

MuleSoft.MCIA-Level-1.v2022-07-26.q109

Exam Code:	MCIA-Level-1
Exam Name:	MuleSoft Certified Integration Architect - Level 1
Certification Provider:	MuleSoft
Free Question Number:	109
Version:	v2022-07-26
# of views:	2519
# of Questions views:	1090
https://www.freeqas.com/qa/MuleSoft/MCIA-Level-1/MuleSoft.MCIA-Level-1.v2022-07-26.q109.html	

NEW QUESTION: 1

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

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

Answer: C (LEAVE A REPLY)

Context of the question is about managing and governing mule applications deployed on Anypoint platform.

Anypoint API Manager (API Manager) is a component of Anypoint Platform that enables you to manage, govern, and secure APIs. It leverages the runtime capabilities of API Gateway and Anypoint Service Mesh, both of which enforce policies, collect and track analytics data, manage proxies, provide encryption and authentication, and manage applications.

Prerequisite of managing an API is that the API must be published to Anypoint Exchange. Hence the correct option is C Mule Ref Doc : <https://docs.mulesoft.com/api-manager/2.x/getting-started-proxy>

NEW QUESTION: 2

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

problem with any of the connectors Besides the database connector , what other connector can be used in the transaction.

A. VM

B. Anypoint MQ

C. SFTP

D. ObjectStore

Answer: ([SHOW ANSWER](#))

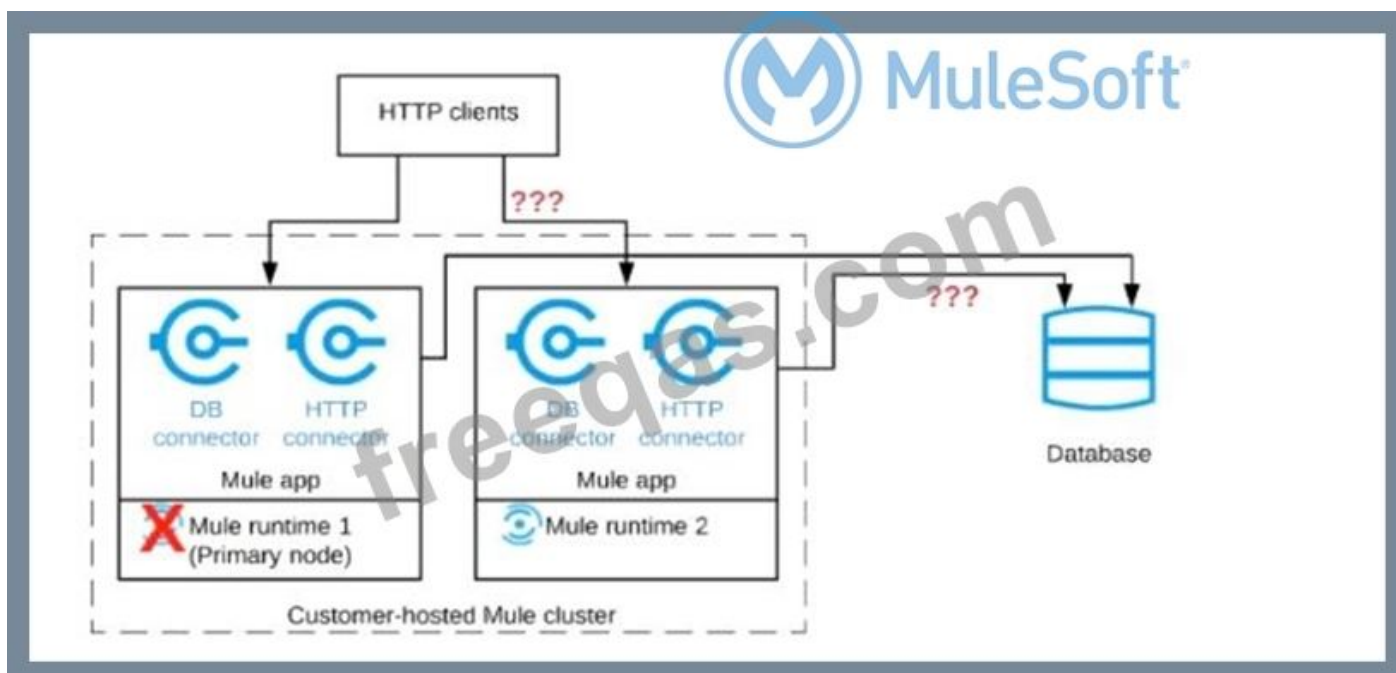
Correct answer is VM VM support Transactional Type. When an exception occur, The transaction rolls back to its original state for reprocessing. This feature is not supported by other connectors.

Here is additional information about Transaction management:

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

NEW QUESTION: 3

Refer to the exhibit.



A Mule application is deployed to a cluster of two customer-hosted Mule runtimes. The Mule application has a flow that polls a database and another flow with an HTTP Listener.

HTTP clients send HTTP requests directly to individual cluster nodes.

What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has failed, but before that node is restarted?

- A. Database polling continues All HTTP requests continue to be accepted, but requests to the failed node incur increased latency
- B. Database polling stops All HTTP requests are rejected
- C. Database polling stops All HTTP requests continue to be accepted
- D. Database polling continues Only HTTP requests sent to the remaining node continue to be accepted

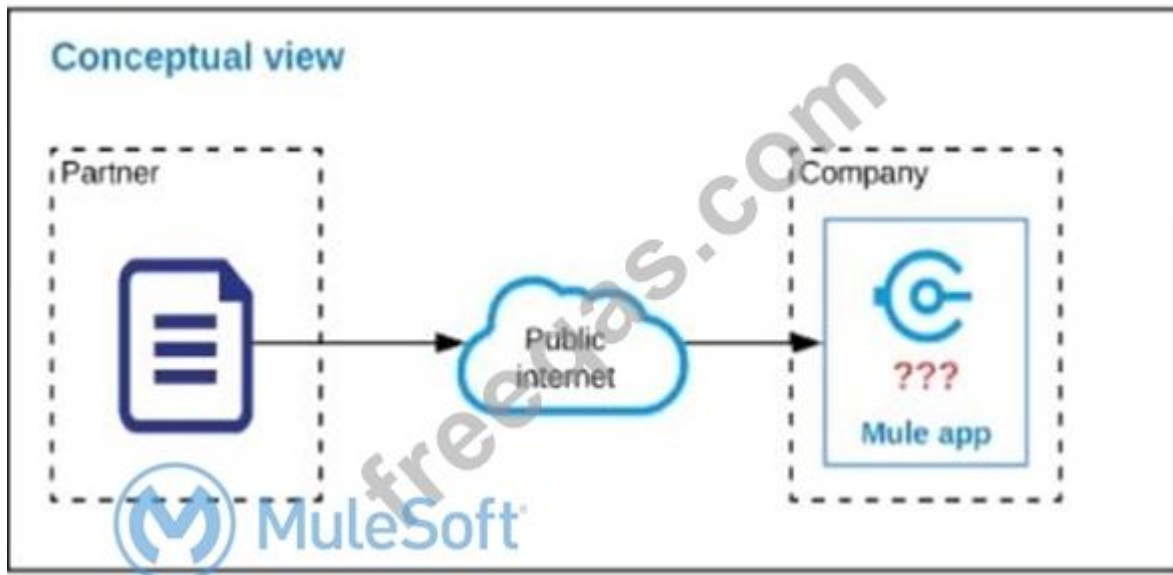
Answer: ([SHOW ANSWER](#))

NEW QUESTION: 4

Refer to the exhibit.

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

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



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

Answer: (SHOW ANSWER)

* Object Store and VM Store is used for sharing data inter or intra mule applications in same setup. Can't be used with external Business Partner

* Also File connector will not be useful as the two companies currently have no shared IT infrastructure. It's specific for local use.

* Correct answer is SFTP connector. The SFTP Connector implements a secure file transport channel so that your Mule application can exchange files with external resources. SFTP uses the SSH security protocol to transfer messages. You can implement the SFTP endpoint as an inbound endpoint with a one-way exchange pattern, or as an outbound endpoint configured for either a one-way or request-response exchange pattern.

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

Answer: A (LEAVE A REPLY)

NEW QUESTION: 6

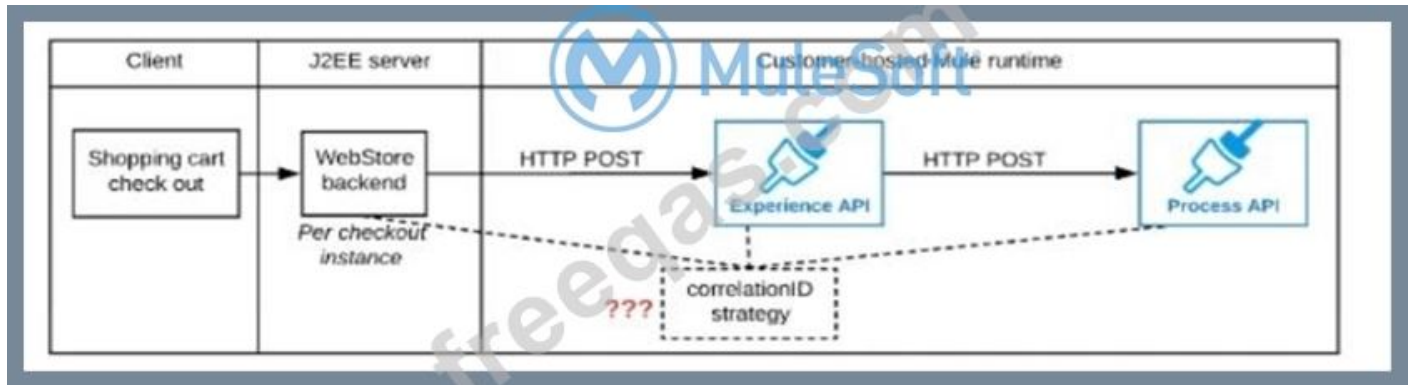
What requires configuration of both a key store and a trust store for an HTTP Listener?

- A. Encryption of both HTTP request and HTTP response bodies for all HTTP clients
- B. Encryption of requests to both subdomains and API resource endpoints
fhttps://aDi.customer.com/ and
https://customer.com/api)
- C. Encryption of both HTTP request header and HTTP request body for all HTTP clients
- D. Support for TLS mutual (two-way) authentication with HTTP clients

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 7

Refer to the exhibit.

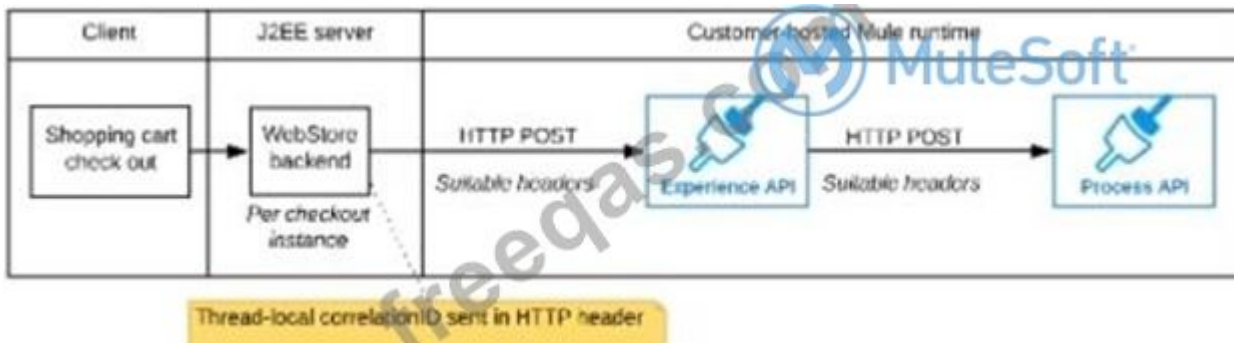


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations are Mule applications executing in a customer -hosted Mule runtime. End-to-end correlation of all HTTP requests and responses belonging to each individual checkout Instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the same correlation ID for all requests and responses belonging to the same checkout instance.

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

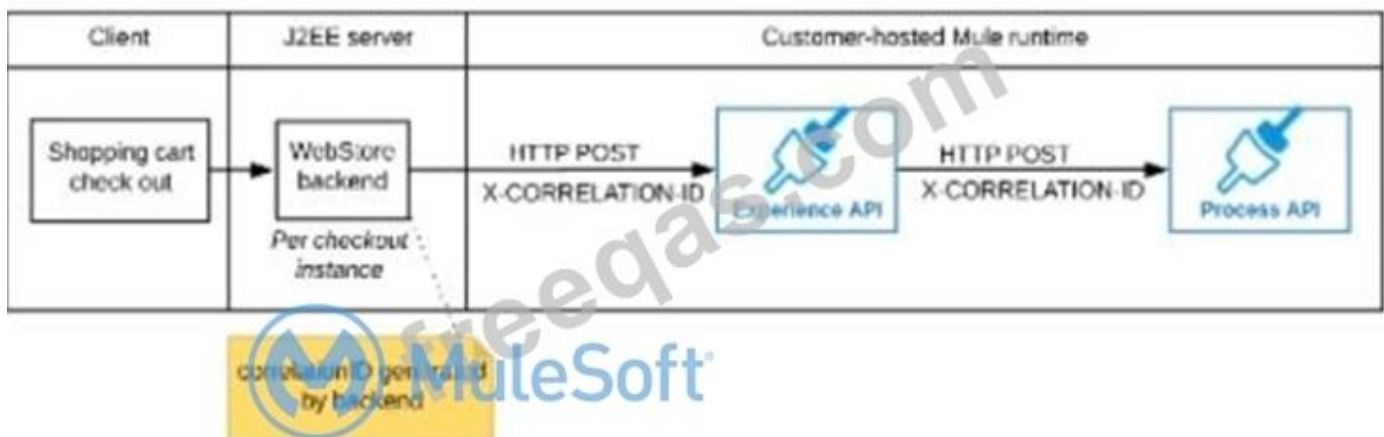
A)

The web store backend, being a Java EE application, automatically makes use of the thread-local correlation ID generated by the Java EE application server and automatically transmits that to the Experience API using HTTP-standard headers. No special code or configuration is included in the web store backend, Experience API, and Process API implementations to generate and manage the correlation ID.



B)

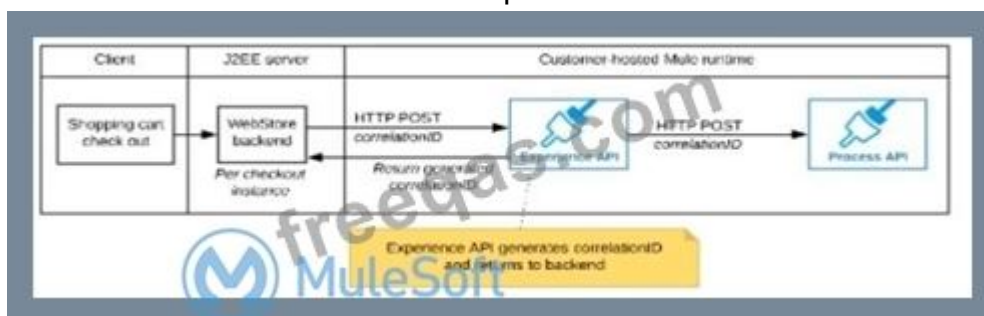
The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header. In each API invocation belonging to that checkout, no special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID.



C)

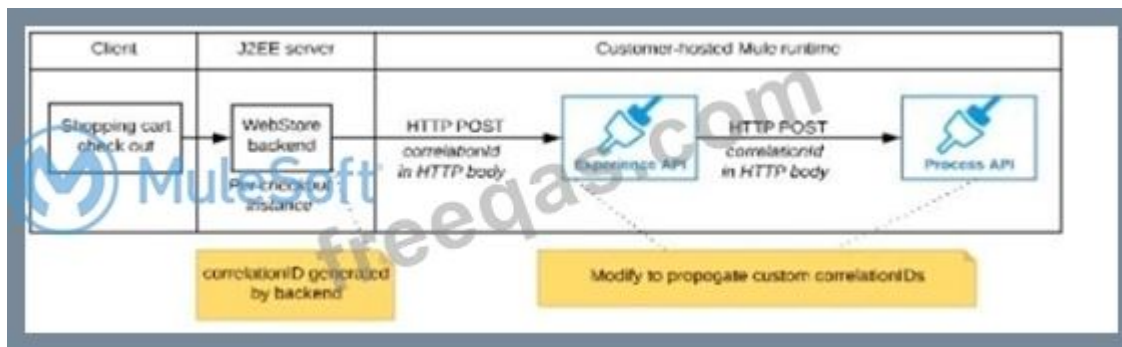
The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

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



D)

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



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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 8

An API client is implemented as a Mule application that includes an HTTP Request operation using a default configuration. The HTTP Request operation invokes an external API that follows standard HTTP status code conventions, which causes the HTTP Request operation to return a 4xx status code.

