

MCIA-Level-1 Dumps

MuleSoft Certified Integration Architect - Level 1

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



NEW QUESTION 1

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 2

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 3

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

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

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

Answer: A

NEW QUESTION 4

Additional nodes are being added to an existing customer-hosted Mule runtime cluster to improve performance. Mule applications deployed to this cluster are invoked by API clients through a load balancer.

What is also required to carry out this change?

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

Answer: B

Explanation:

* Clustering is a group of servers or mule runtime which acts as a single unit.

* Mulesoft Enterprise Edition supports scalable clustering to provide high availability for the Mulesoft application.

* In simple terms, virtual servers composed of multiple nodes and they communicate and share information through a distributed shared memory grid.

* By default, Mulesoft ensures the High availability of applications if clustering implemented.

* Let's consider the scenario one of the nodes in cluster crashed or goes down and under maintenance. In such cases, Mulesoft will ensure that requests are processed by other nodes in the cluster. Mulesoft clustering also ensures that the request is load balanced between all the nodes in a cluster.

* Clustering is only supported by on-premise Mule runtime and it is not supported in Cloudhub.

Correct answer is External monitoring tools or log aggregators must be configured to recognize the new nodes

* Rest of the options are automatically taken care of when a new node is added in cluster.

NEW QUESTION 5

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

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

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

- A. A MuleSoft-managed repository at repository.mulesoft.org

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

Answer: C

NEW QUESTION 6

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

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

Answer: A

NEW QUESTION 7

An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to custom hosted Mule runtimes (on-premises) while others execute in the MuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems.

How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access the backend systems?

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

Answer: D

Explanation:

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

NEW QUESTION 8

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

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

Answer: D

Explanation:

Correct answer is Client Management

* Anypoint Platform acts as a client provider by default, but you can also configure external client providers to authorize client applications.

* As an API owner, you can apply an OAuth 2.0 policy to authorize client applications that try to access your API. You need an OAuth 2.0 provider to use an OAuth 2.0 policy.

* You can configure more than one client provider and associate the client providers with different environments. If you configure multiple client providers after you have already created environments, you can associate the new client providers with the environment.

* You should review the existing client configuration before reassigning client providers to avoid any downtime with existing assets or APIs.

* When you delete a client provider from your master organization, the client provider is no longer available in environments that used it.

* Also, assets or APIs that used the client provider can no longer authorize users who want to access them.

-----MuleSoft

Reference: <https://docs.mulesoft.com/access-management/managing-api-clients>

<https://www.folkstalk.com/2019/11/mulesoft-integration-and-platform.html>

NEW QUESTION 9

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 10

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

The organization does not currently use AWS in any way.

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 10

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 15

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 18


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 19

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

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

Answer: D

NEW QUESTION 21

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 25

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

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

Answer: B

Explanation:

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

NEW QUESTION 30

Organization wants to achieve high availability goal for Mule applications in customer hosted runtime plane. Due to the complexity involved, data cannot be shared

among of different instances of same Mule application. What option best suits to this requirement considering high availability is very much critical to the organization?

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

Answer: B

Explanation:

High availability is about up-time of your application

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

NEW QUESTION 31

A Mule application is deployed to a cluster of two(2) cusomter-hosted Mule runtimes. Currently the node name Alice is the primary node and node named bob is the secondary node. The mule application has a flow that polls a directory on a file system for new files.

The primary node Alice fails for an hour and then restarted.

After the Alice node completely restarts, from what node are the files polled, and what node is now the primary node for the cluster?

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

Answer: D


Explanation:

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

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

NEW QUESTION 32

Refer to the exhibit.



An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a server-side certificate that rotates often.

What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

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

Answer: B

Explanation:

In this scenario, both A & C will work, but A is better as it does not require repackage to the domain project at all.

Correct answer is Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint. Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION.

What is Mule Domain Project?

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


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

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

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

You can start using domains in only three steps:

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



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

NEW QUESTION 37

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

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

Answer: A

Explanation:

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

NEW QUESTION 40

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 43

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 48

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 50

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 52

What is maximum vCores can be allocated to application deployed to CloudHub?

- A. 1 vCores
- B. 2 vCores
- C. 4 vCores
- D. 16 vCores

Answer: D

NEW QUESTION 54

A mule application designed to fulfil two requirements

a) Processing files are synchronously from an FTPS server to a back-end database using VM intermediary queues for load balancing VM events

b) Processing a medium rate of records from a source to a target system using batch job scope

Considering the processing reliability requirements for FTPS files, how should VM queues be configured for processing files as well as for the batch job scope if


the application is deployed to Cloudhub workers?

