MuleSoft.MCIA-Level-1-Maintenance.v2023-06-17.q42

Exam Code:	MCIA-Level-1-Maintenance
Exam Name:	MuleSoft Certified Integration Architect - Level 1 MAINTENANCE
Certification Provider:	MuleSoft
Free Question Number:	42
Version:	v2023-06-17
# of views:	1042
# of Questions views:	420
https://www.freeqas.com/qa/MuleSoft/MCIA-Level-1-Maintenance/MuleSoft.MCIA-Level-1-Maintenance.v2023-06-17.q42.html	

NEW QUESTION: 1
What condition requires using a CloudHub Dedicated Load Balancer?
A. When cross-region load balancing is required between separate deployments of the same Mule application
B. When custom DNS names are required for API implementations deployed to customer-hosted Mule runtimes
C. When API invocations across multiple CloudHub workers must be load balanced
D. When server-side load-balanced TLS mutual authentication is required between API implementations and API clients
Answer: D (LEAVE A REPLY)
Explanation
Correct answer is When server-side load-balanced TLS mutual authentication is required between API implementations and API clients CloudHub dedicated load balancers (DLBs) are an optional component of Anypoint Platform that enable you to route external HTTP and HTTPS traffic to multiple Mule applications deployed to CloudHub workers in a Virtual Private Cloud (VPC). Dedicated load balancers enable you to: * Handle load balancing among the different CloudHub workers that run your application. * Define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication. * Configure proxy rules that map your applications to custom domains. This enables you to host your applications under a single domain

NEW QUESTION: 2
As an enterprise architect, what are the two reasons for which you would use a canonical data model in the new integration project using Mulesoft Anypoint platform (choose two answers)
A. There are multiple canonical definitions of each data type
B. Because the model isolates the back and systems and support mule applications from change

C. To isolate areas within a bounded context
D. To incorporate industry standard data formats
E. To have consistent data structure aligned in processes
Answer: C,E (LEAVE A REPLY)

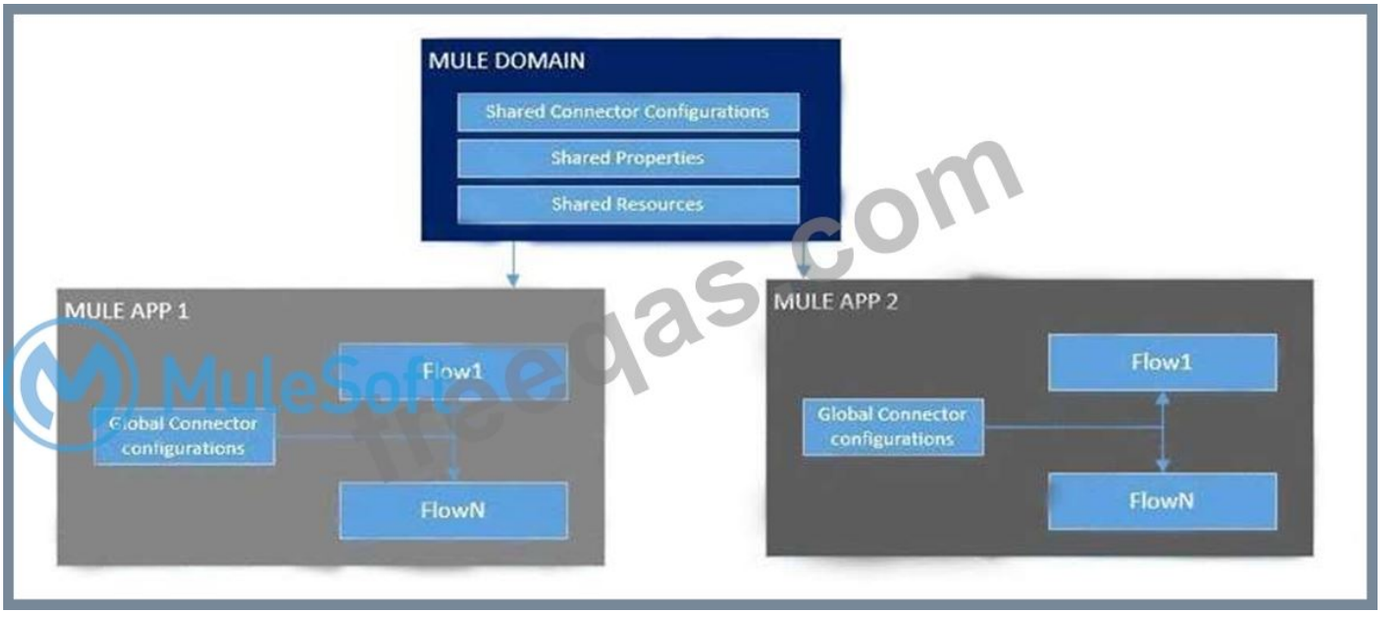
NEW QUESTION: 3
A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should 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: (SHOW ANSWER)
Explanation
Set the Anypoint MQ connector operation to publish or consume messages, or to accept (ACK) or not accept (NACK) a message.

