	Yes
A.C.I.D Properties	No	Yes
Performance	Better than XA	Latency Increases
Thread Pooling	BLOCKING_IO	BLOCKING_IO
Recovery in case of system failure	No	Using Bitronix

NEW QUESTION 83

Refer to the exhibit.
 A Mule application is being designed to expose a SOAP web service to its clients.
 What language is typically used inside the web service's interface definition to define the data structures that the web service is expected to exchange with its clients?



- A. WSDL
- B. XSD
- C. JSON Schema
- D. RAML

Answer: B

Explanation:

Correct Answer XSD In this approach to developing a web service, you begin with an XML schema (XSD file) that defines XML data structures to be used as parameters and return types in the web service operations.

----- Reference:
https://www.w3schools.com/xml/schema_intro.asp

NEW QUESTION 84

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 88

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

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

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

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

Answer: B

NEW QUESTION 89

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

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

Answer: C

NEW QUESTION 94

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

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

Answer: D

Explanation:

We need to note few things about the scenario which will help us in reaching the correct solution.

Point 1 : The APIs are all publicly available and are associated with several mobile applications and web applications. This means Apply an IP blacklist policy is not viable option. as blacklisting IPs is limited to partial web traffic. It can't be useful for traffic from mobile application

Point 2 : The organization does NOT want to use any authentication or compliance policies for these APIs. This means we can not apply HTTPS mutual authentication scheme.

Header injection or removal will not help the purpose.

By its nature, JSON is vulnerable to JavaScript injection. When you parse the JSON object, the malicious code inflicts its damages. An inordinate increase in the size and depth of the JSON payload can indicate injection. Applying the JSON threat protection policy can limit the size of your JSON payload and thwart recursive additions to the JSON hierarchy.

Hence correct answer is Apply a JSON threat protection policy to all APIs to detect potential threat vectors

NEW QUESTION 96

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

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

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

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

Answer: D

NEW QUESTION 97

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

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

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

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

Answer: D

Explanation:

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

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

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

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

Object Store Stats API

NEW QUESTION 101

A Mule application contains a Batch Job scope with several Batch Step scopes. The Batch Job scope is configured with a batch block size of 25.


A payload with 4,000 records is received by the Batch Job scope.

When there are no errors, how does the Batch Job scope process records within and between the Batch Step scopes?

- A. The Batch Job scope processes multiple record blocks in parallel, and a block of 25 records can jump ahead to the next Batch Step scope over an earlier block of records Each Batch Step scope is invoked with one record in the payload of the received Mule event For each Batch Step scope, all 25 records within a block are processed in parallel All the records in a block must be completed before the block of 25 records is available to the next Batch Step scope
- B. The Batch Job scope processes each record block sequentially, one at a time Each Batch Step scope is invoked with one record in the payload of the received Mule event For each Batch Step scope, all 25 records within a block are processed sequentially, one at a time All 4000 records must be completed before the blocks of records are available to the next Batch Stepscope
- C. The Batch Job scope processes multiple record blocks in parallel, and a block of 25 records can jump ahead to the next Batch Step scope over an earlier block of records Each Batch Step scope is invoked with one record in the payload of the received Mule event For each Batch Step scope, all 25 records within a block are processed sequentially, one record at a time All the records in a block must be completed before the block of 25 records is available to the next Batch Step scope
- D. The Batch Job scope processes multiple record blocks in parallel Each Batch Step scope is invoked with a batch of 25 records in the payload of the received Mule event For each Batch Step scope, all 4000 records are processed in parallel Individual records can jump ahead to the next Batch Step scope before the rest of the records finish processing in the current Batch Step scope

Answer: A

NEW QUESTION 103
 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 108


What metrics about API invocations are available for visualization in custom charts using Anypoint Analytics?

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

Answer: C

Explanation:

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



NEW QUESTION 113

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 118

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

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

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

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

Answer: C

Explanation:

* On Error Propagate halts execution and sends error to the client. In this scenario it's mentioned that "processing of other rows must not be affected" so Option B and C are ruled out.

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

* Correct answer is For Each scope & On Error Continue scope Below requirement can be fulfilled in the below way

1) Using For Each scope , which will send each row from csv file sequentially. each row needs to be sent sequentially as requirement is to send the message in exactly the same way as it is mentioned in the csv file

2) Also other part of requirement is if any processing step for a row fails then it should log an error but should not affect other record processing . This can be achieved using On error Continue scope on these set of activities. so that error will not halt the processing. Also logger needs to be added in error handling section so that it can be logged.

* Attaching diagram for reference. Here it's try scope, but similar would be the case with For Each loop. Diagram Description automatically generated



NEW QUESTION 123

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 S. 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 nowhere 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 won't 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 124

An organization has deployed runtime fabric on an eight node cluster with performance profile. An API uses a non-persistent object store for maintaining some of its state data. What will be the impact to the state data if a 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 125

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

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

these interface technologies?