What is a possible cause of this status code response?

- A. The external API reported an error with the HTTP request that was received from the outbound HTTP Request operation of the Mule application
- B. An error occurred inside the external API implementation when processing the HTTP request that was received from the outbound HTTP Request operation of the Mule application
- C. The external API reported that the API implementation has moved to a different external endpoint
- D. The HTTP response cannot be interpreted by the HTTP Request operation of the Mule application after it was received from the external API

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 9

An Order microservice and a Fulfillment microservice are being designed to communicate with their clients through message-based integration (and NOT through API invocations).

The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mule application, the Fulfillment microservice.

The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message (a kind of event message). Each OrderFulfilled message can be consumed by any interested Mule application, and the Order microservice is one such Mule application.

What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

A. Order messages are sent to an Anypoint MQ exchange

OrderFulfilled messages are sent to an Anypoint MQ queue

Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices

B. Order messages are sent directly to the Fulfillment microservices

OrderFulfilled messages are sent directly to the Order microservice

The Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load each microservice

C. Order messages are sent to a JMS queue OrderFulfilled messages are sent to a JMS topic

Both microservices interact with the same JMS provider (message broker) Instance, which must therefore scale to support the load of both microservices

D. Order messages are sent to a JMS queue OrderFulfilled messages are sent to a JMS topic

The Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice

Answer: D (LEAVE A REPLY)

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 route of a Parallel For Each scope

B. Call each API request in a separate lookup call from Dataweave reduce operator

C. Call each API request in a separate route of a Scatter-Gather

D. Call each API request in a separate Async scope

Answer: B (LEAVE A REPLY)

NEW QUESTION: 11

An organization currently uses a multi-node Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications. The organization is planning to transition to a deployment model based on Docker containers in a Kubernetes cluster. The organization has already created a standard Docker image containing a Mule runtime and all required dependencies (including a JVM), but excluding the Mule application itself.

What is an expected outcome of this transition to container-based Mule application deployments?

A. Required redesign of Mule applications to follow microservice architecture principles

B. Required change to the URL endpoints used by clients to send requests to the Mule applications

C. Required migration to the Docker and Kubernetes-based Anypoint Platform - Private Cloud Edition

D. Guaranteed consistency of execution environments across all deployments of a Mule application

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 12

A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application.

What non-functional requirement can be expected to be negatively affected when implementing this change?

A. Availability

B. Response time

C. Consistency

D. Throughput

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 13

A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL `https://api.acme.com/partners/v1`. The API implementation is estimated to require deployment to 5 CloudHub workers.

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

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

A. Add the X.509 certificate to the Mule application's deployable archive, then configure a CloudHub Dedicated Load Balancer (DLB) for each of the Mule application's CloudHub workers
Create a CNAME for `api.acme.com` pointing to the DLB's A record

B. Add the X.509 certificate to a CloudHub Dedicated Load Balancer (DLB), not to the Mule application
Create a CNAME for `api.acme.com` pointing to the DLB's A record

C. Add the X.509 certificate to the Mule application's deployable archive, then configure the CloudHub Shared Load Balancer (SLB) for each of the Mule application's CloudHub workers
Create a CNAME for `api.acme.com` pointing to the SLB's A record

D. Add the X.509 certificate to the CloudHub Shared Load Balancer (SLB), not to the Mule application
Create a CNAME for `api.acme.com` pointing to the SLB's A record

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 14

An organization has decided on a cloud migration strategy that aims to minimize the organization's own IT resources. Currently, the organization has all of its Mule applications running on its own premises and uses an on-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 Mule application deployment and load balancing that preserves the public URLs?

A. Deploy the Mule applications to CloudHub

Update the CNAME record for `api.acme.com` in the organization's 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

B. For each migrated Mule application, deploy an API Proxy Mule application to CloudHub, with all applications under the control of a CloudHub dedicated load balancer (DLB) Update the CNAME record for `api.acme.com` in the organization's 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 application

C. Deploy the Mule applications to CloudHub

Update the CNAME record for `api.acme.com` in the organization's 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

D. Deploy the Mule applications to CloudHub

Create a CNAME record for `api.acme.com` in the CloudHub Shared Load Balancer (SLB) pointing to the A record of the on-premises load balancer Apply mapping rules in the SLB to map URLs to their corresponding Mule applications

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 15

An organization is designing the following two Mule applications that must share data via a common persistent object store instance:

- Mule application P will be deployed within their on-premises datacenter.
- Mule application C will run on CloudHub in an Anypoint VPC.

The object store implementation used by CloudHub is the Anypoint Object Store v2 (OSv2).

what type of object store(s) should be used, and what design gives both Mule applications access to the same object store instance?

A. Application C uses the Object Store connector to access a persistent object Application P accesses the persistent object store via the Object Store REST API

B. Application C and P both use the Object Store connector to access the Anypoint Object Store v2

C. Application C and P both use the Object Store connector to access a persistent object store

D. Application P uses the Object Store connector to access a persistent object store Application C accesses this persistent object store via the Object Store REST API through an IPsec tunnel

Answer: A (LEAVE A REPLY)

NEW QUESTION: 16

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

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

Answer: D (LEAVE A REPLY)

Valid MCIA-Level-1 Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1 Exam! PrepPdf.com now offer the **newest MCIA-Level-1 exam dumps**, the PrepPdf.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here: <https://www.preppdf.com/MuleSoft/MCIA-Level-1-prepaway-exam-dumps.html> (275 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 17

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 (LEAVE A REPLY)

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: 18

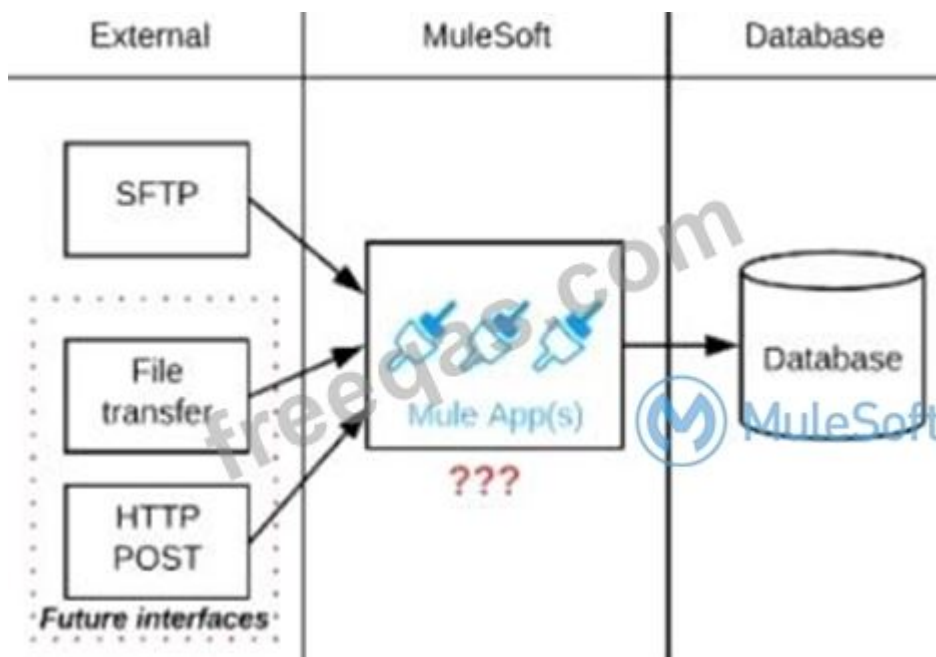
What is required before an API implemented using the components of Anypoint Platform can be managed and governed (by applying API policies) on Anypoint Platform?

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

Answer: D (LEAVE A REPLY)

NEW QUESTION: 19

Refer to the exhibit.



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

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

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

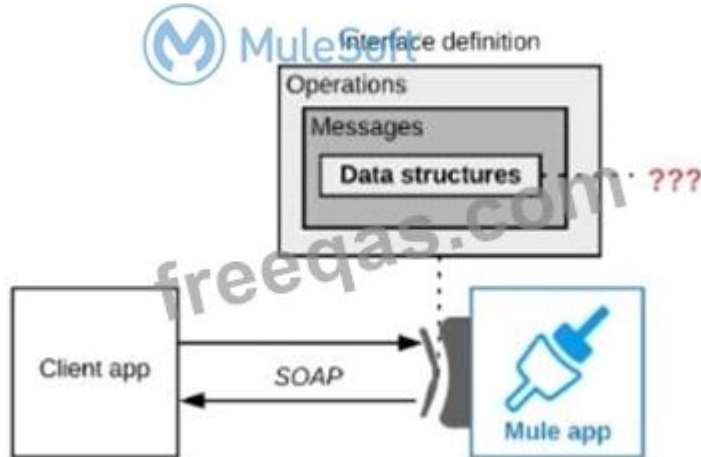
Answer: ([SHOW ANSWER](#))

NEW QUESTION: 20

Refer to the exhibit.

A Mule application is being designed to expose a SOAP web service to its clients.

What language is typically used inside the web service's interface definition to define the data structures that the web service is expected to exchange with its clients?



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

Answer: B ([LEAVE A REPLY](#))

Correct answer: XSD In this approach to developing a web service, you begin with

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

NEW QUESTION: 21

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

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

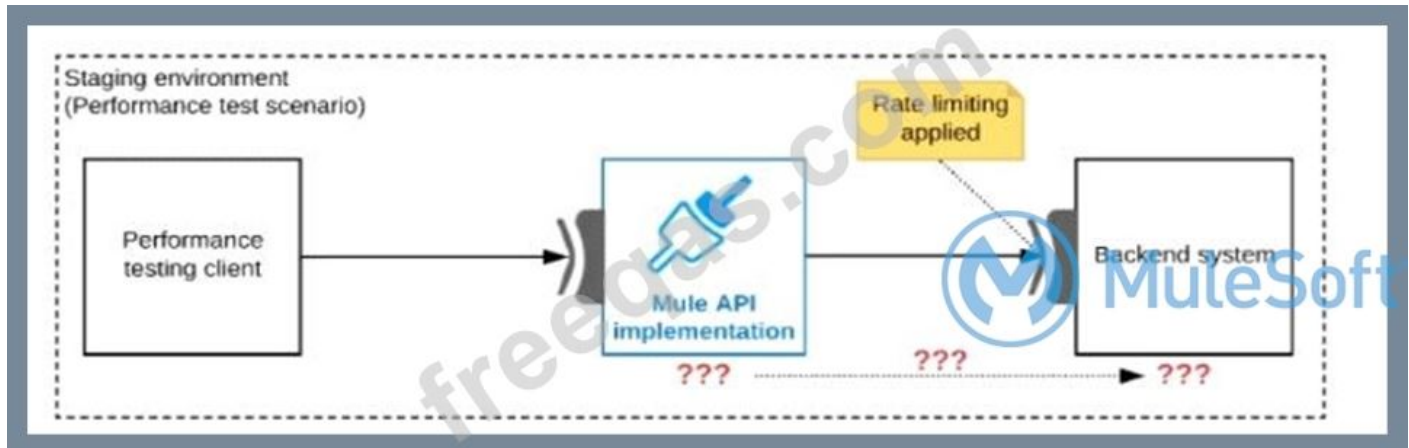
What combination of Mule components is most idiomatic (used according to their intended purpose) when implementing the above requirements?

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 22

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

Answer: D (LEAVE A REPLY)

NEW QUESTION: 23

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

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

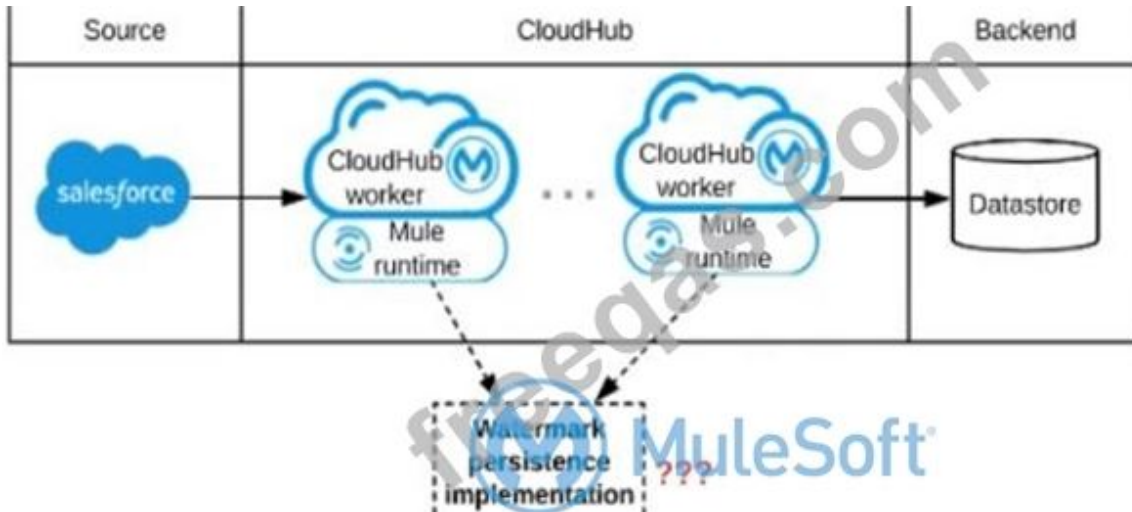
- A.** The credentials to access the backend systems and contact details for the administrator of each system
- B.** A list of current and future consumers of the Mule application and their contact details

- C. Snapshots of the Mule application's flows, including their error handling
- D. SLAs and non-functional requirements to access the backend systems

Answer: D (LEAVE A REPLY)

NEW QUESTION: 24

Refer to the exhibit.



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

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

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

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

Answer: (SHOW ANSWER)

NEW QUESTION: 25

A Mule application uses an HTTP Request operation to invoke an external API.

The external API follows the HTTP specification for proper status code usage.

What is a possible cause when a 3xx status code is returned to the HTTP Request operation from the external API?

- A. The request was ACCEPTED by the external API
- B. The request was REDIRECTED to a different URL by the external API
- C. The request was NOT RECEIVED by the external API
- D. The request was NOT ACCEPTED by the external API

Answer: B (LEAVE A REPLY)

Explanation/Reference: <https://www.w3.org/Protocols/rfc2616/rfc2616-sec10.html>

NEW QUESTION: 26

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)

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: 27

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 (LEAVE A REPLY)

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: 28

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 a DataWeave reduce operator
- B. Call each API request in a separate Async scope
- C. Call each API request in a separate route of a Scatter-Gather
- D. Call each API request in a separate route of a Parallel For Each scope

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 29

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 30

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

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

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

Answer: (SHOW ANSWER)

Security Assertion Markup Language (SAML) is an open standard that allows identity providers (IdP) to pass authorization credentials to service providers (SP).

SAML transactions use Extensible Markup Language (XML) for standardized communications between the identity provider and service providers.

SAML is the link between the authentication of a user's identity and the authorization to use a service.

WS-Security is the key extension that supports many authentication models including: basic username/password credentials, SAML, OAuth and more.