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

Answer: C

NEW QUESTION 55

Refer to the exhibit.



A Mule application has an HTTP Listener that accepts HTTP DELETE requests. This Mule application is deployed to three CloudHub workers under the control of the CloudHub Shared Load Balancer.

A web client makes a sequence of requests to the Mule application's public URL.

How is this sequence of web client requests distributed among the HTTP Listeners running in the three CloudHub workers?

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

Answer: C

Explanation:

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

NEW QUESTION 58

An insurance company is using a CloudHub runtime plane. As a part of requirement, email alert should be sent to internal operations team every time of policy applied to an API instance is deleted As an integration architect suggest on how this requirement be met?

- A. Use audit logs in Anypoint platform to detect a policy deletion and configure the Audit logs alert feature to send an email to the operations team
- B. Use Anypoint monitoring to configure an alert that sends an email to the operations team every time a policy is deleted in API manager
- C. Create a custom connector to be triggered every time of policy is deleted in API manager
- D. Implement a new application that uses the Audit log REST API to detect the policy deletion and send an email to operations team the SMTP connector

Answer: D

NEW QUESTION 60

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 62

An organization is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to rejections from the back-end system will need to be processed manually (outside the back-end system). The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located


inside the organization's firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages. What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

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

Answer: A

NEW QUESTION 64

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 68

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



Experience API generates correlationID and returns to backend

B)


The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id 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



correlationID generated by backend

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



Thread-local correlationID sent in HTTP 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



correlationID generated by backend

Modify to propogate custom correlationIDs

- 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 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.

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 72

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



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

Answer: C

Explanation:


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

Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize APIs as per the requirements of different applications. It is not a recommended approach.

Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.

Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system

A screenshot of a computer Description automatically generated with low confidence




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

Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.

Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.

Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system

Diagram Description automatically generated



NEW QUESTION 74

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 77

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 78

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 82

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 87

An organization has defined a common object model in Java to mediate the communication between different Mule applications in a consistent way. A Mule application is being built to use this common object model to process responses from a SOAP API and a REST API and then write the processed results to an order management system.

The developers want Anypoint Studio to utilize these common objects to assist in creating mappings for various transformation steps in the Mule application.

What is the most idiomatic (used for its intended purpose) and performant way to utilize these common objects to map between the inbound and outbound systems in the Mule application?

- A. Use JAXB (XML) and Jackson (JSON) data bindings
- B. Use the WSS module
- C. Use the Java module
- D. Use the Transform Message component

Answer: A

NEW QUESTION 88

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 92

An integration Mute application consumes and processes a list of rows from a CSV file. Each row must be read from the CSV file, validated, and the row data sent

to a JMS queue, in the exact order as in the CSV file.

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

What combination of 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 93

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

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

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

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

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

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

Answer: D

Explanation:

: Using Anypoint MQ, you can create two types of queues: Standard queue These queues don't guarantee a specific message order. Standard queues are the best fit for applications in which messages must be delivered quickly. FIFO (first in, first out) queue These queues ensure that your messages arrive in order. FIFO queues are the best fit for applications requiring strict message ordering and exactly-once delivery, but in which message

delivery speed is of less importance. Use of FIFO queue is nowhere in the option and also it decreased throughput. Similarly persistent object store is not the preferred solution approach when you are 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 to 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 REQUEST.

NEW QUESTION 98

In one of the critical payment related mule application, transaction is being used. As an enhancement to implementation, scatter gather route is introduced which is also the part of transaction group. Scatter gather route has 4 routes.

What will be the behavior of the Mule application in case of error occurs in 4th route of the scatter-gather router and transaction needs to be rolled back?

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

Answer: B

Explanation:

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

NEW QUESTION 102

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

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

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

Answer: B

NEW QUESTION 105

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 node is unaffected by the crash

Answer: D

NEW QUESTION 107

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

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

- A. Keep the default logging configuration in RuntimeManager. Define the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support all Mule application deployments.
- B. Disable CloudHub logging in Runtime Manager. Define the splunk appender in EACH Mule application's log4j2.xml file.
- C. Disable CloudHub logging in Runtime Manager. Define the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support all Mule application deployments.
- D. Keep the default logging configuration in Runtime Manager. Define the Splunk appender in EACH Mule application log4j2.xml file.

Answer: B

Explanation:

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

NEW QUESTION 112

A rate limiting policy has been applied to a SOAP V1.2 API published in CloudHub. The API implementation catches errors in a global error handler on error propagation in the main flow for HTTP: RETRY_EXHAUSTED with HTTP status set to 429 and any with the HTTP status set to 500.