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

Answer: C

Explanation:

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

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

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

Considering the above points only correct option is

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

NEW QUESTION 129

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

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

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

- A. Deploy the Mule application to CloudhubCreate a CNAME record for base URL(<https://api.rutujar.com>) in the Cloudhub shared load balancer that points to the A record of the on-premises load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications
- B. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL (<https://api.rutujar.com>) in the organization's DNS server to point to the A record of the Cloudhub dedicated load balancerApply mapping rules in DLB to map URLto their corresponding Mule applications
- C. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL (<https://api.rutujar.com>) in the organization's DNS server to point to the A record of the CloudHub shared load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications
- D. For each migrated Mule application, deploy an API proxy application to Cloudhub with all traffic to themule applications routed through a Cloud Hub Dedicated load balancer (DLB)Update a CNAME record for base URL (<https://api.rutujar.com>) in the organization's DNS server to point to the A record of the CloudHub dedicated load balancerApply mapping rules in DLB to map each API proxy application who is responding new application

Answer: B

NEW QUESTION 134

A rate limiting policy has been applied to a soap VI.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 135

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 136

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 137

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 140


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 142

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 144

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 149

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-SAMLSecurityProfile.pdf>

NEW QUESTION 151

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 152

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 153

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 157


A Mule application contains a Batch Job with two Batch Steps (Batch_Step_1 and Batch_Step_2). A payload with 1000 records is received by the Batch Job. How many threads are used by the Batch Job to process records, and how does each Batch Step process records within the Batch Job?

- A. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and RECORDS are processed IN PARALLEL within and between the two Batch Steps
- B. Each Batch Job uses a SINGLE THREAD for all Batch steps Each Batch step instance receives ONE record at a time as the payload, and RECORDS are processed IN ORDER, first through Batch_Step_1 and then through Batch_Step_2
- C. Each Batch Job uses a SINGLE THREAD to process a configured block size of record Each Batch Step instance receives A BLOCK OF records as the payload, and BLOCKS of records are processed IN ORDER
- D. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and BATCH STEP INSTANCES execute IN PARALLEL to process records and Batch Steps in ANY order as fast as possible

Answer: A

Explanation:

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



NEW QUESTION 162

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 164

The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure. Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).

Given the company's current setup and requirements, what is the most idiomatic (used for its intended purpose) way to send Mule application logs to the external LMT?

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

Answer: A

NEW QUESTION 166

49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58. <https://api.acme.com/partnefs/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 apl.acme.com, signed by a reputable CA, to be used as the server certificate.

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

- A. Add the X.509 certificate to the Mule application's deployable archive, then configure a CloudHub Dedicated Load Balancer (DLB) for each of the Mule application's CloudHub 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 171

One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make.

Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments.

The production environment however does not have any rate limiting.

What is the cost-effective approach to conduct performance test of the API implementation in the non-production staging environment?

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

Answer: C

NEW QUESTION 174

A marketing organization is designing a Mule application to process campaign data. The Mule application will periodically check for a file in a SFTP location and process the records in the file. The size of the file can vary from 10MB to 5GB. Due to the limited availability of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.

The application must transform and send different formats of this file to three different downstream SFTP locations.

What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

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

Answer: A

NEW QUESTION 179

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

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

Answer: A

Explanation:

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

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

Explanation


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

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

Mule provides two types of object stores:

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

Diagram Description automatically generated



NEW QUESTION 182

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

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

Answer: A

Explanation:

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

NEW QUESTION 187

A travel company wants to publish a well-defined booking service API to be shared with its business partners. These business partners have agreed to ONLY consume SOAP services and they want to get the service contracts in an easily consumable way before they start any development. The travel company will publish the initial design documents to Anypoint Exchange, then share those documents with the business partners. When using an API-led approach, what is the first design document the travel company should deliver to its business partners?

- A. Create a WSDL specification using any XML editor
- B. Create a RAML API specification using any text editor
- C. Create an OAS API specification in Design Center
- D. Create a SOAP API specification in Design Center

Answer: A

Explanation:

SOAP API specifications are provided as WSDL. Design center doesn't provide the functionality to create WSDL file. Hence WSDL needs to be created using XML editor

NEW QUESTION 189

A Mule application is synchronizing customer data between two different database systems. What is the main benefit of using XA transaction over local transactions to synchronize these two database system?

- A. Reduce latency
- B. Increase throughput
- C. Simplifies communication
- D. Ensure consistency

Answer: D

Explanation:

- * XA transaction add tremendous latency so "Reduce Latency" is incorrect option XA transactions define "All or No" commit protocol.
- * Each local XA resource manager supports the A.C.I.D properties (Atomicity, Consistency, Isolation, and Durability).