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

NEW QUESTION: 31

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 (LEAVE A REPLY)

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

Valid MCIA-Level-1 Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1 Exam! PrepPdf.com now offer the **newest MCIA-Level-1 exam dumps**, the PrepPdf.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here:
<https://www.preppdf.com/MuleSoft/MCIA-Level-1-prepaway-exam-dumps.html> (275 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 32

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 (LEAVE A REPLY)

* 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:

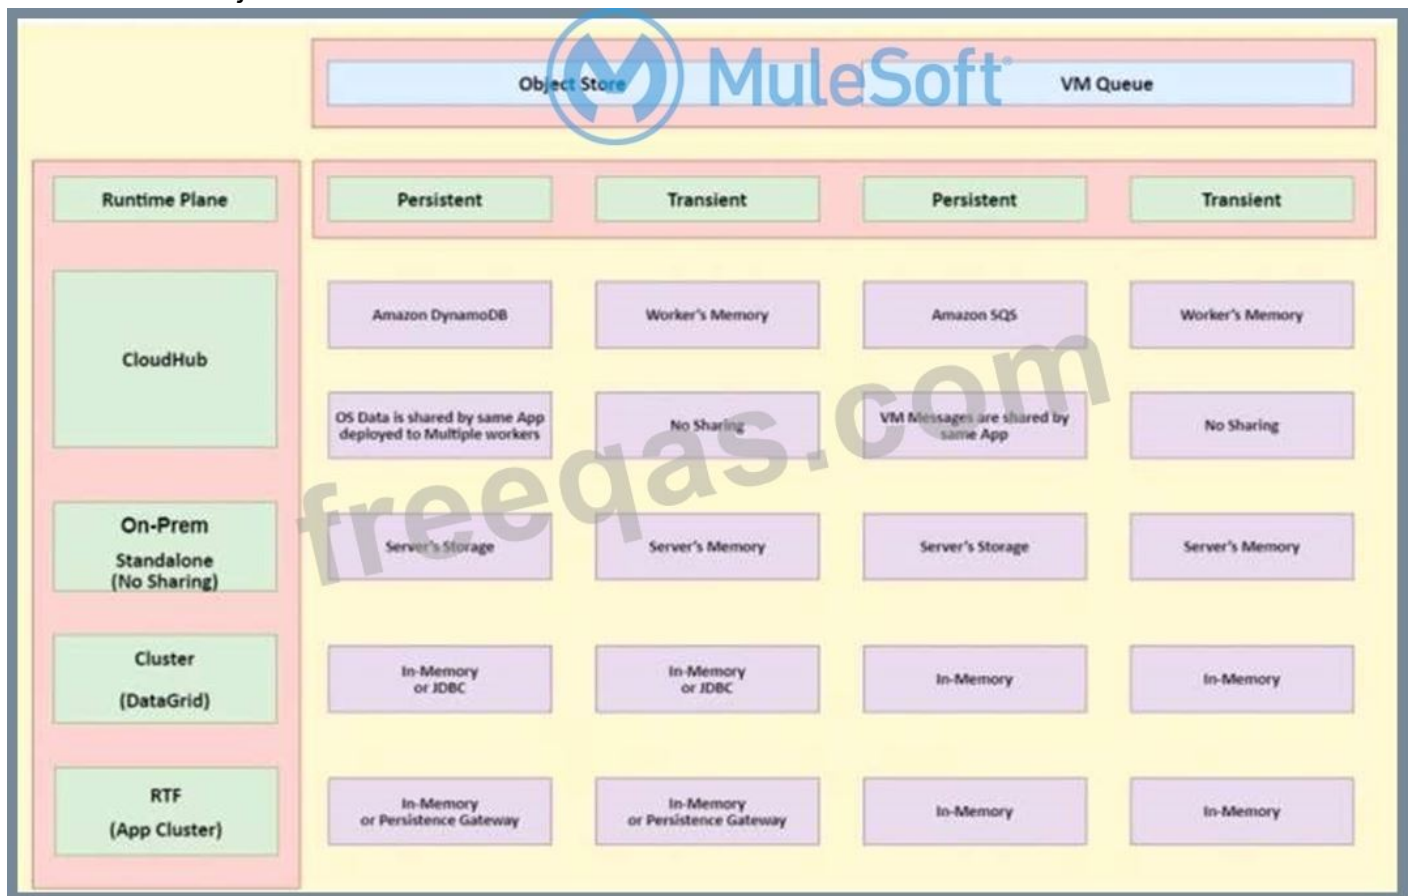
Object Store API

Object Store Stats API

Reference:

Additional Info:

When to use Object Store and when to use VM



NEW QUESTION: 33

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

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

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

Answer: (SHOW ANSWER)

* As per definition of API by Mulesoft , it is Application Programming Interface using HTTP-based protocols. Non-HTTP-based programmatic interfaces are not APIs.

* HTTP-based programmatic interfaces are APIs even if they don't use REST or JSON. Hence implementation based on Java RMI, CORBA/IIOP, raw TCP/IP interfaces are not API's as they are not using HTTP.

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

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

Considering the above points only correct option is

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

NEW QUESTION: 34

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 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

B. 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

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 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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 35

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: ([SHOW ANSWER](#))

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: 36

An organization is designing a Mule application to periodically poll an SFTP location for new files containing sales order records and then process those sales orders. Each sales order must be processed exactly once.

To support this requirement, the Mule application must identify and filter duplicate sales orders on the basis of a unique ID contained in each sales order record and then only send the new sales orders to the downstream system.

What is the most idiomatic (used for its intended purpose) Anypoint connector, validator, or scope that can be configured in the Mule application to filter duplicate sales orders on the basis of the unique ID field contained in each sales order record?

A. Configure a watermark In an On New or Updated File event source to filter unique records by the order ID

B. Configure a Database connector to filter and store each record by the order ID

C. Configure a Cache scope to filter and store each record from the received file by the order ID

D. Configure an Idempotent Message Validator component to filter each record by the order ID

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 37

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

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

Answer: (SHOW ANSWER)

Correct answer is Workers are randomly distributed across available AZs within that region. This ensure high availability for deployed mule applications Mulesoft documentation reference : <https://docs.mulesoft.com/runtime-manager/cloudhub-hadr>

NEW QUESTION: 38

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 (LEAVE A REPLY)

* "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: 39

What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

- A.** Only a DLB allows the configuration of a custom TLS server certificate
- B.** Only the SLB can forward HTTP traffic to the VPC-internal ports of the CloudHub workers
- C.** Both a DLB and the SLB allow the configuration of access control via IP whitelists
- D.** Both a DLB and the SLB implement load balancing by sending HTTP requests to workers with the lowest workloads

Answer: (SHOW ANSWER)

- * Shared load balancers don't allow you to configure custom SSL certificates or proxy rules
- * Dedicated Load Balancer are optional but you need to purchase them additionally if needed.
- * TLS is a cryptographic protocol that provides communications security for your Mule app. TLS offers many different ways of exchanging keys for authentication, encrypting data, and guaranteeing message integrity.
- * The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
- * Only a DLB allows the configuration of a custom TLS server certificate
- * DLB enables you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
- * To use a DLB in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.
- * MuleSoft Reference: <https://docs.mulesoft.com/runtime-manager/dedicated-load-balancer-tutorial> Additional Info on SLB Vs DLB:

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

NEW QUESTION: 40

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 BATCH STEP INSTANCES execute IN PARALLEL to process records and Batch Steps in ANY order as fast as possible

B. 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

C. 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

D. 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

Answer: A (LEAVE A REPLY)

Explanation/Reference: <https://docs.mulesoft.com/mule-runtime/3.8/batch-processing>

NEW QUESTION: 41

A Mule application uses the Database connector.

What condition can the Mule application automatically adjust to or recover from without needing to restart or redeploy the Mule application?

A. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade

B. The database server was unavailable for four hours due to a major outage but is now fully operational again

C. One of the stored procedures being called by the Mule application has been renamed

D. The credentials for accessing the database have been updated and the previous credentials are no longer valid

Answer: D (LEAVE A REPLY)

NEW QUESTION: 42

A Mule application uses the Database connector.

What condition can the Mule application automatically adjust to or recover from without needing to restart or redeploy the Mule application?

A. One of the stored procedures being called by the Mule application has been renamed

B. The database server was unavailable for four hours due to a major outage but is now fully operational again

C. The credentials for accessing the database have been updated and the previous credentials are no longer valid

D. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade

Answer: B (LEAVE A REPLY)

- * Any change in the application will require a restart except when the issue outside the app. For below situations , you would need to redeploy the code after doing necessary changes
- One of the stored procedures being called by the Mule application has been renamed. In this case, in the Mule application you will have to do changes to accommodate the new stored procedure name.
- Required redesign of Mule applications to follow microservice architecture principles. As code is changed, deployment is must
- If the credentials changed and you need to update the connector or the properties.
- The credentials for accessing the database have been updated and the previous credentials are no longer valid. In this situation you need to restart or redeploy depending on how credentials are configured in Mule application.
- * So Correct answer is The database server was unavailable for four hours due to a major outage but is now fully operational again as this is the only external issue to application.

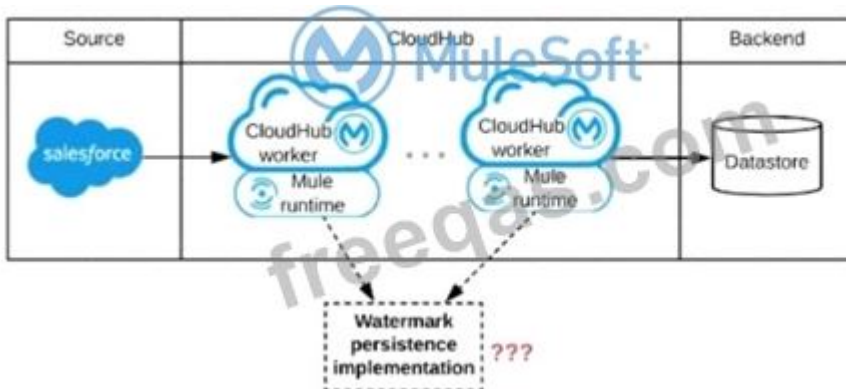
NEW QUESTION: 43

Refer to the exhibit.

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

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

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



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

Answer: C (LEAVE A REPLY)

NEW QUESTION: 44

An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform. What feature of Anypoint Platform should be used to meet this requirement?

- A. Federated Client Management
- B. Federated Identity Management
- C. Enterprise Security Module
- D. Client ID Secret

Answer: B (LEAVE A REPLY)

NEW QUESTION: 45

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.

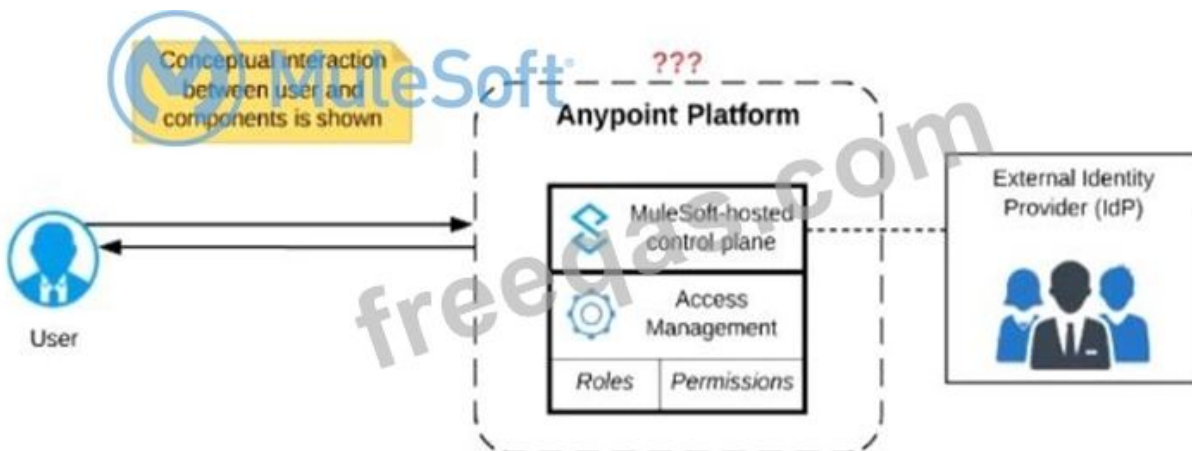
To minimize the overall latency. What is the most idiomatic (used for its intended purpose) design to call each API request in the Mule application?

- 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 Parallel For Each scope
- C. Call each API request in a separate route of a Scatter-Gather
- D. Call each API request in a separate Async scope

Answer: C (LEAVE A REPLY)

NEW QUESTION: 46

Refer to the exhibit.



Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.

What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled via the external Identity Provider?

- A. Controlling the business group within Anypoint Platform to which the user belongs
- B. Assigning Anypoint Platform permissions to a role
- C. Assigning Anypoint Platform role(s) to a user
- D. Removing a user's access to Anypoint Platform when they no longer work for the organization

Answer: B (LEAVE A REPLY)

* By default, Anypoint Platform performs its own user management

- For user management, one external IdP can be integrated with the Anypoint Platform organization (note: not at business group level)
- Permissions and access control are still enforced inside Anypoint Platform and CANNOT be controlled via the external Identity Provider * As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). * You can map users in a federated organization's group to a role which also gives the flexibility of controlling the business group within Anypoint Platform to which the user belongs to. Also user can nbe removed from external identity management system when they no longer work for the organization. So they wont be able to authenticate using SSO to login to Anypoint Platform. * Using external identity we can no change permissions of a particular role in Mulesoft Anypoint platform.
- * So Correct answer is Assigning Anypoint Platform permissions to a role

Valid MCIA-Level-1 Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1 Exam! PrepPdf.com now offer the **newest MCIA-Level-1 exam dumps**, the PrepPdf.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here:
<https://www.preppdf.com/MuleSoft/MCIA-Level-1-prepaway-exam-dumps.html> (275 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 47

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 (LEAVE A REPLY)

* 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: 48

Insurance organization is planning to deploy Mule application in MuleSoft Hosted runtime plane. As a part of requirement , application should be scalable . highly available. It also has regulatory requirement which demands logs to be retained for at least 2 years. As an Integration Architect what step you will recommend in order to achieve this?

- A.** It is not possible to store logs for 2 years in CloudHub deployment. External log management system is required.
- B.** When deploying an application to CloudHub , logs retention period should be selected as 2 years
- C.** When deploying an application to CloudHub, worker size should be sufficient to store 2 years data
- D.** Logging strategy should be configured accordingly in log4j file deployed with the application.

Answer: A (LEAVE A REPLY)

Correct answer is It is not possible to store logs for 2 years in CloudHub deployment. External log management system is required. CloudHub has a specific log retention policy, as described in the documentation: the platform stores logs of up to 100 MB per app & per worker or for up to 30 days, whichever limit is hit first. Once this limit has been reached, the oldest log information is deleted in chunks and is irretrievably lost. The recommended approach is to persist your logs to a external logging system of your choice (such as Splunk, for instance) using a log appender. Please note that this solution results in the logs no longer being stored on our platform, so any support cases you lodge will require for you to provide the appropriate logs for review and case resolution

NEW QUESTION: 49

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

- A.** Request size, request HTTP verbs, response time

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

Answer: (SHOW ANSWER)

Explanation/Reference: <https://docs.mulesoft.com/monitoring/api-analytics-dashboard>

NEW QUESTION: 50

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

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

Answer: D (LEAVE A REPLY)

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

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

The public CloudHub (shared) load balancer already redirects these requests, where myApp is the name of the Mule application deployment to CloudHub: HTTP requests to Error! Hyperlink reference not valid. redirects to Error! Hyperlink reference not valid.

HTTPS traffic to Error! Hyperlink reference not valid. redirects to Error! Hyperlink reference not valid.

NEW QUESTION: 51

An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform.

What feature of Anypoint Platform should be used to meet this requirement?

- A. Federated Client Management
- B. Federated Identity Management
- C. Enterprise Security Module
- D. Client ID Secret

Answer: (SHOW ANSWER)

Explanation/Reference: <https://docs.mulesoft.com/access-management/external-identity>

NEW QUESTION: 52

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

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

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 53

An Order microservice and a Fulfillment microservice are being designed to communicate with their clients through message-based integration (and NOT through API invocations).

The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mule application, the Fulfillment microservice.

The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message (a kind of event message). Each OrderFulfilled message can be consumed by any interested Mule application, and the Order microservice is one such Mule application.