NEW QUESTION: 4
A finance giant is planning to migrate all its Mule applications to Runtime fabric (RTF). Currently all Mule applications are deployed cloud hub using automated CI/CD scripts.
As an integration architect, which of the below step would you suggest to ensure that the applications from cloudhub are migrated properly to Runtime Fabric (RTF) with an assumption that organization is keen on keeping the same deployment strategy.
A. runtimeFabric dependency should be added as a mule plug-in to POM.xml file and CI/CD script should be modified as per the RTF configurations
B. No changes need to be made to POM.xml file and CI/CD script should be modified as per the RTF configurations
C. runtimeFabric deployment should be added to POM.xml file in all
the mule applications and CI/CD script should be modified as per the RTF configurations
D. runtimeFabric profile should be added mule configuration files in the mule applications and CI/CD script should be modified as per the RTF configurations
Answer: C (LEAVE A REPLY)

NEW QUESTION: 5
A mule application is required to periodically process large data set from a back-end database to Salesforce CRM using batch job scope configured properly process the higher rate of records.
The application is deployed to two cloudhub workers with no persistence queues enabled.
What is the consequence if the worker crashes during records processing?
A. Remaining records be processed by second worker
B. Remaining records will be processed by a new replacement worker
C. All the records will be processed from scratch by the second worker leading to duplicate processing
D. Remaining records will be left and processed
Answer: D (LEAVE A REPLY)

NEW QUESTION: 6
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 (LEAVE A REPLY)
Explanation
* Mule Domain Project is ONLY available for customer-hosted Mule runtimes, but not for Anypoint Runtime Fabric
* Mule domain project is available for Hybrid and Private Cloud (PCE). Rest all provide application isolation and can't support domain project.
What is Mule Domain Project?
* A Mule Domain Project is implemented to configure the resources that are shared among different projects.
These resources can be used by all the projects associated with this domain. Mule applications can be associated with only one domain, but a domain can be associated with multiple projects. Shared resources allow multiple development teams to work in parallel using the same set of reusable connectors. Defining these connectors as shared resources at the domain level allows the team to: - Expose multiple services within the domain through the same port. - Share the connection to persistent storage. - Share services between apps through a well-defined interface. - Ensure consistency between apps upon any changes because the configuration is only set in one place.
* Use domains Project to share the same host and port among multiple projects. You can declare the http connector within a domain project and associate the domain project with other projects. Doing this also allows to control thread settings, keystore configurations, time outs for all the requests made within multiple applications. You may think that one can also achieve this by duplicating the http connector configuration across all the applications. But, doing this may pose a nightmare if you have to make a change and redeploy all the applications.

* If you use connector configuration in the domain and let all the applications use the new domain instead of a default domain, you will maintain only one copy of the http connector configuration. Any changes will require only the domain to the redeployed instead of all the applications.
You can start using domains in only three steps:
1) Create a Mule Domain project
2) Create the global connector configurations which needs to be shared across the applications inside the Mule Domain project
3) Modify the value of domain in mule-deploy.properties file of the applications Graphical user interface Description automatically generated



NEW QUESTION: 7
An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/'CD pipeline.
Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?
A. Use Anypoint CLI or Anypoint Platform REST apis with scripting language such as groovy
B. Use runtime manager rest apis for API management and mavenforAPI deployment
C. Use Exchange rest api's for API management and MavenforAPI deployment
D. Use Maven with a custom configuration required for the API lifecycle
Answer: C (LEAVE A REPLY)

NEW QUESTION: 8
What API policy would LEAST likely be applied to a Process API?
A. Custom circuit breaker
B. Client ID enforcement
C. Rate limiting

D. JSON threat protection

Answer: D (LEAVE A REPLY)

Explanation

Key to this question lies in the fact that Process API are not meant to be accessed directly by clients. Lets analyze options one by one. Client ID enforcement : This is applied at process API level generally to ensure that identity of API clients is always known and available for API-based analytics Rate Limiting : This policy is applied on Process Level API to secure API's against degradation of service that can happen in case load received is more than it can handle Custom circuit breaker : This is also quite useful feature on process level API's as it saves the API client the wasted time and effort of invoking a failing API. JSON threat protection :

This policy is not required at Process API and rather implemented as Experience API's. This policy is used to safeguard application from malicious attacks by injecting malicious code in JSON object. As ideally Process API's are never called from external world , this policy is never used on Process API's Hence correct answer is JSON threat protection MuleSoft Documentation Reference :

https://docs.mulesoft.com/api-manager/2.x/policy-mule3-json-threat

NEW QUESTION: 9

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

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

A. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back

B. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations

C. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations

D. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back

Answer: (SHOW ANSWER)

NEW QUESTION: 10

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

A. the message broker must be hosted on premises

B. support for point-to-point messaging

C. ability for a third party outside the company's network to consume events from the queue

D. payloads must be encrypted

E. single message payload can be up to 15 MB

Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 11
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 (LEAVE A REPLY)
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: 12
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 (LEAVE A REPLY)
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: 13
Organization wants to achieve high availability goal for Mule applications in customer hosted runtime plane.
Due to the complexity involved, data cannot be shared among of different instances of same Mule application.