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

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

Answer: A

NEW QUESTION 116

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 119

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 123

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 124

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 129

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 130

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 131

A manufacturing company is planning to deploy Mule applications to its own Azure Kubernetes Service infrastructure.

The organization wants to make the Mule applications more available and robust by deploying each Mule application to an isolated Mule runtime in a Docker container while managing all the Mule applications from the MuleSoft-hosted control plane.

What is the most idiomatic (used for its intended purpose) choice of runtime plane to meet these organizational requirements?

A. Anypoint Platform Private Cloud Edition

B. Anypoint Runtime Fabric


C. CloudHub

D. Anypoint Service Mesh

Answer: B

NEW QUESTION 132

Refer to the exhibit.



Error handling

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

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

Answer: C

NEW QUESTION 137

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

How should the JMS Listener be configured to maximize performance and concurrent message processing of the JMS queue?

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

Answer: D

NEW QUESTION 141

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

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

Answer: C

Explanation:

- * Anypoint Runtime Fabric is a container service that automates the deployment and orchestration of your Mule applications and gateways.
- * Runtime Fabric runs on customer-managed infrastructure on AWS, Azure, virtual machines (VMs) or bare-metal servers.
- * As none of the Mule applications use Mule domain projects. applications are not required to be rewritten. Also when applications are deployed on RTF, by default ingress is allowed only on 8081.
- * Hence port conflicts are not required to be managed by DevOps team

NEW QUESTION 144

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

- A. An On Error scope Non-persistent VM ActiveMQ Dead Letter Queue for manual processing

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

Answer: D

Explanation:

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

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

NEW QUESTION 149

One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make. Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments. The production environment however does not have any rate limiting.

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

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

Answer: C

NEW QUESTION 152

An organization has just developed a Mule application that implements a REST API. The mule application will be deployed to a cluster of customer hosted Mule runtimes.

What additional infrastructure component must the customer provide in order to distribute inbound API requests across the Mule runtimes of the cluster?

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

Answer: B

Explanation:

Correct answer is An HTTP Load Balancer.

Key thing to note here is that we are deploying application to customer hosted Mule runtime. This means we will need load balancer to route the requests to different instances of the cluster.

NEW QUESTION 153

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 availabilty 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 157

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

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

- A. Tests for flows cannot be created with Mule errors raised inside the flow or already existing in the incoming event
- B. Recorder supports smoking a message before or inside a ForEach processor
- C. The recorder support loops where the structure of the data been tested changes inside the iteration

- D. A recorded flow execution ends successfully but the result does not reach its destination because the application is killed
- E. Mocking values resulting from parallel processes are possible and will not affect the execution of the processes that follow in the test

Answer: AD

NEW QUESTION 162

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 167

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 169

An insurance company has an existing API which is currently used by customers. API is deployed to customer hosted Mule runtime cluster. The load balancer that is used to access any APIs on the mule cluster is only configured to point to applications hosted on the server at port 443.

Mule application team of a company attempted to deploy a second API using port 443 but the application will not start and checking logs shows an error indicating the address is already in use.

Which steps must the organization take to resolve this error and allow customers to access both the API's?

- A. Change the base path of the HTTP listener configuration in the second API to a different one from the first API
- B. Set HTTP listener configuration in both API's to allow for connections from multiple ports
- C. Move the HTTP listener configurations from the API's and package them in a mule domain project using port 443
- D. Set the HTTP listener of the second API to use different port than the one used in the first API

Answer: C

NEW QUESTION 170

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL <https://api.acme.com>

As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub

What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

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

Answer: A

Explanation:

<https://help.mulesoft.com/s/feed/0D52T000055pzgsSAA>.

NEW QUESTION 175

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 179

An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger of flow for processing a high-volume set of records from the source system and synchronize with the SaaS system using the Batch job scope

After processing 1000 records in a periodic synchronization of 1 lakh records, the replicas in which batch job instance was started went down due to unexpected failure in the runtime fabric environment

What is the consequence of losing the replicas that run the Batch job instance?

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

Answer: B

NEW QUESTION 183

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 188

An insurance provider is implementing Anypoint platform to manage its application infrastructure and is using the customer hosted runtime for its business due to certain financial requirements it must meet. It has built a number of synchronous API's and is currently hosting these on a mule runtime on one server. These applications make use of a number of components including heavy use of object stores and VM queues. Business has grown rapidly in the last year and the insurance provider is starting to receive reports of reliability issues from its applications.

The DevOps team indicates that the API's are currently handling too many requests and this is over loading the server. The team has also mentioned that there is a significant downtime when the server is down for maintenance.

As an integration architect, which option would you suggest to mitigate these issues?

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

Answer: B

NEW QUESTION 190

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 191

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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