What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

- A. Order messages are sent directly to the Fulfillment microservices
OrderFulfilled messages are sent directly to the Order microservice
The Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load of each microservice
- B. Order messages are sent to a JMS queue
OrderFulfilled messages are sent to a JMS topic
Both microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices
- C. Order messages are sent to a JMS queue
OrderFulfilled messages are sent to a JMS topic
The Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice
- D. Order messages are sent to an Anypoint MQ exchange
OrderFulfilled messages are sent to an Anypoint MQ queue
Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 54

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)

* 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:

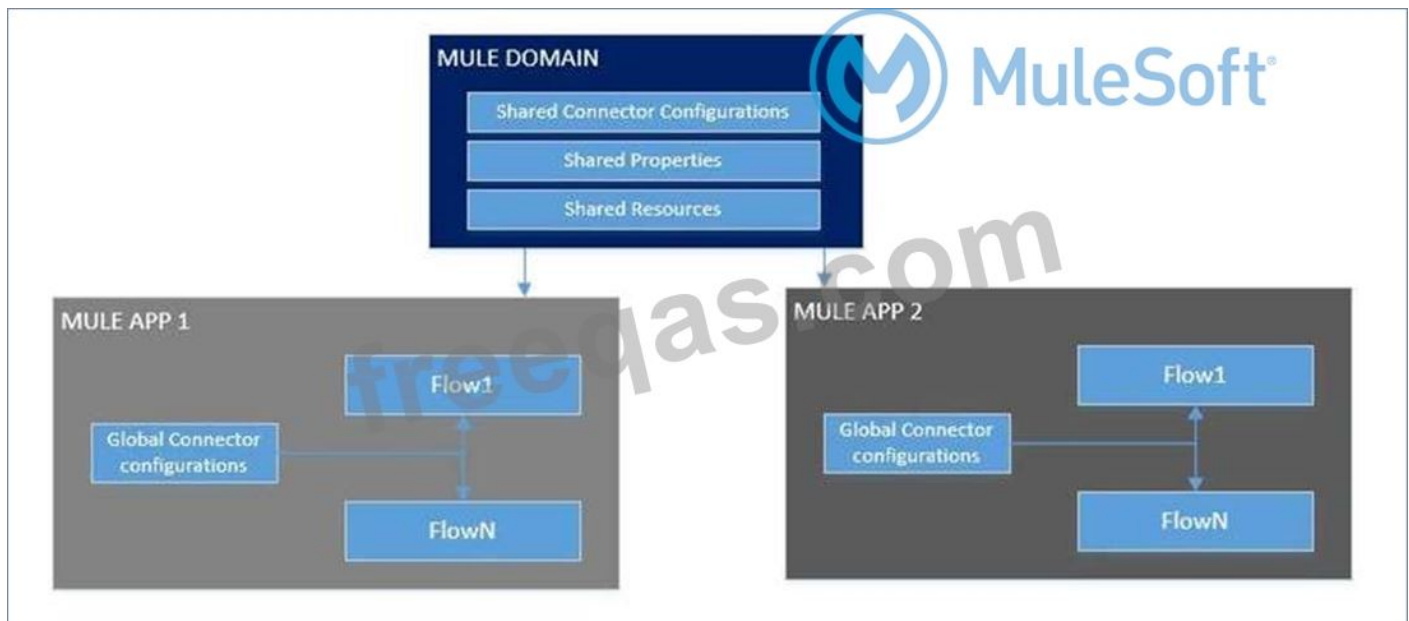
- Expose multiple services within the domain through the same port.
- Share the connection to persistent storage.
- Share services between apps through a well-defined interface.
- Ensure consistency between apps upon any changes because the configuration is only set in one place.

* 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



NEW QUESTION: 55

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

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

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

Answer: B (LEAVE A REPLY)

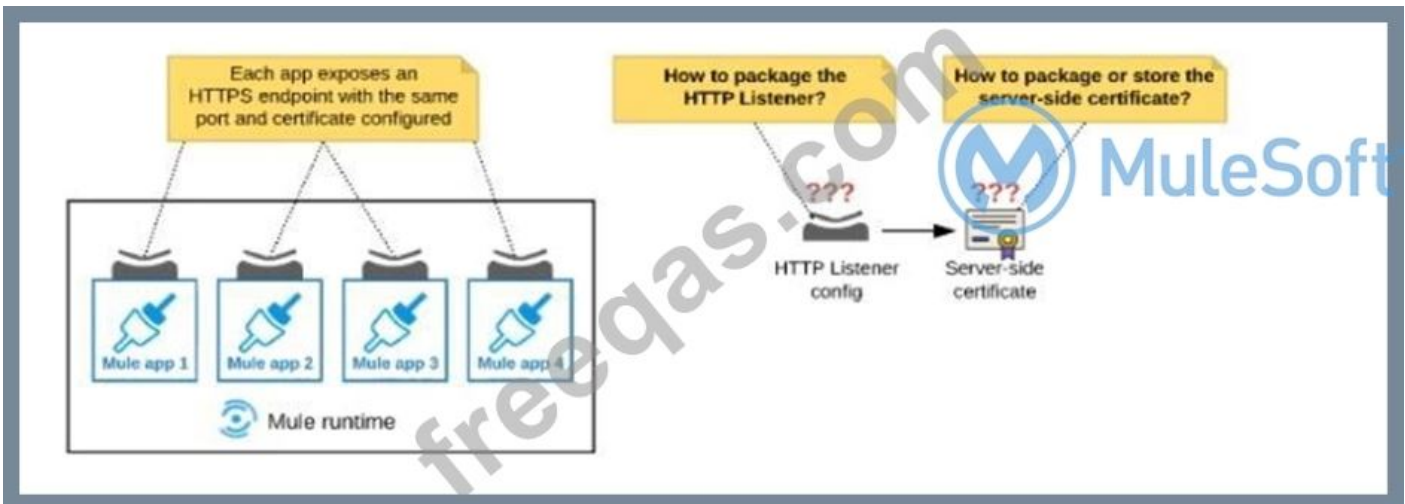
* Organization should continue to use SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching.

* Reason is that Anypoint exchange is not full fledged version repositories like GitHub.

* But at same time it is tightly coupled with Mule assets

NEW QUESTION: 56

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.

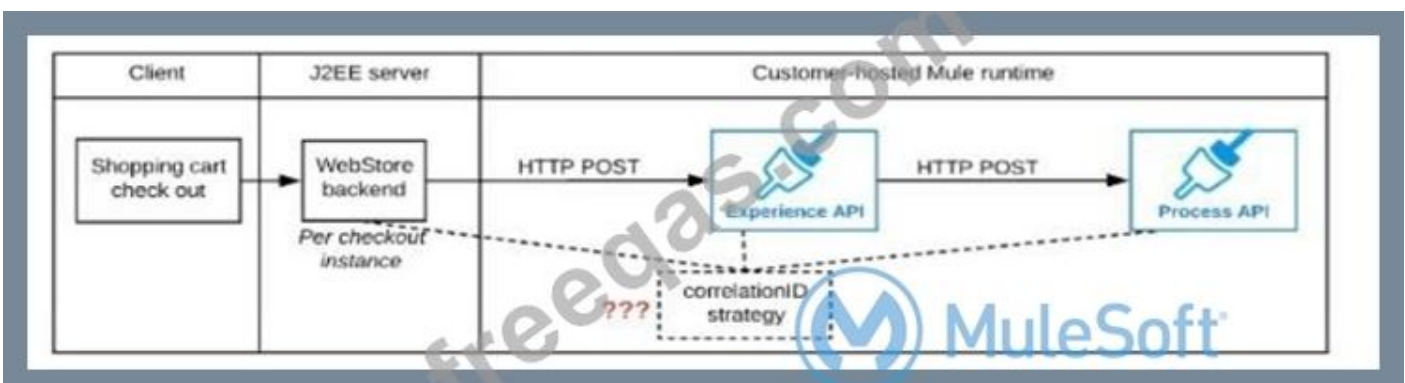
What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

- 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. Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION
- C. 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 the SAME Mule DOMAIN project Go to Set
- 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

Answer: (SHOW ANSWER)

NEW QUESTION: 57

Refer to the exhibit.

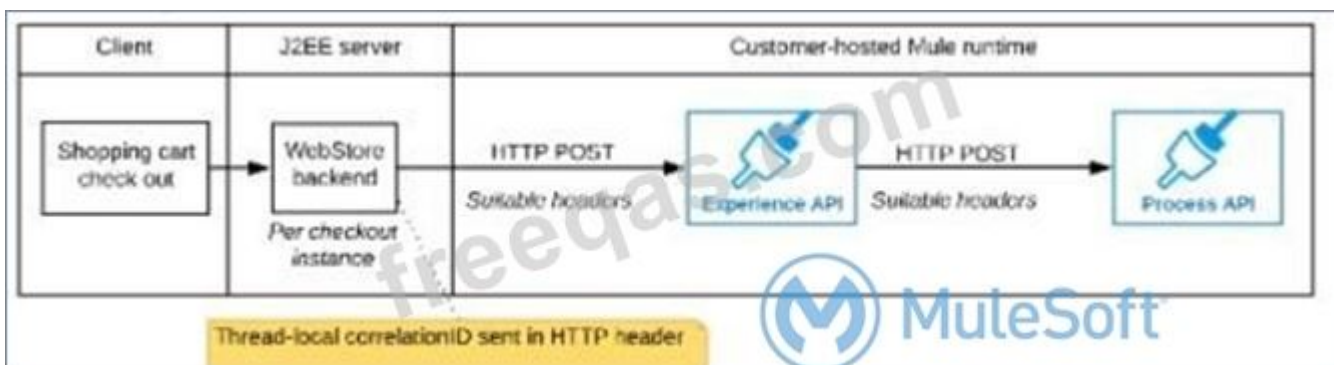


A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations are Mule applications executing in a customer -hosted Mule runtime. End-to-end correlation of all HTTP requests and responses belonging to each individual checkout Instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the same correlation ID for all requests and responses belonging to the same checkout instance.

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

A)

The web store backend, being a Java EE application, automatically makes use of the thread-local correlation ID generated by the Java EE application server and automatically transmits that to the Experience API using HTTP-standard headers. No special code or configuration is included in the web store backend, Experience API, and Process API implementations to generate and manage the correlation ID



B)

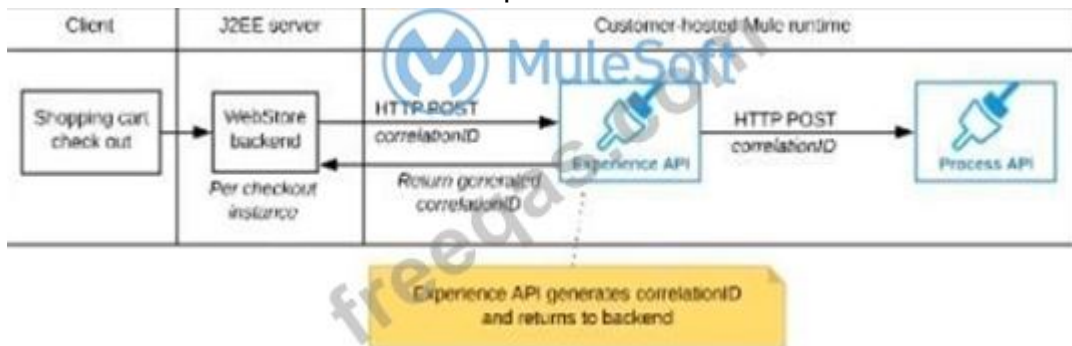
The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header. In each API invocation belonging to that checkout. No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID



C)

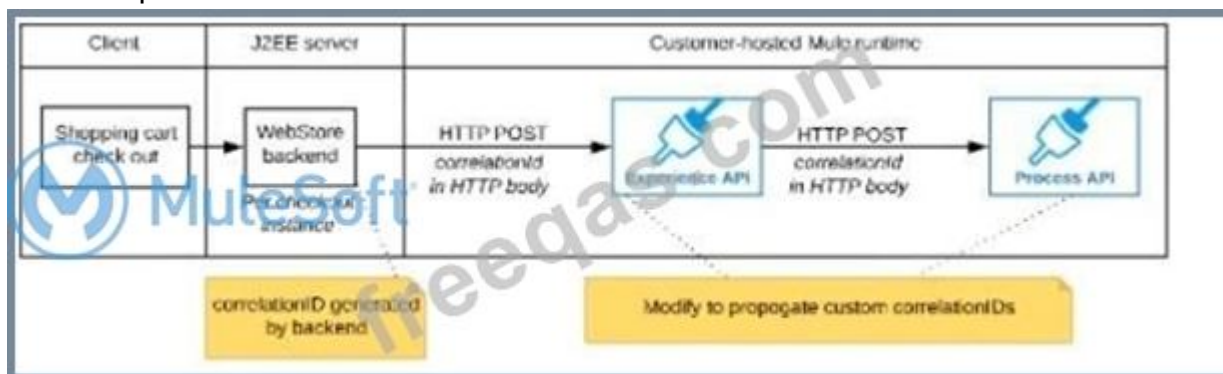
The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

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



D)

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



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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 58

What Mule application can have API policies applied by Anypoint Platform to the endpoint exposed by that Mule application?

- A. A Mule application that accepts requests over HTTP/1x
- B. A Mule application that accepts JSON requests over TCP but is NOT required to provide a response.
- C. A Mule application that accepts JSON requests over WebSocket
- D. A Mule application that accepts gRPC requests over HTTP/2

Answer: A ([LEAVE A REPLY](#))

* HTTP/1.1 keeps all requests and responses in plain text format.

* HTTP/2 uses the binary framing layer to encapsulate all messages in binary format, while still maintaining HTTP semantics, such as verbs, methods, and headers. It came into use in 2015, and offers several methods to decrease latency, especially when dealing with mobile platforms and server-intensive graphics and videos

* Currently, Mule application can have API policies only for Mule application that accepts requests over HTTP/1x

NEW QUESTION: 59

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 (LEAVE A REPLY)

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: 60

A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.

The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.

What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

- A.** Retrieve the RAML file(s) from the deployed Mule application
- B.** Retrieve the OpenAPI Specification file(s) from API Manager

- C. Retrieve the XML file(s) from Runtime Manager
- D. Retrieve the WSDL file(s) from the deployed Mule application

Answer: C (LEAVE A REPLY)

NEW QUESTION: 61

An API has been updated in Anypoint Exchange by its API producer from version 3.1.1 to 3.2.0 following accepted semantic versioning practices and the changes have been communicated via the API's public portal. The API endpoint does NOT change in the new version. How should the developer of an API client respond to this change?

- A. The update should be identified as a project risk and full regression testing of the functionality that uses this API should be run.
- B. The API producer should be contacted to understand the change to existing functionality.
- C. The API producer should be requested to run the old version in parallel with the new one.
- D. The API client code ONLY needs to be changed if it needs to take advantage of new features.

Answer: (SHOW ANSWER)

* Semantic Versioning is a 3-component number in the format of X.Y.Z, where :

X stands for a major version.

Y stands for a minor version:

Z stands for a patch.

So, SemVer is of the form Major.Minor.Patch Coming to our question , minor version of the API has been changed which is backward compatible. Hence there is no change required on API client end. If they want to make use of new featured that have been added as a part of minor version change they may need to change code at their end. Hence correct answer is The API client code ONLY needs to be changed if it needs to take advantage of new features.



Valid MCIA-Level-1 Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1 Exam! PrepPdf.com now offer the **newest MCIA-Level-1 exam dumps**, the PrepPdf.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here:

NEW QUESTION: 62

An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform. What feature of Anypoint Platform should be used to meet this requirement?

- A. Enterprise Security Module
- B. Client ID Secret
- C. Federated Identity Management
- D. Federated Client Management

Answer: (SHOW ANSWER)

Correct answer is Federated Identity Management 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: OpenID Connect: End user identity verification by an authorization server including SSO SAML 2.0: Web-based authorization including cross-domain SSO Where as Client Management is where 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 Reference:

<https://docs.mulesoft.com/access-management/external-identity>

NEW QUESTION: 63

An external REST client periodically sends an array of records in a single POST request to a Mule application API endpoint.

The Mule application must validate each record of the request against a JSON schema before sending it to a downstream system in the same order that it was received in the array Record processing will take place inside a router or scope that calls a child flow. The child flow has its own error handling defined. Any validation or communication failures should not prevent further processing of the remaining records.

To best address these requirements what is the most idiomatic(used for it intended purpose) router or scope to used in the parent flow, and what type of error handler should be used in the child flow?

- A. Until Successful router in the parent flow
On Error Propagate error handler in the child flow
- B. First Successful router in the parent flow
On Error Continue error handler in the child flow
- C. Parallel For Each scope in the parent flow
On Error Propagate error handler in the child flow
- D. For Each scope in the parent flow

On Error Continue error handler in the child flow

Answer: D (LEAVE A REPLY)

NEW QUESTION: 64

A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application.

What non-functional requirement can be expected to be negatively affected when implementing this change?