What option best suits to this requirement considering high availability is very much critical to the organization?

A. The cluster can be configured

B. Use third party product to implement load balancer

C. High availability can be achieved only in CloudHub

D. Use persistent object store

Answer: B (LEAVE A REPLY)

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: 14

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 the Java module

B. Use the WSS module

C. Use JAXB (XML) and Jackson (JSON) data bindings

D. Use the Transform Message component

Answer: C (LEAVE A REPLY)

NEW QUESTION: 15

The AnyAirline organization's passenger reservations center is designing an integration solution that combines invocations of three different System APIs (bookFlight, bookHotel, and bookCar) in a business transaction.

Each System API makes calls to a single database.

The entire business transaction must be rolled back when at least one of the APIs fails.

What is the most idiomatic (used for its intended purpose) way to integrate these APIs in near real-time that provides the best balance of consistency, performance, and reliability?

A. Implement local transactions in each API implementation

Coordinate between the API implementations using a Saga pattern

Apply various compensating actions depending on where a failure occurs

B. Implement eXtended Architecture (XA) transactions between the API implementations Coordinate between the API implementations using a Saga pattern Implement caching in each API implementation to improve performance
C. Implement local transactions within each API implementation
Configure each API implementation to also participate in the same eXtended Architecture (XA) transaction Implement caching in each API implementation to improve performance
D. Implement an eXtended Architecture (XA) transaction manager in a Mule application using a Saga pattern Connect each API implementation with the Mule application using XA transactions Apply various compensating actions depending on where a failure occurs
Answer: A (LEAVE A REPLY)

NEW QUESTION: 16
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. The recorder support loops where the structure of the data been tested changes inside the iteration
B. A recorded flow execution ends successfully but the result does
not reach its destination because the application is killed
C. Recorder supports smoking a message before or inside a ForEach processor
D. Tests for flows cannot be created with Mule errors raised inside
the flow or already existing in the incoming event
E. Mocking values resulting from parallel processes are possible and will not affect the execution of the processes that follow in the test
Answer: (SHOW ANSWER)

Valid MCIA-Level-1-Maintenance Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1-Maintenance Exam! PrepPdf.com now offer the **newest MCIA-Level-1-Maintenance exam dumps**, the PrepPdf.com MCIA-Level-1-Maintenance exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1-Maintenance dumps with Test Engine here: https://www.preppdf.com/MuleSoft/MCIA-Level-1-Maintenance-prepaway-exam-dumps.html (**118** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 17
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. Configure Persistent Gateway at the RTF
C. Install the object store pod on one of the cluster nodes
D. Configure Persistence Gateway in any of the servers using Mule Object Store
Answer: (SHOW ANSWER)

NEW QUESTION: 18
A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction.
What is the behavior of the Scatter-Gather when running within this local transaction?
A. Execution of each route within the Scatter-Gather occurs sequentially Any error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations
B. Execution of each route within the Scatter-Gather occurs sequentially Any error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
C. Execution of all routes within the Scatter-Gather occurs in parallel Any error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
D. Execution of each route within the Scatter-Gather occurs in parallel Any error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations
Answer: (SHOW ANSWER)

NEW QUESTION: 19
An organization if struggling frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects.
How can these best practices be achieved with the applications having the least amount of responsibility?
A. Create a mule domain project with all the dependencies define in its POM.xml file and add each application to the domain Project
B. Create a parent POM of all the required dependencies and reference each in each application's POM.xml file
C. Create a Mule plugin project with all the dependencies and add it as a dependency in each application's POM.xml file
D. Add all dependencies in each application's POM.xml file
Answer: B (LEAVE A REPLY)

NEW QUESTION: 20
An ABC Farms project team is planning to build a new API that is required to work with data from different domains across the organization.

The organization has a policy that all project teams should leverage existing investments by reusing existing APIs and related resources and documentation that other project teams have already developed and deployed.
To support reuse, where on Anypoint Platform should the project team go to discover and read existing APIs, discover related resources and documentation, and interact with mocked versions of those APIs?
A. Design Center
B. API Manager
C. Runtime Manager
D. Anypoint Exchange
Answer: D (LEAVE A REPLY)
Explanation
The mocking service is a feature of Anypoint Platform and runs continuously. You can run the mocking service from the text editor, the visual editor, and from Anypoint Exchange. You can simulate calls to the API in API Designer before publishing the API specification to Exchange or in Exchange after publishing the API specification.