So correct choice is "Ensure consistency"

NEW QUESTION 194

A company wants its users to log in to Anypoint Platform using the company's own internal user credentials. To achieve this, the company needs to integrate an external identity provider (IdP) with the company's

Anypoint Platform master organization, but SAML 2.0 CANNOT be used. Besides SAML 2.0, what single-sign-on standard can the company use to integrate the IdP with their Anypoint Platform master organization?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 198

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 202

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 207

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

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

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

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

Answer: B

NEW QUESTION 211

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 213

An organization is designing a mule application to support an all or nothing transaction between several database operations and some other connectors so that they all roll back if there is a problem with any of the connectors

Besides the database connector, what other connector can be used in the transaction.

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

Answer: A

Explanation:

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

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

	Shared Load Balancer	Dedicated Load Balancer
VPC	Shared VPC (Mulesoft)	VPC (Customer)
Default Load Balancer	Cloudhub provides 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 https://docs.mulesoft.com/runtime-manager/ib-whitelists	Yes
Configure Custom Domain	No	Yes
Custom Certificate	No	Yes
Rate Limit	Lower Rate Limit and applied According to Region	Higher Rate Limit Threshold
VPC	Anypoint VPC optional	Can't Use DLB without Anypoint VPC

NEW QUESTION 214

A global, high-volume shopping Mule application is being built and will be deployed to CloudHub. To improve performance, the Mule application uses a Cache scope that maintains cache state in a CloudHub object store. Web clients will access the Mule application over HTTP from all around the world, with peak volume coinciding with business hours in the web client's geographic location. To achieve optimal performance, what Anypoint Platform region should be chosen for the CloudHub object store?

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

Answer: A

Explanation:

CloudHub object store should be in same region where the Mule application is deployed. This will give optimal performance.

Before learning about Cache scope and object store in Mule 4 we understand what is in general Caching is and other related things.

WHAT DOES "CACHING" MEAN?

Caching is the process of storing frequently used data in memory, file system or database which saves processing time and load if it would have to be accessed from original source location every time.

In computing, a cache is a high-speed data storage layer which stores a subset of data, so that future requests for that data are served up faster than is possible by accessing the data's primary storage location. Caching allows you to efficiently reuse previously retrieved or computed data.

How does Caching work?

The data in a cache is generally stored in fast access hardware such as RAM (Random-access memory) and may also be used in correlation with a software component. A cache's primary purpose is to increase data retrieval performance by reducing the need to access the underlying slower storage layer.

Caching in MULE 4

In Mule 4 caching can be achieved in mule using cache scope and/or object-store. Cache scope internally uses Object Store to store the data.

What is Object Store

Object Store lets applications store data and states across batch processes, Mule components, and applications, from within an application. If used on cloud hub, the object store is shared between applications deployed on Cluster.

Cache Scope is used in below-mentioned cases:

Need to store the whole response from the outbound processor

Data returned from the outbound processor does not change very frequently

As Cache scope internally handle the cache hit and cache miss scenarios it is more readable Object Store is used in below-mentioned cases:

Need to store custom/intermediary data To store watermarks

Sharing the data/stage across applications, schedulers, batch.

If CloudHub object store is in same region where the Mule application is deployed it will aid in fast access of data and give optimal performance.

NEW QUESTION 216

An organization has strict unit test requirement that mandate every mule application must have an MUnit test suit with a test case defined for each flow and a minimum test coverage of 80%.

A developer is building Munit test suit for a newly developed mule application that sends API request to an external rest API.

What is the effective approach for successfully executing the Munit tests of this new application while still achieving the required test coverage for the Munit tests?

- A. Invoke the external endpoint of the rest API from the mule floors
- B. Mark the rest API invocations in the Munits and then call the mocking service flow that simulates standard responses from the REST API
- C. Mock the rest API invocation in the Munits and return a mock response for those invocations
- D. Create a mocking service flow to simulate standard responses from the rest API and then configure the mule flows to call the marking service flow

Answer: C

NEW QUESTION 220

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

How are the messages consumed by the Mule application?

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

Answer: C

Explanation:

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

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

This can be done with the primaryNodeOnly parameter:

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

NEW QUESTION 223

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 227

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

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

Answer: C

Explanation:

The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation

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

NEW QUESTION 228

An API implementation is being developed to expose data from a production database via HTTP requests. The API implementation executes a database SELECT statement that is dynamically created based upon data received from each incoming HTTP request. The developers are planning to use various types of testing to make sure the Mule application works as expected, can handle specific workloads, and behaves correctly from an API consumer perspective. What type of testing would typically mock the results from each SELECT statement rather than actually execute it in the production database?

- A. Unit testing (white box)

- B. Integration testing
- C. Functional testing (black box)
- D. Performance testing

Answer: A

Explanation:

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

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

Below are the typical characteristics of unit testing.

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

NEW QUESTION 232

.....

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