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

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 65

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

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

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

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 66

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

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 67

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

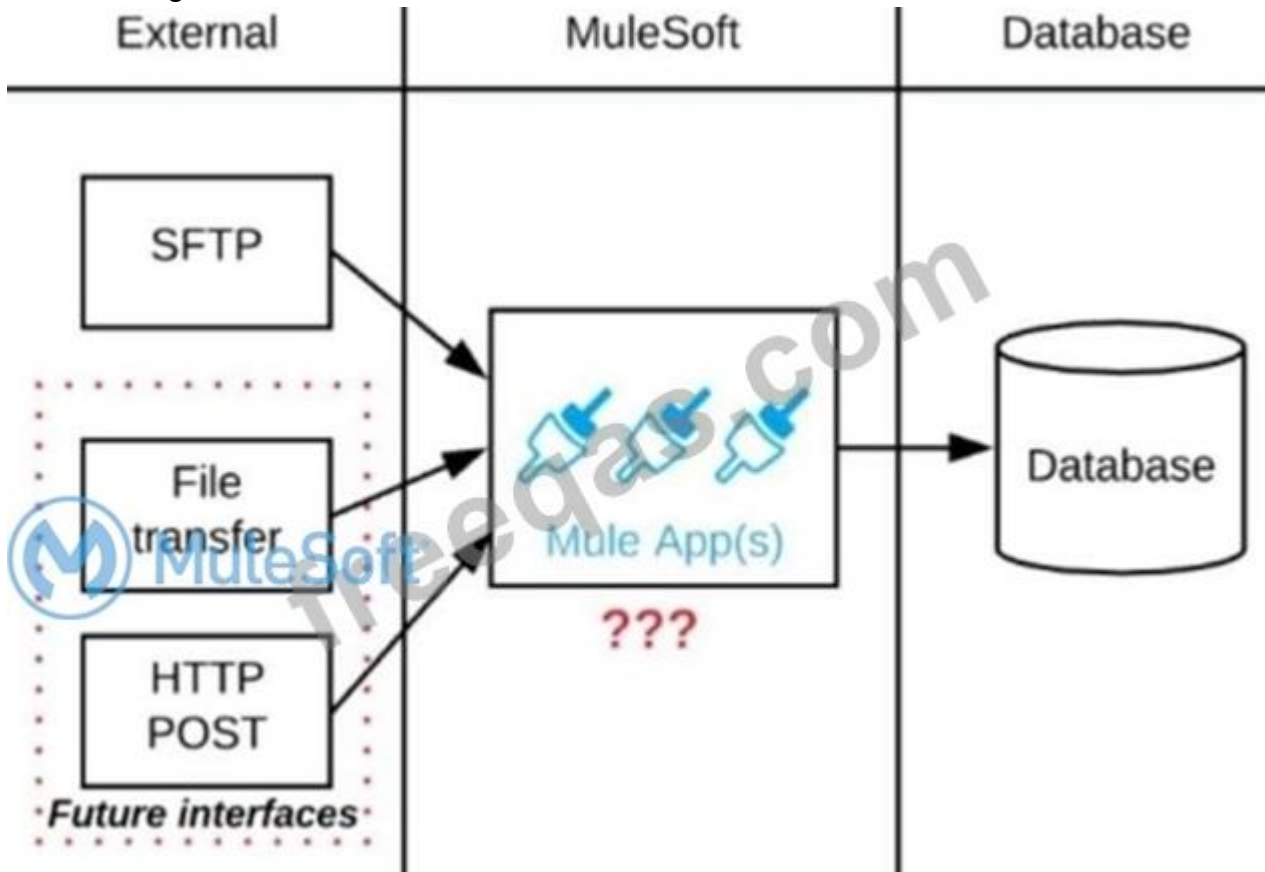
Answer: A ([LEAVE A REPLY](#))

Explanation/Reference:

NEW QUESTION: 68

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

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



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

Answer: A (LEAVE A REPLY)

NEW QUESTION: 69

A mule application uses an HTTP request operation to involve an external API.

The external API follows the HTTP specification for proper status code usage.

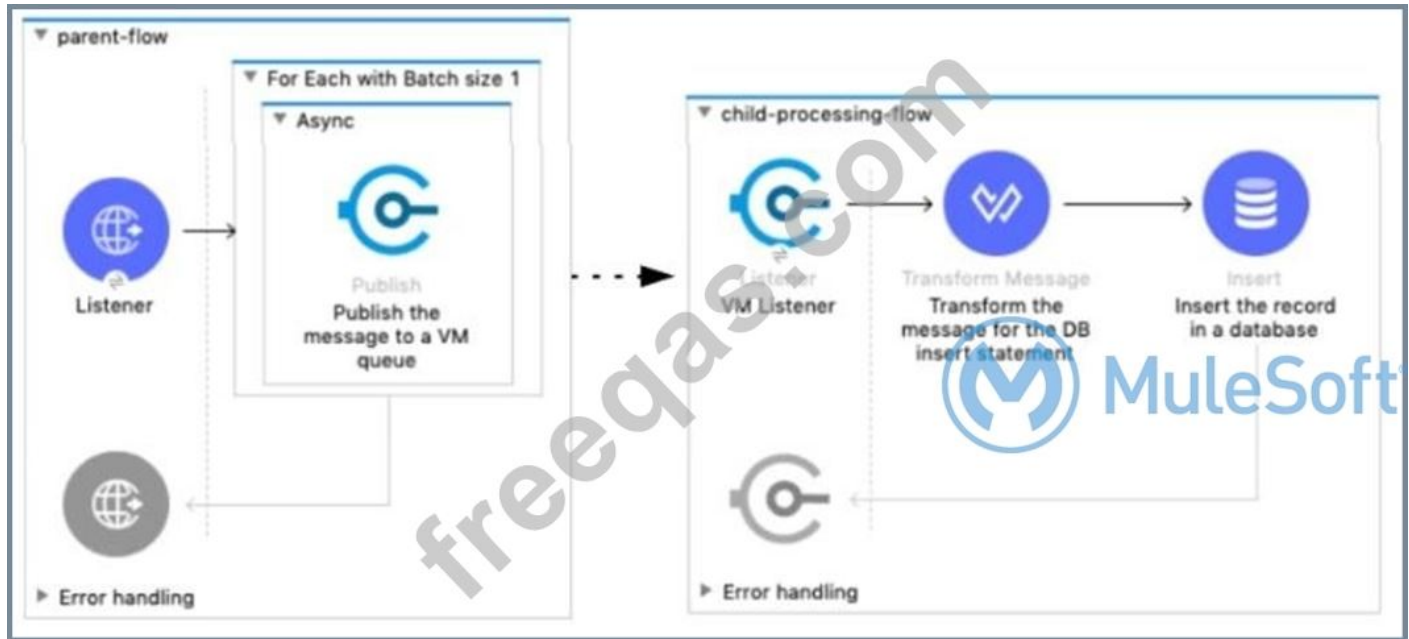
What is possible cause when a 3xx status code is returned to the HTTP Request operation from the external API?

- A. The request was ACCEPTED by the external API
- B. The request was not accepted by the external API
- C. The request was Redirected to a different URL by the external API
- D. The request was NOT RECEIVED by the external API

Answer: C (LEAVE A REPLY)

NEW QUESTION: 70

Refer to the exhibit.



A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items, then sends each item one at a time inside an Async scope to a VM queue.

A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing the item to a database.

This Mule application is deployed to four CloudHub workers with persistent queues enabled.

What message processing guarantees are provided by the VM queue and the CloudHub workers, and how are VM messages routed among the CloudHub workers for each invocation of the parent flow under normal operating conditions where all the CloudHub workers remain online?

- A. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub worker This one CloudHub worker processes ALL 200 item VM messages
- B. ALL Item VM messages are processed AT LEAST ONCE by the SAME CloudHub worker where the parent flow was invoked This one CloudHub worker processes ALL 200 item VM messages
- C. EACH item VM message is processed AT MOST ONCE by ONE CloudHub worker, with workers chosen in a deterministic round-robin fashion Each of the four CloudHub workers can be expected to process 1/4 of the Item VM messages (about 50 items)
- D. EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker Each of the four CloudHub workers can be expected to process some item VM messages

Answer: D (LEAVE A REPLY)

NEW QUESTION: 71

What requires configuration of both a key store and a trust store for an HTTP Listener?

- A. Encryption of both HTTP request and HTTP response bodies for all HTTP clients
- B. Encryption of both HTTP request header and HTTP request body for all HTTP clients
- C. Encryption of requests to both subdomains and API resource endpoints
fhttps://aDi.customer.com/ and https://customer.com/api)
- D. Support for TLS mutual (two-way) authentication with HTTP clients

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 72

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 transaction
3. Acknowledge the JMS message
- B. 1- Read the JMS message (NOT in an XA transaction)
2. Perform EACH DB insert in a SEPARATE DB transaction
3. Acknowledge the JMS message
- C. 1. Read the JMS message in an XA transaction
2. 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: ([SHOW ANSWER](#)**)**

NEW QUESTION: 73

An organization currently uses a multi-node Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications. The organization is planning to transition to a deployment model based on Docker containers in a Kubernetes cluster. The organization has already created a standard Docker image containing a Mule runtime and all required dependencies (including a JVM), but excluding the Mule application itself.

What is an expected outcome of this transition to container-based Mule application deployments?

- A. Required redesign of Mule applications to follow microservice architecture principles
- B. Required migration to the Docker and Kubernetes-based Anypoint Platform - Private Cloud Edition
- C. Required change to the URL endpoints used by clients to send requests to the Mule applications
- D. Guaranteed consistency of execution environments across all deployments of a Mule application

Answer: A (LEAVE A REPLY)

* Organization can continue using existing load balancer even if backend application changes are there. So option A is ruled out.

* As Mule runtime is within their datacenter, this model is RTF and not PCE. So option C is ruled out.

Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications -- This mean PCE or Hybrid not RTF - Also mentioned in Question is that - Mule runtime is hosting several Mule Application, so that also rules out RTF and as for hosting multiple Application it will have Domain project which need redesign to make it microservice architecture

----- Correct answer: Required redesign of Mule applications to follow microservice architecture principles

NEW QUESTION: 74

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 75

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)

* 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: 76

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 at 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 at Mule application deployments.

D. Keep the default logging configuration in Runtime Manager

Define the Splunk appender in EACH Mule application log4j2.xml file

Answer: B (LEAVE A REPLY)

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

Valid MCIA-Level-1 Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1 Exam! PrepPdf.com now offer the **newest MCIA-Level-1 exam dumps**, the PrepPdf.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here:
<https://www.preppdf.com/MuleSoft/MCIA-Level-1-prepaway-exam-dumps.html> (275 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 77

Mule applications need to be deployed to CloudHub so they can access on-premises database systems. These systems store sensitive and hence tightly protected data, so are not accessible over the internet.

What network architecture supports this requirement?

- A.** An Anypoint VPC connected to the on-premises network using an IPsec tunnel or AWS DirectConnect, plus matching firewall rules in the VPC and on-premises network
- B.** Static IP addresses for the Mule applications deployed to the CloudHub Shared Worker Cloud, plus matching firewall rules and IP whitelisting in the on-premises network
- C.** An Anypoint VPC with one Dedicated Load Balancer fronting each on-premises database system, plus matching IP whitelisting in the load balancer and firewall rules in the VPC and on-premises network
- D.** Relocation of the database systems to a DMZ in the on-premises network, with Mule applications deployed to the CloudHub Shared Worker Cloud connecting only to the DMZ

Answer: A (LEAVE A REPLY)

<https://docs.mulesoft.com/runtime-manager/vpc-connectivity-methods-concept>

NEW QUESTION: 78

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

Answer: (SHOW ANSWER)

NEW QUESTION: 79

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

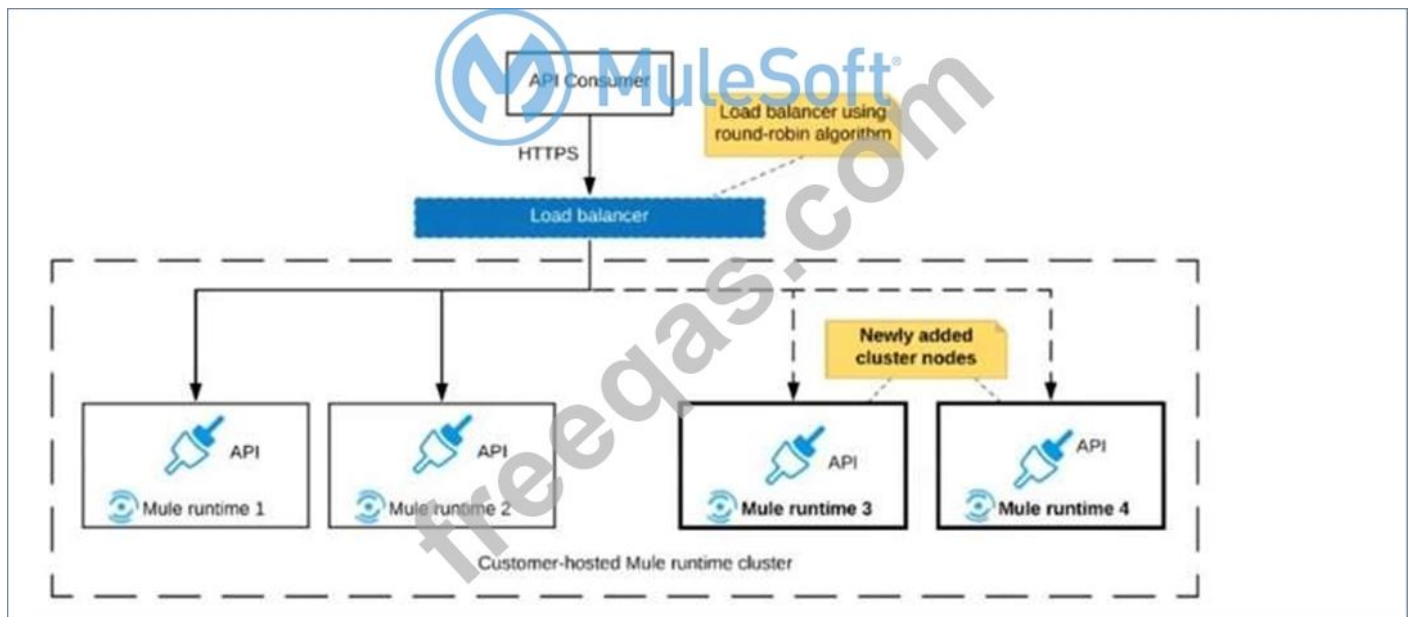
- A.** To receive messages into the Mule application, the JMS provider initiates a network connection to the JMS connector and pushes messages along this connection

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 80

Refer to the exhibit.



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

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

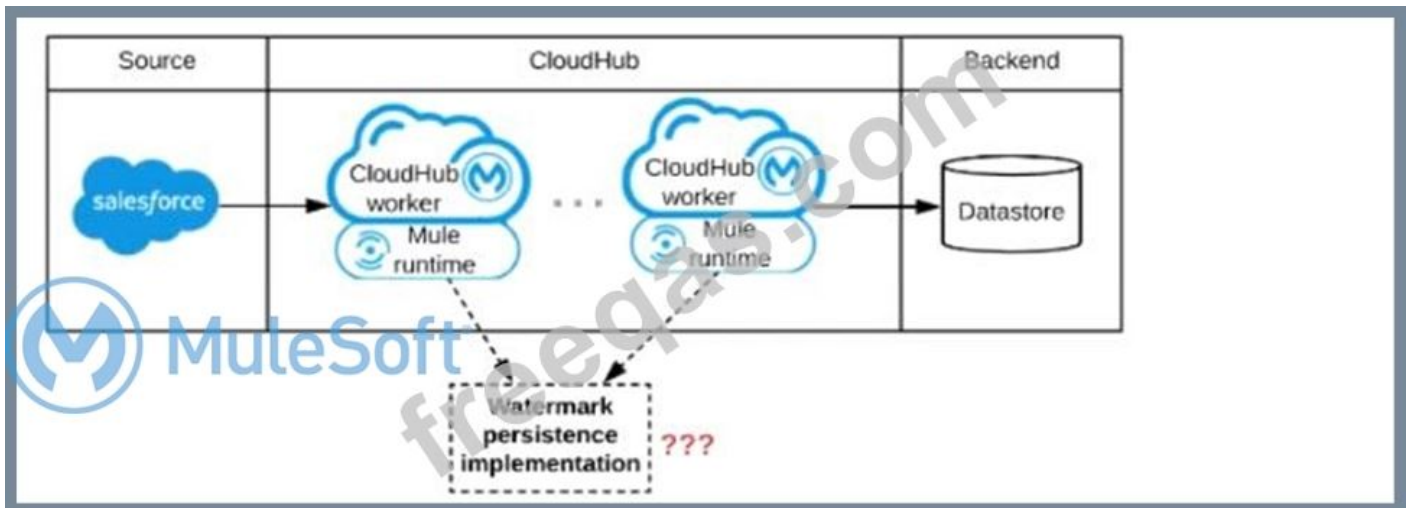
What average performance change is guaranteed to happen, assuming all cluster nodes are fully operational?