NEW QUESTION: 21
As a part of project , existing java implementation is being migrated to Mulesoft. Business is very tight on the budget and wish to complete the project in most economical way possible.
Canonical object model using java is already a part of existing implementation. Same object model is required by mule application for a business use case. What is the best way to achieve this?
A. Make use of Java module
B. Create similar model for Mule applications
C. Create a custom application to read Java code and make it available for Mule application
D. Use Anypoint exchange
Answer: A (LEAVE A REPLY)
Explanation
Mule 4 is built to:
*Minimize the need for custom code.
*Avoid the need for you to know or understand Java.
However, some advanced uses cases require integration with custom Java code, such as:
*Reuse of a library, such as a tax calculation library.
*Reuse of a canonical object model that is standard in the organization.
*Execution of custom logic using Java.
Mule ref doc :https://docs.mulesoft.com/java-module/1.2/

NEW QUESTION: 22
What is maximum vCores can be allocated to application deployed to CloudHub?
A. 2 vCores

B. 4 vCores
C. 16 vCores
D. 1 vCores
Answer: C (LEAVE A REPLY)

NEW QUESTION: 23
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. 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
B. An Until Successful scope to call the back-end system
One or more ActiveMQ long-retry queues
One or more ActiveMQ dead-letter queues for manual processing
C. One or more On Error scopes to assist calling the back-end system
An Until Successful scope containing VM components for long retries
A persistent dead-letter VM queue configured in CloudHub
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: (SHOW ANSWER)

NEW QUESTION: 24
A mule application must periodically process a large dataset which varies from 6 GB lo 8 GB from a back-end database and write transform data lo an FTPS server using a properly configured bad job scope.
The performance requirements of an application are approved to run in the cloud hub 0.2 vCore with 8 GB storage capacity and currency requirements are met.
How can the high rate of records be effectively managed in this application?

A. Use streaming with an in-memory reputable store strategy for reading records from the database and batch aggregator with streaming to write to FTPS
B. Use streaming with a file store repeatable strategy reading records from the database and batch aggregator without any required configuration
C. Use streaming with a file store repeatable strategy for reading records from the database and batch aggregator with an optimal size
D. Use streaming with a file storage repeatable strategy for reading records from the database and batch aggregator with streaming to write to FTPS
Answer: D (LEAVE A REPLY)

NEW QUESTION: 25
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 (LEAVE A REPLY)
Explanation
The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation This is based on below definitions API client * An application component * that accesses a service * by invoking an API of that service - by definition of the term API over HTTP API consumer * A business role, which is often assigned to an individual * that develops API clients, i.e., performs the activities necessary for enabling an API client to invoke APIs API implementation * An application component * that implements the functionality

NEW QUESTION: 26
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 (LEAVE A REPLY)

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: 27

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. The Anypoint Object Store service at cloudhub.io

B. A MuleSoft-managed repository atrepository.mulesoft.org

C. A binary artifact repository

D. API Community Manager

Answer: (SHOW ANSWER)

NEW QUESTION: 28

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

A. Mule event ID

B. Mule correlation ID

C. Client's IP address

D. DataWeave UUID

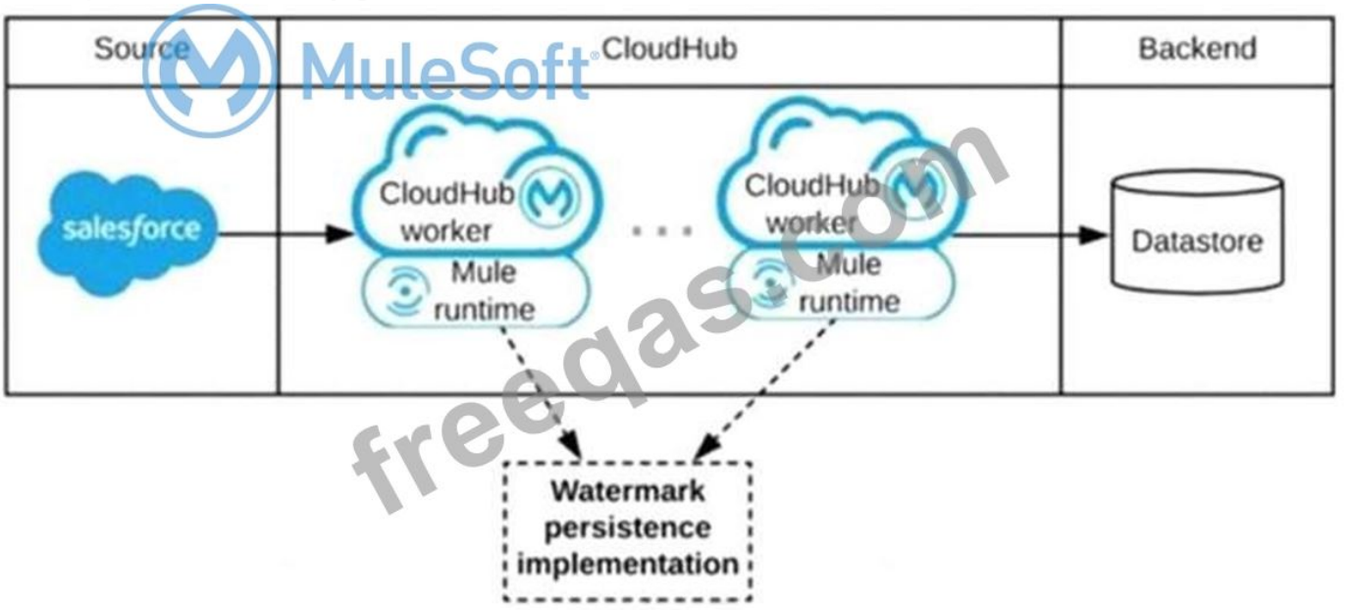
Answer: B (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 29
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: (SHOW ANSWER)
Explanation
Correct answer is "Compile, package, unit test, validate unit test coverage, deploy" Explanation : 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 CICD pipeline which can ne 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 cant 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