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

Answer: C (LEAVE A REPLY)

NEW QUESTION: 81

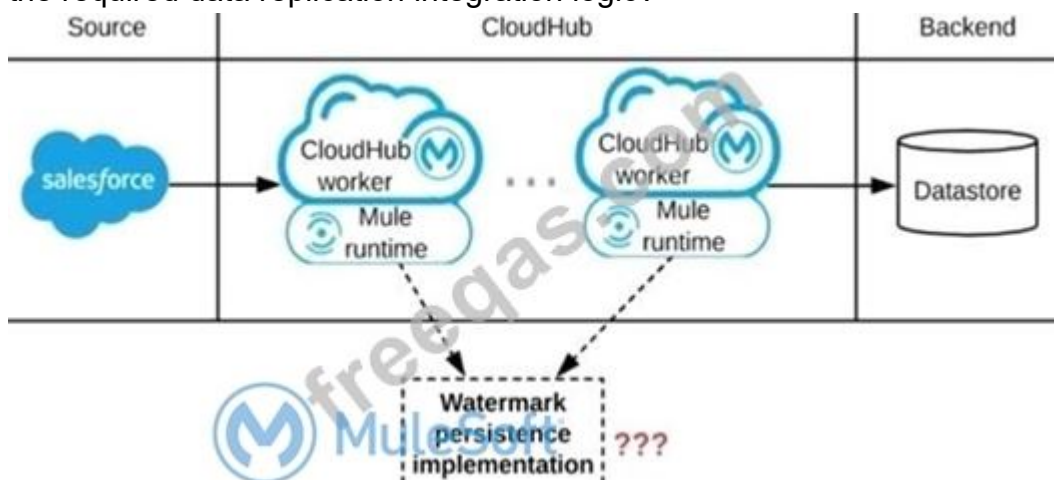
Refer to the exhibit.



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

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

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



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

Answer: (SHOW ANSWER)

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

* Mule provides two types of object stores:

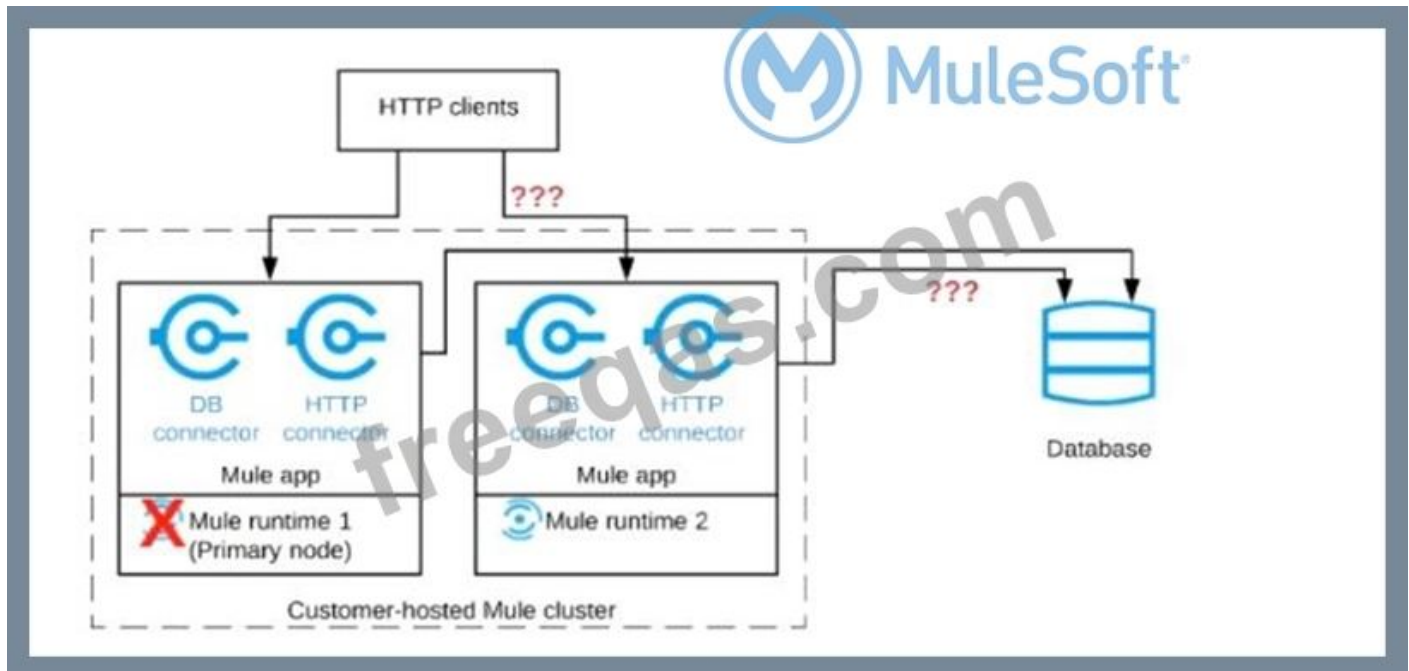
- 1) In-memory store - stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime.
- 2) Persistent store - Mule persists data when an object store is explicitly configured to be persistent.

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

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

NEW QUESTION: 82

Refer to the exhibit.



A Mule application is deployed to a cluster of two customer-hosted Mule runtimes. The Mule application has a flow that polls a database and another flow with an HTTP Listener.

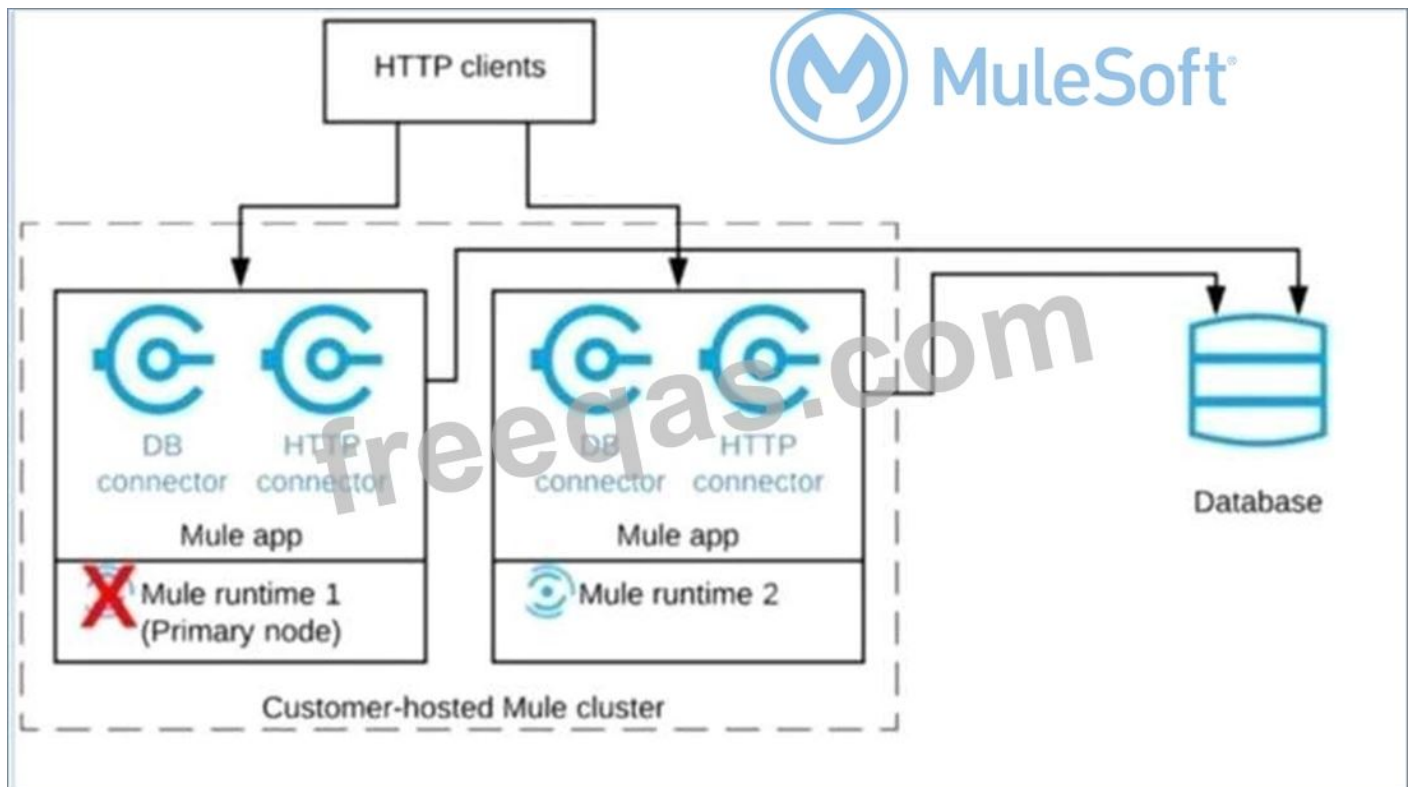
HTTP clients send HTTP requests directly to individual cluster nodes.

What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has failed, but before that node is restarted?

- A. Database polling continues Only HTTP requests sent to the remaining node continue to be accepted
- B. Database polling stops All HTTP requests continue to be accepted
- C. Database polling continues All HTTP requests continue to be accepted, but requests to the failed node incur increased latency
- D. Database polling stops All HTTP requests are rejected

Answer: (SHOW ANSWER)

Correct answer is Database polling continues Only HTTP requests sent to the remaining node continue to be accepted. : Architecture described in the question could be described as follows. When node 1 is down, DB polling will still continue via node 2. Also requests which are coming directly to node 2 will also be accepted and processed in BAU fashion. Only thing that won't work is when requests are sent to Node 1 HTTP connector. The flaw with this architecture is HTTP clients are sending HTTP requests directly to individual cluster nodes. By default, clustering Mule runtime engines ensures high system availability. If a Mule runtime engine node becomes unavailable due to failure or planned downtime, another node in the cluster can assume the workload and continue to process existing events and messages



NEW QUESTION: 83

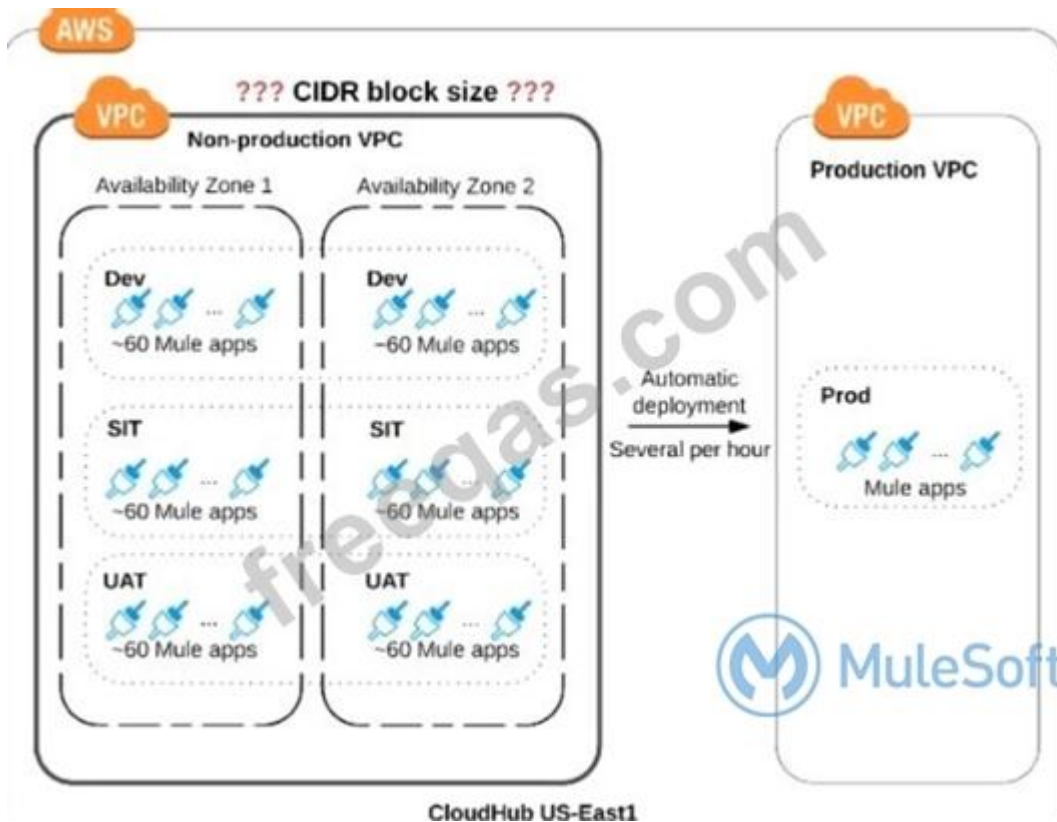
An automation engineer needs to write scripts to automate the steps of the API lifecycle, including steps to create, publish, deploy and manage APIs and their implementations in Anypoint Platform. What Anypoint Platform feature can be used to automate the execution of all these actions in scripts in the easiest way without needing to directly invoke the Anypoint Platform REST APIs?

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

Answer: A (LEAVE A REPLY)

NEW QUESTION: 84

Refer to the exhibit.



An organization is sizing an Anypoint VPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization currently has three non-production environments (DEV, SIT and UAT) that share this VPC. The AWS region of the VPC has two AZs.

The organization has a very mature DevOps approach which automatically progresses each application through all non-production environments before automatically deploying to production. This process results in several Mule application deployments per hour, using CloudHub's normal zero-downtime deployment feature.

What is a CIDR block for this VPC that results in the smallest usable private IP address range?

- A. 10.0.0.0/24 (256 IPs)
- B. 10.0.0.0/22 (1024 IPs)
- C. 10.0.0.0/25 (128 IPs)
- D. 10.0.0.0/26 (64 IPS)

Answer: B (LEAVE A REPLY)

NEW QUESTION: 85

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. New firewall rules must be configured to accommodate communication between API clients and the new nodes
- 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. A new load balancer must be provisioned to allow traffic to the new nodes in a round-robin fashion

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 86

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

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

A. Automated Policies in API Manager

B. Runtime Manager agent

C. The Mule Maven Plugin

D. Anypoint CLI

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 87

What Anypoint Connectors support transactions?

A. Database, JMS, VM

B. Database, JMS, HTTP

C. Database, JMS, VM, SFTP

D. Database, VM, File

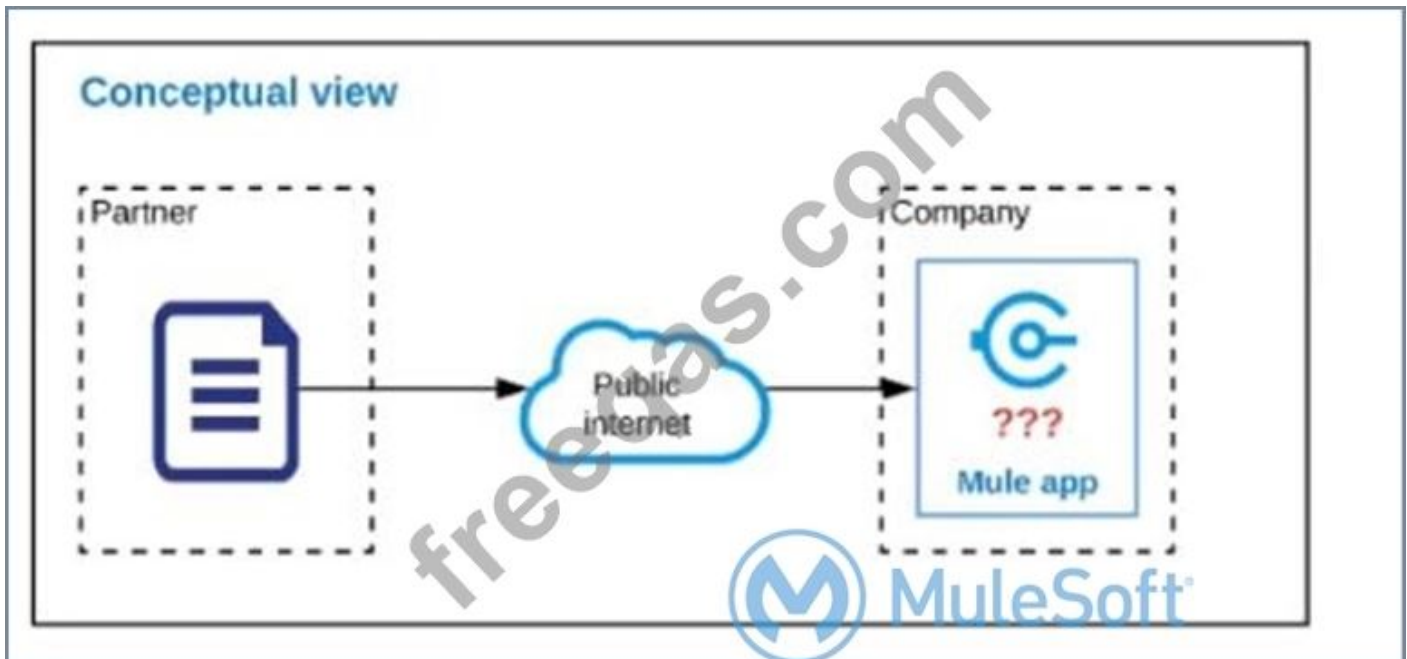
Answer: A ([LEAVE A REPLY](#))

Below Anypoint Connectors support transactions * JMS - Publish - Consume * VM - Publish - Consume * Database - All operations

NEW QUESTION: 88

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

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

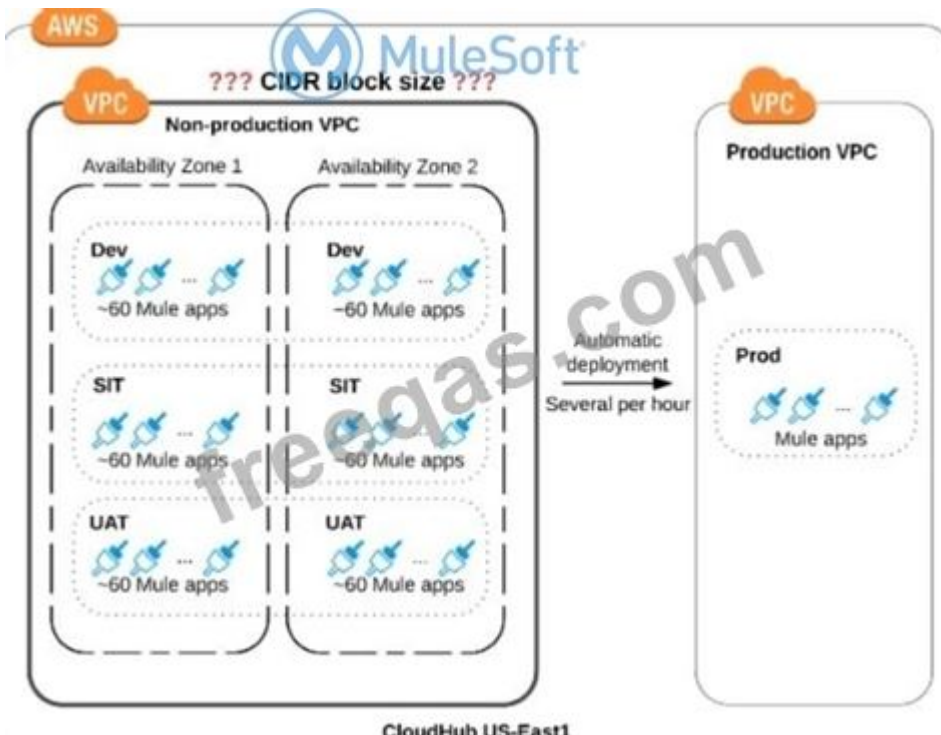


- A. VM connector
- B. Object Store connector
- C. File connector
- D. SFTP connector

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 89

Refer to the exhibit.



An organization is sizing an AnypointVPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization

currently has three non-production environments (DEV, SIT and UAT) that share this VPC. The AWS region of the VPC has two AZs.

The organization has a very mature DevOps approach which automatically progresses each application through all non-production environments before automatically deploying to production. This process results in several Mule application deployments per hour, using CloudHub's normal zero-downtime deployment feature.

What is a CIDR block for this VPC that results in the smallest usable private IP address range?

- A. 10.0.0.0/24 (256 IPs)
- B. 10.0.0.0/25 (128 IPs)
- C. 10.0.0.0/22 (1024 IPs)
- D. 10.0.0.0/26 (64 IPS)

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 90

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

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

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

Answer: ([SHOW ANSWER](#))

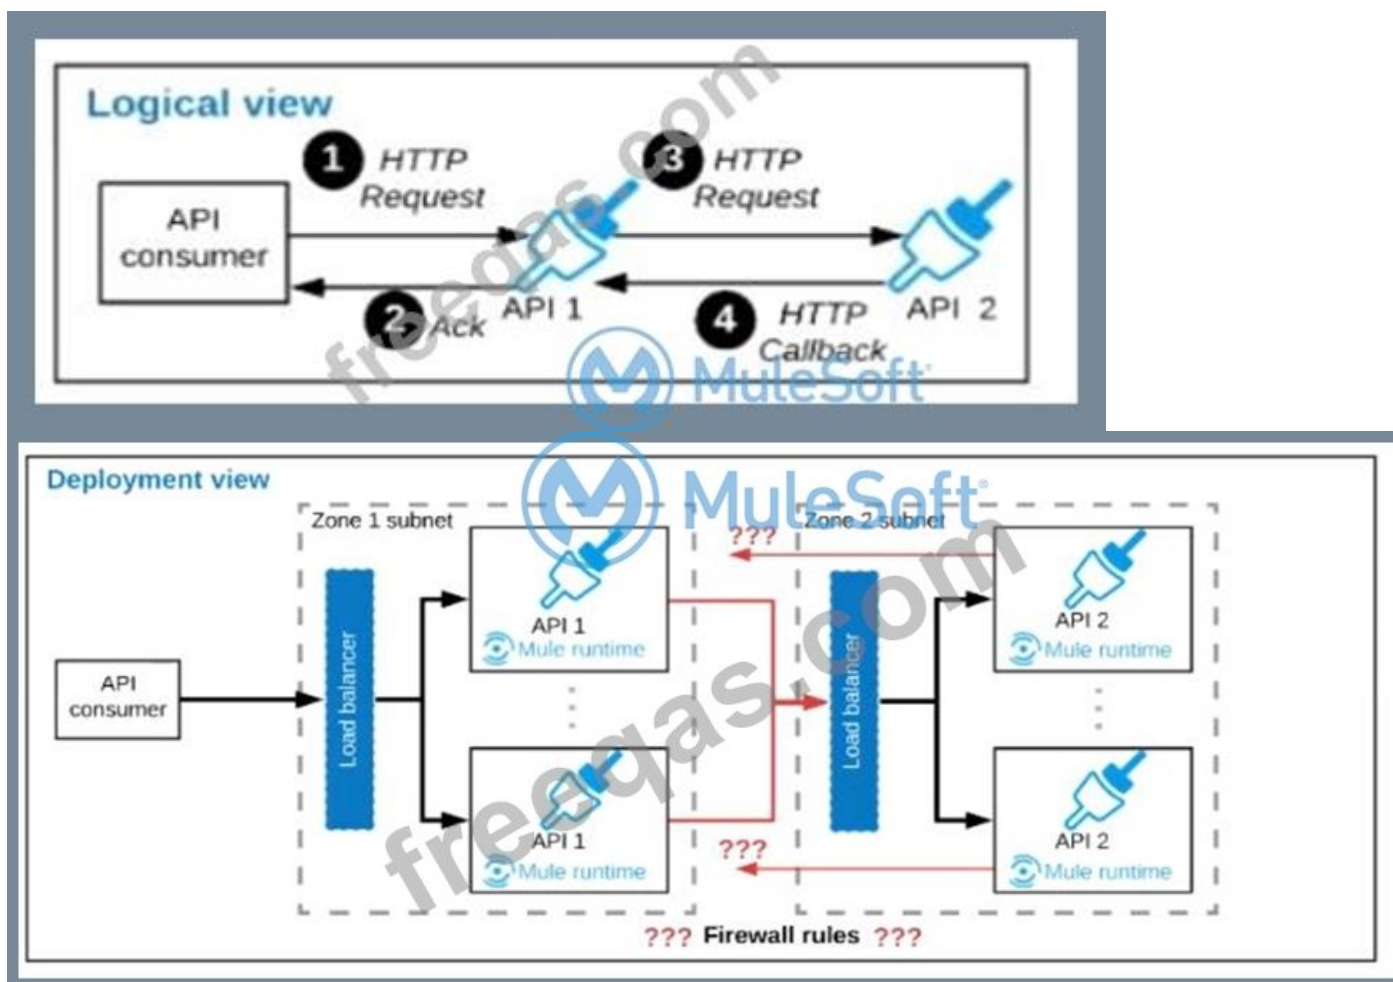
Correct answer is: Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC * Connecting to your Anypoint VPC extends your corporate network and allows CloudHub workers to access resources behind your corporate firewall.

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

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

NEW QUESTION: 91

Refer to the exhibit.



A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API 1 to notify about completion of the asynchronous process.

Each API is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.

In the network architecture, how must the firewall rules be configured to enable the above interaction between API 1 and API 2?

- A. To open direct two-way communication between the Mule runtimes of both APIs
- B. To allow communication between the load balancers used by each API
- C. To authorize the certificates used by both the APIs
- D. To enable communication from each API's Mule runtimes and network zone to the load balancer of the other API

Answer: ([SHOW ANSWER](#))

Valid MCIA-Level-1 Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1 Exam! PrepPdf.com now offer the **newest MCIA-Level-1 exam dumps**, the PrepPdf.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here:

NEW QUESTION: 92

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

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

Answer: D (LEAVE A REPLY)

We need to note few things about the scenario which will help us in reaching the correct solution.
Point 1 : The APIs are all publicly available and are associated with several mobile applications and web applications. This means Apply an IP blacklist policy is not viable option. as blacklisting IPs is limited to partial web traffic. It can't be useful for traffic from mobile application
Point 2 : The organization does NOT want to use any authentication or compliance policies for these APIs. This means we can not apply HTTPS mutual authentication scheme.

Header injection or removal will not help the purpose.

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

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

NEW QUESTION: 93

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

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 94

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)

* Correct answer is 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

NEW QUESTION: 95

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

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

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

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 96

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 (LEAVE A REPLY)

* 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.

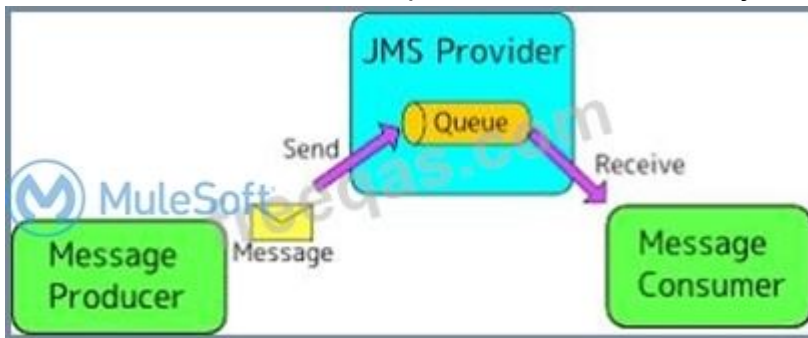
Correct answer: The JMS connector supports both sending and receiving of JMS

* 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/>

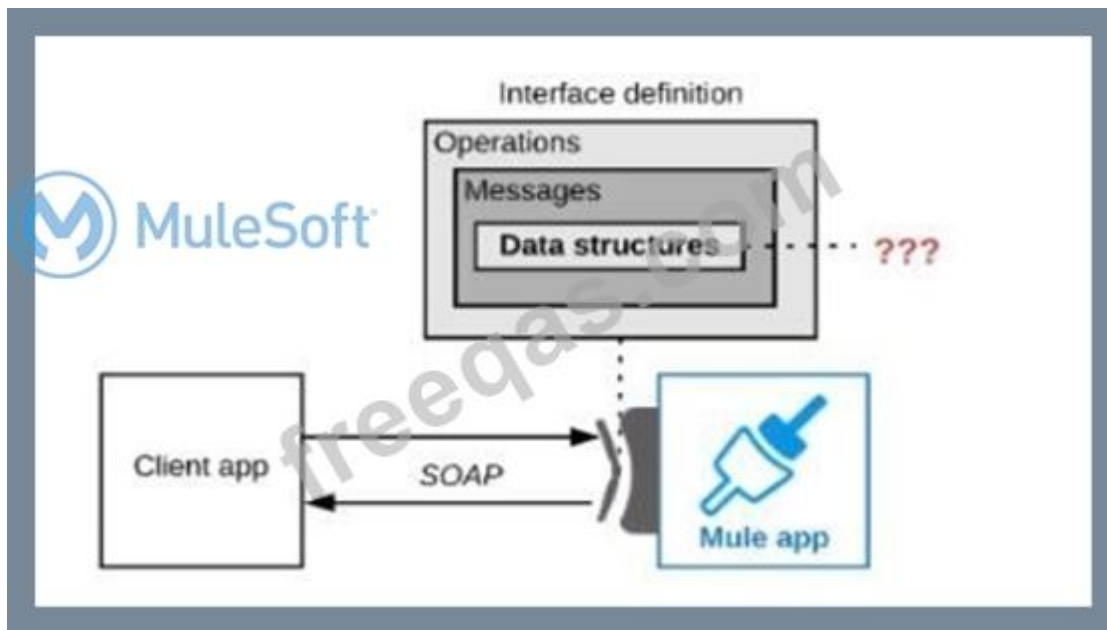


NEW QUESTION: 97

Refer to the exhibit.

A Mule application is being designed to expose a SOAP web service to its clients.

What language is typically used inside the web service's interface definition to define the data structures that the web service is expected to exchange with its clients?



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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 98

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

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 99

An Order microservice and a Fulfillment microservice are being designed to communicate with their clients through message-based integration (and NOT through API invocations).

The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mule application, the Fulfillment microservice.

The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message (a kind of event message). Each OrderFulfilled message can be consumed by any interested Mule application, and the Order microservice is one such Mule application.