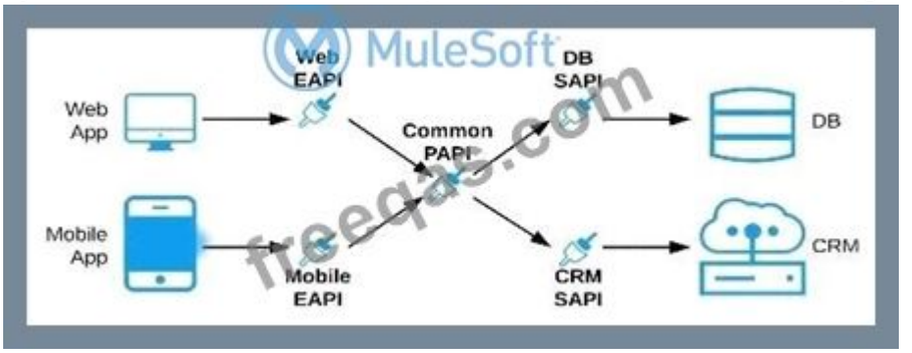


NEW QUESTION: 30
A company is designing a mule application to consume batch data from a partner's ftps server The data files have been compressed and then digitally signed using PGP.
What inputs are required for the application to securely consumed these files?
A. ATLS context first store containing a public certificate for partner ftps server and the PGP public key of the partner TLS contact Key Store containing the FTP credentials
B. The PGP public key of the partner
The PGP private key for the company The FTP username and password
C. ATLS context Key Store requiring the private key and certificate for the company PGP public key of partner PGP private key for the company
D. TLS context trust or containing a public certificate for the ftps server The FTP username and password The PGP public key of the partner
Answer: B (LEAVE A REPLY)

NEW QUESTION: 31
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?

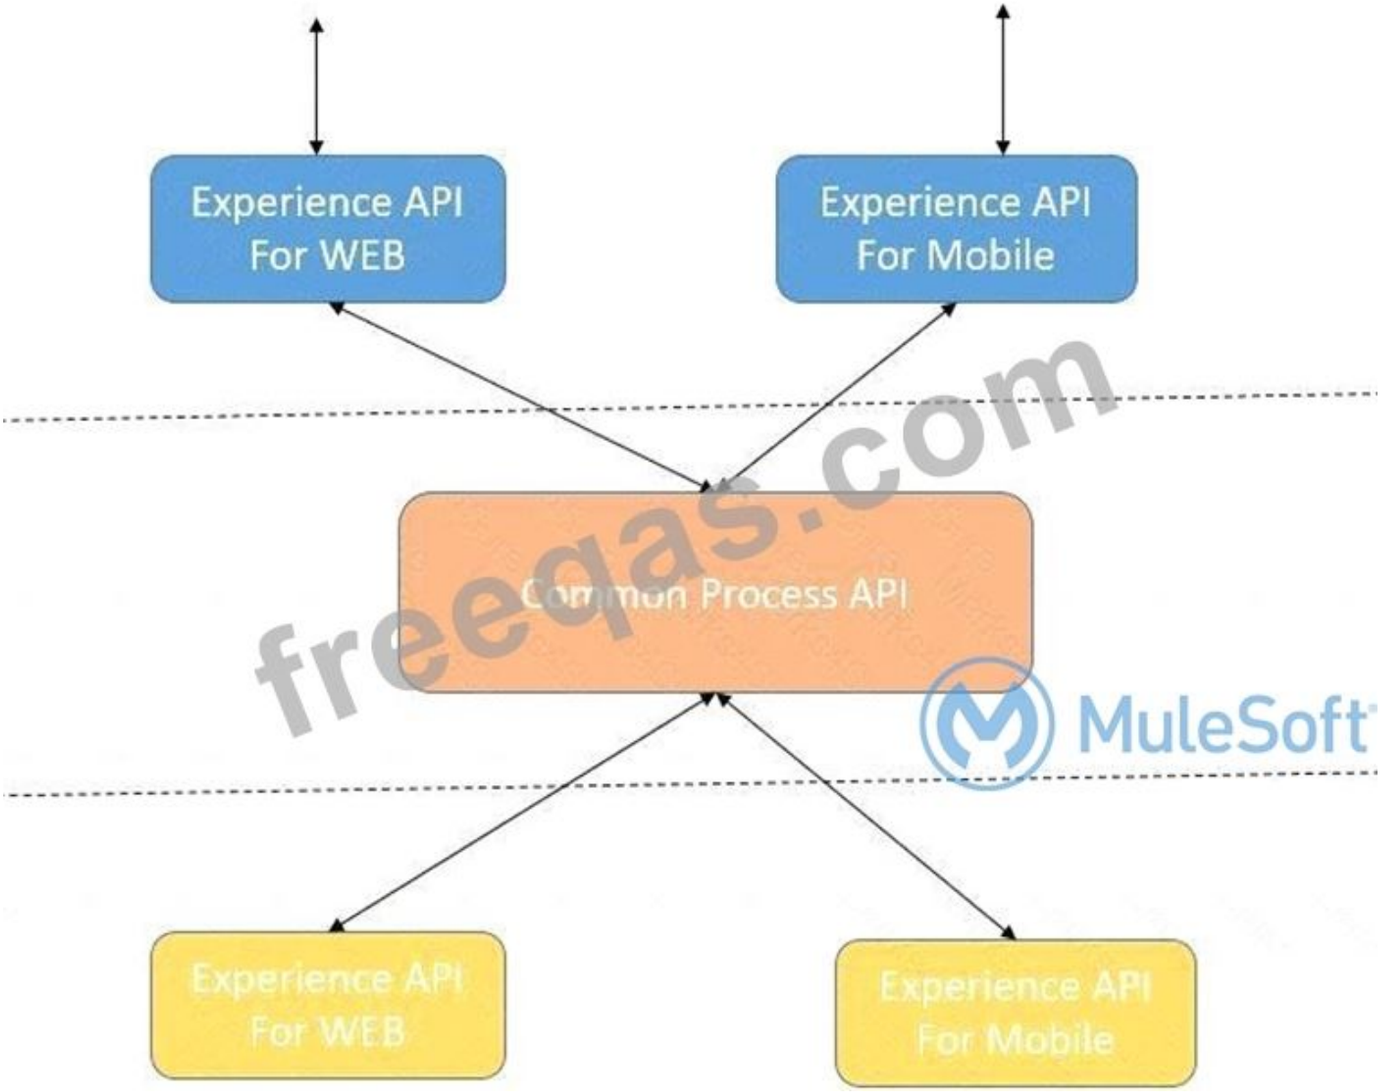


A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
B. One set of APIs (Experience API, Process 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 (LEAVE A REPLY)
Explanation
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.
Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.
Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system A screenshot of a computer Description automatically generated with low confidence



Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having

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



Valid MCIA-Level-1-Maintenance Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1-Maintenance Exam! PrepPdf.com now offer the **newest MCIA-Level-1-Maintenance exam dumps**, the PrepPdf.com MCIA-Level-1-Maintenance exam **questions have been**

updated and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1-Maintenance dumps with Test Engine here: https://www.preppdf.com/MuleSoft/MCIA-Level-1-Maintenance-prepaway-exam-dumps.html (**118** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 32
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 placement replica will be available and will take or processing the remaining 99,000 records
D. A new replacement replica will be available and will be process all 1,00,000 records from scratch leading to duplicate record processing
Answer: B (LEAVE A REPLY)

NEW QUESTION: 33
A rale limiting policy has been applied to a soap VI.2 API published in Clondhub. The API implementation catches errors in a global error handler on error propagate in the main flow for HTTP: RETRY_EXHAUSTED with HTTP status set to 429 and any with the HTTP status set to 500.
What is the expected H1TP status when the client exceeds the quota of the API calls?
A. HTTP status 400 from the rate-limiting policy violation since the call does not reach the back-end
B. 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
Answer: (SHOW ANSWER)

NEW QUESTION: 34
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 (LEAVE A REPLY)
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: 35
A new upstream API Is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity. The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms. If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API's desired SLA?
A. Set a timeout of 100 ms; that leaves 400 ms for the other two downstream APIs to complete
B. Do not set a timeout; the Invocation of this API Is mandatory and so we must wait until it responds
C. Set a timeout of 50 ms; this times out more invocations of that API but gives additional room for retries
D. No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API
Answer: (SHOW ANSWER)
Explanation
Before we answer this question , we need to understand what median (50th percentile) and 80th percentile means. If the 50th percentile (median) of a response time is 500ms that means that 50% of my transactions are either as fast or faster than 500ms.
If the 90th percentile of the same transaction is at 1000ms it means that 90% are as fast or faster and only 10% are slower. Now as per upstream SLA , 99th percentile is 800 ms which means 99% of the incoming requests should have response time less than or equal to 800 ms. But as per one of the backend API , their 95th percentile is 1000 ms which means that backend API will take 1000 ms or less than that for 95% of. requests.
As there are three API invocation from upstream API , we can not conclude a timeout that can be set to meet the desired SLA as backend SLA's do not support it.
Let see why other answers are not correct.
1) Do not set a timeout --> This can potentially violate SLA's of upstream API
2) Set a timeout of 100 ms; ---> This will not work as backend API has 100 ms as median meaning only 50% requests will be answered in this time and we will get timeout for 50% of the requests. Important thing to note here is, All APIs need to be executed sequentially, so if you get