What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

A. Order messages are sent to an Anypoint MQ exchange

OrderFulfilled messages are sent to an Anypoint MQ queue

Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices

B. Order messages are sent directly to the Fulfillment microservices

OrderFulfilled messages are sent directly to the Order microservice

The Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load each microservice

C. Order messages are sent to a JMS queue OrderFulfilled messages are sent to a JMS topic

Both microservices interact with the same JMS provider (message broker) Instance, which must therefore scale to support the load of both microservices

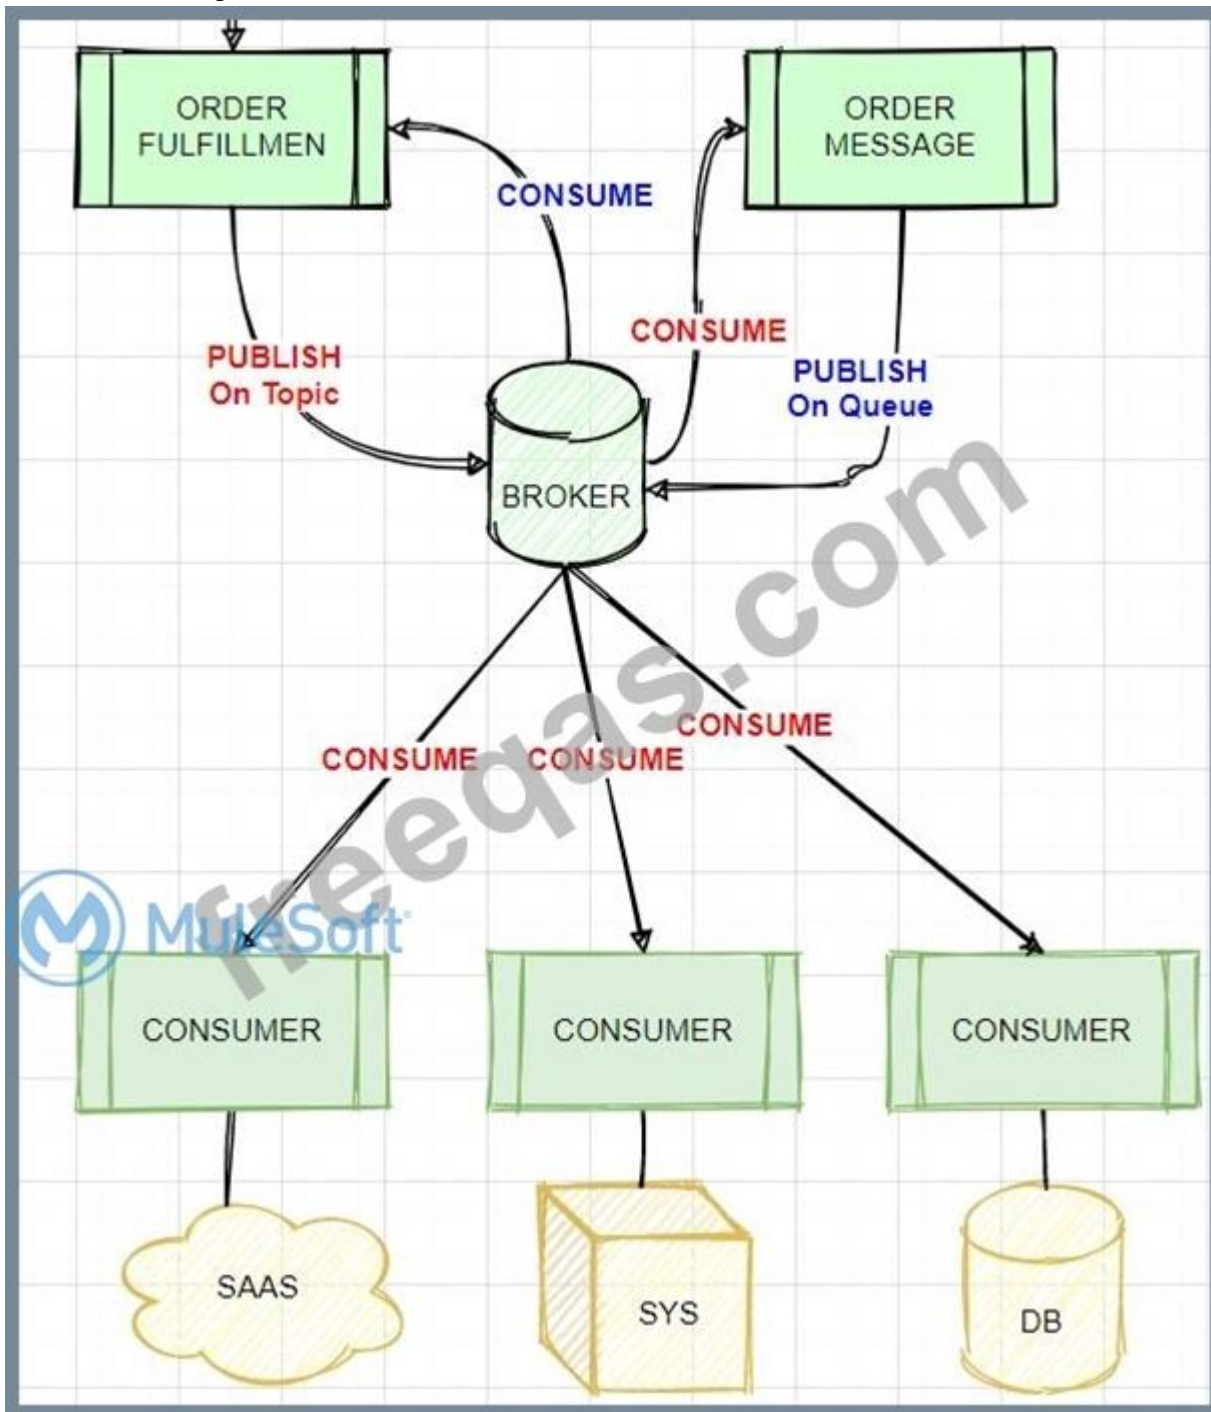
D. Order messages are sent to a JMS queue OrderFulfilled messages are sent to a JMS topic

The Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice

Answer: (SHOW ANSWER)

* If you need to scale a JMS provider/ message broker, - add nodes to scale it horizontally or - add memory to scale it vertically * Cons of adding another JMS provider/ message broker: - adds cost. - adds complexity to use two JMS brokers - adds Operational overhead if we use two brokers, say, ActiveMQ and IBM MQ * So Two options that mention to use two brokers are not best choice. * It's mentioned that "The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message. Each OrderFulfilled message can be consumed by any interested Mule application." - When you publish a message on a topic, it goes to all the subscribers who are interested - so zero to many subscribers will receive a copy of the message. - When you send a message on a queue, it will be received by exactly one consumer. * As we need multiple consumers to consume the message below option is not valid choice: "Order messages are sent to an Anypoint MQ exchange. OrderFulfilled messages are sent to an Anypoint MQ queue. Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices" * Order messages are only consumed by one Mule application, the Fulfillment microservice, so we will publish it on queue and OrderFulfilled message can be consumed by any interested Mule application so it need to be published on Topic using same broker. * Correct answer: Best choice in this scenario is: "Order messages are sent to a JMS queue. OrderFulfilled messages are sent to a JMS topic. Both microservices interact with the same JMS provider (message broker)

instance, which must therefore scale to support the load of both microservices" Tried to depict scenario in diagram:



NEW QUESTION: 100

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: ([SHOW ANSWER](#))

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: 101

An organization uses Mule runtimes which are managed by Anypoint Platform - Private Cloud Edition.

What MuleSoft component is responsible for feeding analytics data to non-MuleSoft analytics platforms?

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

Answer: A (LEAVE A REPLY)

Explanation/Reference: <https://docs.mulesoft.com/runtime-manager/sending-data-from-arm-to-external-analytics-software>

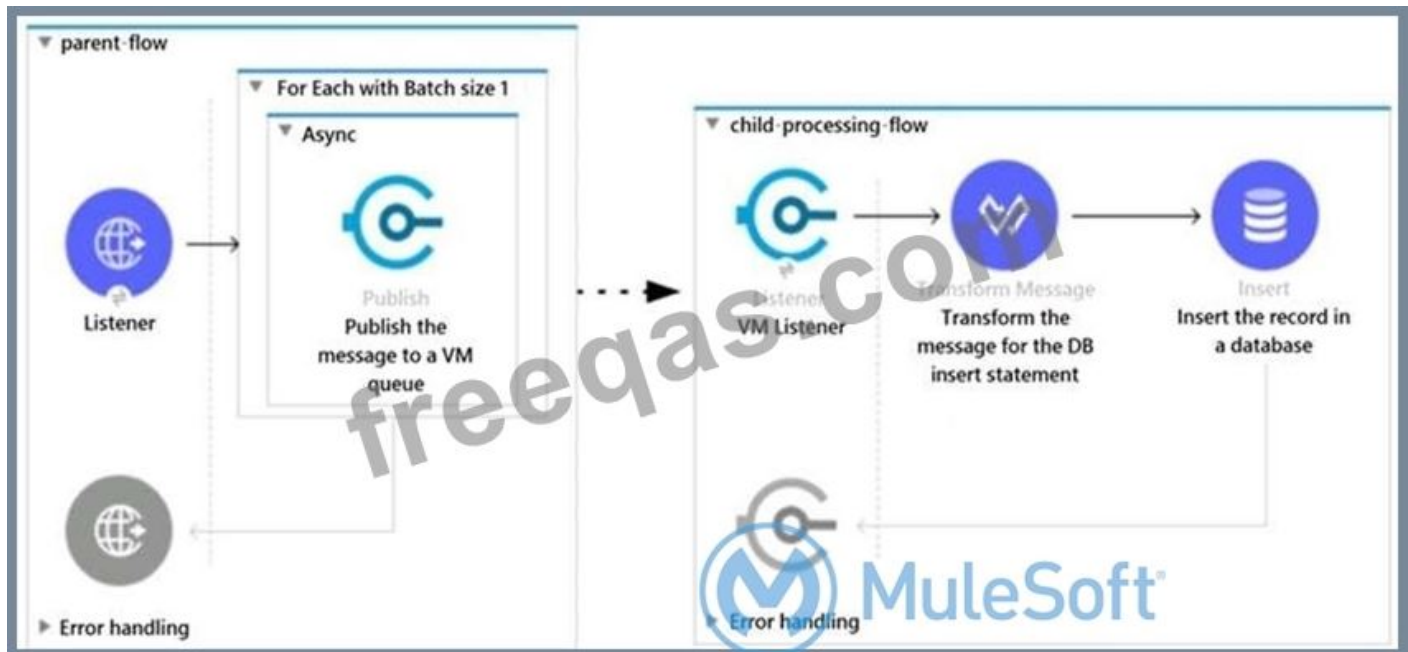
NEW QUESTION: 102

Refer to the exhibit. A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items, then sends each item one at a time inside an Async scope to a VM queue.

A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing the item to a database.

This Mule application is deployed to four CloudHub workers with persistent queues enabled.

What message processing guarantees are provided by the VM queue and the CloudHub workers, and how are VM messages routed among the CloudHub workers for each invocation of the parent flow under normal operating conditions where all the CloudHub workers remain online?



- A. EACH item VM message is processed AT MOST ONCE by ONE CloudHub worker, with workers chosen in a deterministic round-robin fashion Each of the four CloudHub workers can be expected to process 1/4 of the item VM messages (about 50 items)
- B. ALL item VM messages are processed AT LEAST ONCE by the SAME CloudHub worker where the parent flow was invoked This one CloudHub worker processes ALL 200 item VM messages
- C. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub worker This one CloudHub worker processes ALL 200 item VM messages
- D. EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker Each of the four CloudHub workers can be expected to process some item VM messages

Answer: B (LEAVE A REPLY)

NEW QUESTION: 103

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

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

Answer: (SHOW ANSWER)

NEW QUESTION: 104

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

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

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 105

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 106

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

If any processing step for a row fails, 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. For Each scope

On Error Continue scope

- C. VM connector

First Successful scope

On Error Propagate scope

- D. Async scope

On Error Propagate scope

Answer: ([SHOW ANSWER](#))

Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here:

<https://www.preppdf.com/MuleSoft/MCIA-Level-1-prepaway-exam-dumps.html> (275 Q&As

Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 107

What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

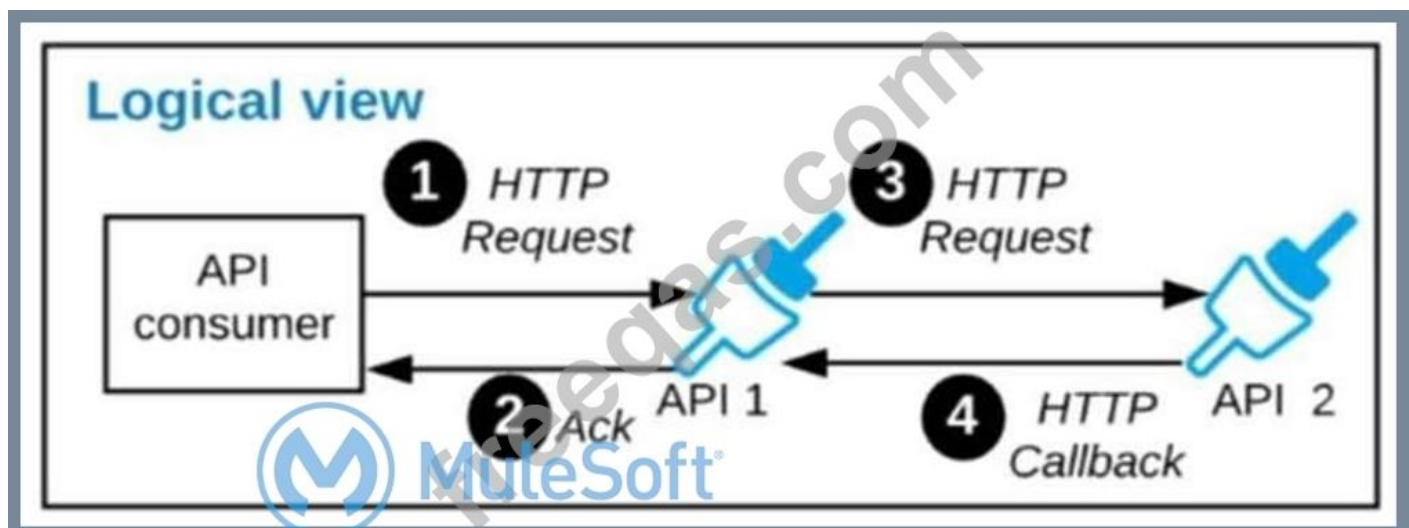
- A. Only a DLB allows the configuration of a custom TLS server certificate
- B. Only the SLB can forward HTTP traffic to the VPC-internal ports of the CloudHub workers
- C. Both a DLB and the SLB implement load balancing by sending HTTP requests to workers with the lowest workloads
- D. Both a DLB and the SLB allow the configuration of access control via IP whitelists

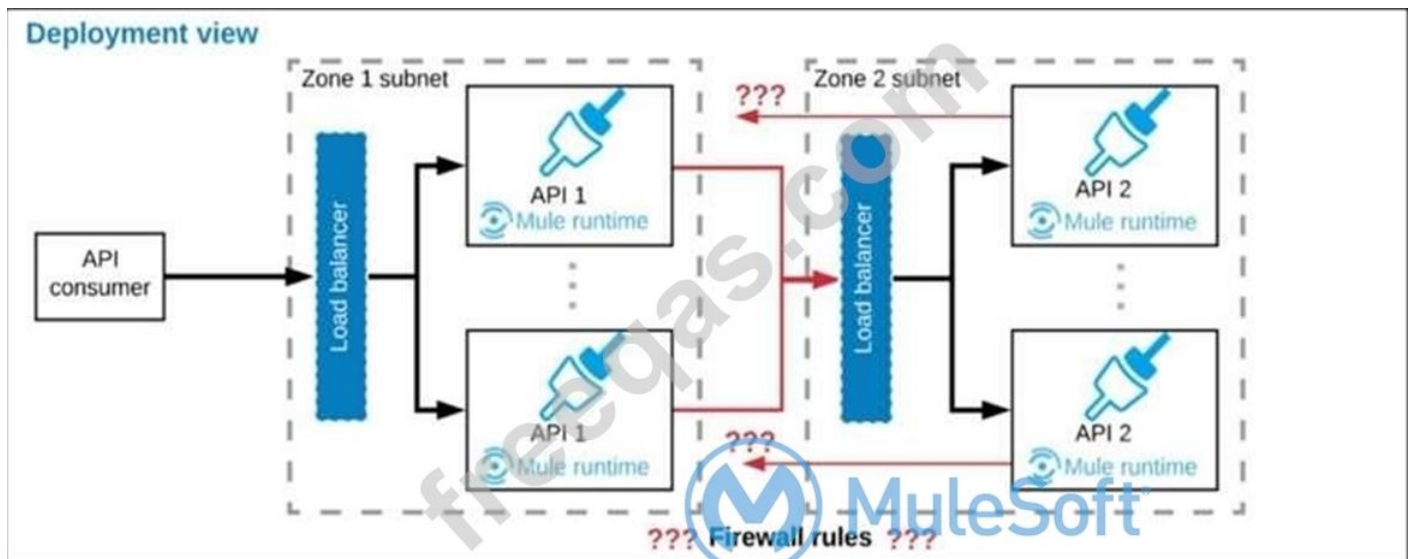
Answer: ([SHOW ANSWER](#))

NEW QUESTION: 108

Refer to the exhibit. A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API 1 to notify about completion of the asynchronous process. Each API is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.

In the network architecture, how must the firewall rules be configured to enable the above interaction between API 1 and API 2?





- A. To allow communication between the load balancers used by each API
- B. To enable communication from each API's Mule runtimes and network zone to the load balancer of the other API
- C. To authorize the certificates used by both the APIs
- D. To open direct two-way communication between the Mule runtimes of both APIs

Answer: D (LEAVE A REPLY)

NEW QUESTION: 109

An organization is designing an integration solution to replicate financial transaction data from a legacy system into a data warehouse (DWH).

The DWH must contain a daily snapshot of financial transactions, to be delivered as a CSV file. Daily transaction volume exceeds tens of millions of records, with significant spikes in volume during popular shopping periods.

What is the most appropriate integration style for an integration solution that meets the organization's current requirements?

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

Answer: (SHOW ANSWER)

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

Valid MCIA-Level-1 Dumps shared by PrepPdf.com for Helping Passing MCIA-Level-1 Exam! PrepPdf.com now offer the **newest MCIA-Level-1 exam dumps**, the PrepPdf.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com MCIA-Level-1 dumps with Test Engine here:
<https://www.preppdf.com/MuleSoft/MCIA-Level-1-prepaway-exam-dumps.html> (275 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)