timeout in first API, there is no use of going to second and third API. As a service provider you wouldn't want to keep 50% of your consumers dissatisfied. So not the best option to go with.
*To quote an example: Let's assume you have built an API to update customer contact details.
- First API is fetching customer number based on login credentials
- Second API is fetching Info in 1 table and returning unique key
- Third API, using unique key provided in second API as primary key, updating remaining details
* Now consider, if API times out in first API and can't fetch customer number, in this case, it's useless to call API 2 and 3 and that is why question mentions specifically that all APIs need to be executed sequentially.
3) Set a timeout of 50 ms --> Again not possible due to the same reason as above Hence correct answer is No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

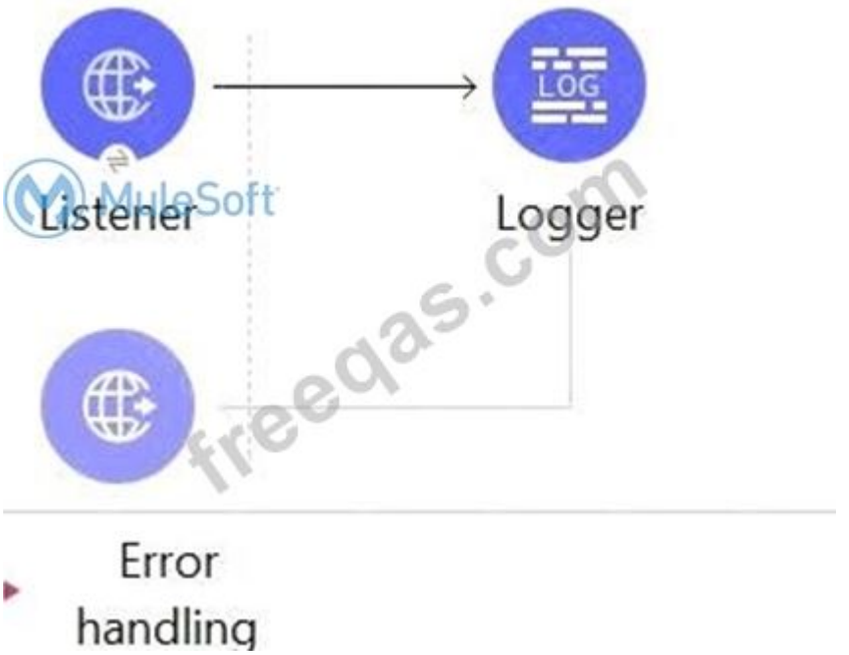
NEW QUESTION: 36
Customer has deployed mule applications to different customer hosted mule run times. Mule applications are managed from Anypoint platform.
What needs to be configured to monitor these Mule applications from Anypoint monitoring and what sends monitoring data to Anypoint monitoring?
A. Enable monitoring of individual applications from runtime manager application settings
Runtime manager agent sends monitoring data from the mule applications to Anypoint monitoring
B. By default, Anypoint monitoring agent will be installed on each Mule run time Anypoint Monitoring agent automatically sends monitoring data from the Mule applications to Anypoint monitoring
C. Install runtime manager agent on each mule runtime
Runtime manager agent since monitoring data from the mule applications to Anypoint monitoring
D. Anypoint monitoring agent on each mule runtime
Anypoint monitoring agent sends monitoring data from the mule applications to Anypoint monitoring
Answer: D (LEAVE A REPLY)

NEW QUESTION: 37
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. 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
B. 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
C. The configuration is built-in and nothing extra is required for the TLS context

D. Request a private key from Google and create a PKCS12 file with it and add it in keyStore as a part of TLS context
Answer: C (LEAVE A REPLY)

NEW QUESTION: 38
A company is using Mulesoft to develop API's and deploy them to Cloudhub and on premises targets.
Recently it has decided to enable Runtime Fabric deployment option as well and infrastructure is set up for this option.
What can be used to deploy Runtime Fabric?
A. AnypointCLI
B. Anypoint platform REST API's
C. Directly uploading ajar file from the Runtime manager
D. Mule maven plug-in
Answer: (SHOW ANSWER)

NEW QUESTION: 39
Refer to the exhibit.



The HTTP Listener and the Logger are being handled from which thread pools respectively?
A. UBER and NONBLOCKING
B. Shared Selector Pool and CPU LITE
C. BLOCKING _IO and UBER
D. CPU_INTENSIVE and Dedicated Selector pool
Answer: B (LEAVE A REPLY)

NEW QUESTION: 40

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

NEW QUESTION: 41
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 cluster configuration
B. Increase physical specifications of server CPU memory and network
C. Add a load balancer and add additional servers in a server group configuration
D. Change applications by use an event-driven model
Answer: (SHOW ANSWER)

NEW QUESTION: 42
What is a key difference between synchronous and asynchronous logging from Mule applications?
A. Synchronous logging writes log messages in a single logging thread but does not block the Mule event being processed by the next event processor
B. Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event
C. Asynchronous logging produces more reliable audit trails with more accurate timestamps

D. Synchronous logging within an ongoing transaction writes log messages in the same thread that processes the current Mule event

Answer: B (LEAVE A REPLY)

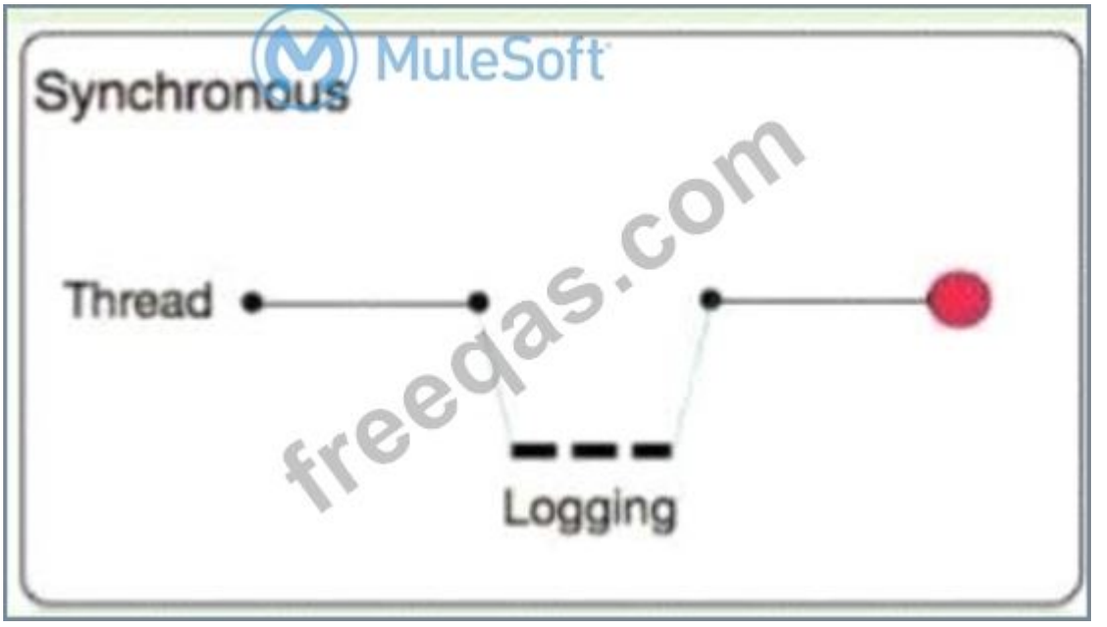
Explanation

Types of logging:

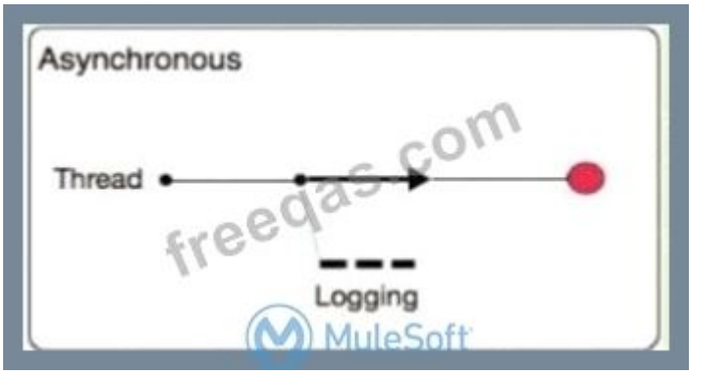
A) Synchronous: The execution of thread that is processing messages is interrupted to wait for the log message to be fully handled before it can continue.

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

Chart, diagram Description automatically generated



Chart, diagram, box and whisker chart Description automatically generated



B) Asynchronous:

The logging operation occurs in a separate thread, so the actual processing of your message won't be delayed to wait for the logging to complete Substantial improvement in throughput and latency of message processing Mule runtime engine (Mule) 4 uses Log4j 2 asynchronous logging by default The disadvantage of asynchronous logging is error handling.

If the logger fails to write to disk, the thread doing the processing won't be aware of any issues writing to the disk, so you won't be able to rollback anything. Because the actual writing of the log gets differed, there's a chance that log messages might never make it to disk and get lost, if Mule were to crash before the buffers are flushed.

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

Valid MCIA-Level-1-Maintenance Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1-Maintenance Exam! PrepPdf.com now offer the **newest MCIA-Level-1-Maintenance exam dumps**, the PrepPdf.com MCIA-Level-1-Maintenance exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1-Maintenance dumps with Test Engine here: https://www.preppdf.com/MuleSoft/MCIA-Level-1-Maintenance-prepaway-exam-dumps.html (**118** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)