

MuleSoft.MCIA-Level-1.v2024-03-12.q178

Exam Code:	MCIA-Level-1
Exam Name:	MuleSoft Certified Integration Architect - Level 1
Certification Provider:	MuleSoft
Free Question Number:	178
Version:	v2024-03-12
# of views:	1371
# of Questions views:	1780
https://www.freepdf.dumps.com/MuleSoft.MCIA-Level-1.v2024-03-12.q178.html	

NEW QUESTION: 1

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


How are the messages consumed by the Mule application?


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

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

NEW QUESTION: 2

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

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

NEW QUESTION: 3

A system administrator needs to determine when permissions were last changed for an Anypoint Platform user.

Which Anypoint Platform component should the administrator use to obtain this information?

- A. Anypoint Monitoring
- B. Mule Stack Traces
- C. Audit Logging
- D. Anypoint Studio

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

NEW QUESTION: 4

An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments.

What must the testing team do before they can start integration testing the API in the Staging environment?

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

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

NEW QUESTION: 5


What is a defining characteristic of an integration-Platform-as-a-Service (iPaaS)?

- A. Code-first
- B. A Cloud-based
- C. On-premises
- D. No-code

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

NEW QUESTION: 6

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

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

NEW QUESTION: 7

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

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

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

Answer: ([SHOW ANSWER](#)**)**


Correct answer is An HTTP Load Balancer.

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

Rest all options are distractors and their requirement depends on project use case.

NEW QUESTION: 8

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.** Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing
- C.** Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- D.** 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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 9

An integration Mule application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mule application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system. Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system).

The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed.

The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages.

What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

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

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

NEW QUESTION: 10

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

of testing would typically mock the results from each SELECT statement rather than actually execute it in the production database?

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

Answer: ([SHOW ANSWER](#))

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

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

Below are the typical characteristics of unit testing.

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

NEW QUESTION: 11

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 and P both use the Object Store connector to access a persistent object store
- B. 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
- C. Application C and P both use the Object Store connector to access the Anypoint Object Store v2
- D. 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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 12

An organization is designing multiple new applications to run on CloudHub in a single Anypoint VPC and that must share data using a common persistent Anypoint object store V2 (OSv2).

Which design gives these mule applications access to the same object store instance?

- A. AVM connector configured to directly access the persistence queue of the persistent object store
- B. Object store V2 can be shared across cloudhub applications with the configured osv2 connector
- C. The object store V2 rest API configured to access the persistent object store
- D. An Anypoint MQ connector configured to directly access the persistent object store

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

NEW QUESTION: 13

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

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

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

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

Correct answer is IBM: CICS CRM: SOAP

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

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

NEW QUESTION: 14

During a planning session with the executive leadership, the development team director presents plans for a new API to expose the data in the company's order database. An earlier effort to build an API on top of this data failed, so the director is recommending a design-first approach.

Which characteristics of a design-first approach will help make this API successful?

- A. Building MUnit tests so administrators can confirm code coverage percentage during deployment
- B. Developing a specification so consumers can test before the implementation is built
- C. Publishing the fully implemented API to Exchange so all developers can reuse the API

D. Adding global policies to the API so all developers automatically secure the implementation before coding anything

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 15

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: Required redesign of Mule applications to follow microservice

NEW QUESTION: 16

As a part of project , existing java implementation is being migrated to Mulesoft. Business is very tight on the budget and wish to complete the project in most economical way possible.

Canonical object model using java is already a part of existing implementation. Same object model is required by mule application for a business use case. What is the best way to achieve this?

- A.** Make use of Java module
- B.** Create similar model for Mule applications
- C.** Create a custom application to read Java code and make it available for Mule application

D. Use Anypoint exchange

Answer: A (LEAVE A REPLY)

Mule 4 is built to:

- * Minimize the need for custom code.
- * Avoid the need for you to know or understand Java.

However, some advanced uses cases require integration with custom Java code, such as:

- * Reuse of a library, such as a tax calculation library.
- * Reuse of a canonical object model that is standard in the organization.
- * Execution of custom logic using Java.

Mule ref doc : <https://docs.mulesoft.com/java-module/1.2/>

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here: https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

An integration team follows MuleSoft's recommended approach to full lifecycle API development. Which activity should this team perform during the API implementation phase?

- A. Validate the API specification
- B. Use the API specification to build the MuleSoft application
- C. Use the API specification to monitor the MuleSoft application
- D. Design the API specification

Answer: B (LEAVE A REPLY)

NEW QUESTION: 18

What is an advantage that Anypoint Platform offers by providing universal API management and Integration-Platform-as-a-Service (iPaaS) capabilities in a unified platform?

- A. Ability to use a single iPaaS to manage and integrate all API gateways
- B. Ability to use a single connector to manage and integrate all APIs
- C. Ability to use a single control plane for both full-lifecycle API management and integration
- D. Ability to use a single iPaaS to manage all API developer portals

Answer: C (LEAVE A REPLY)

Anypoint Platform offers universal API management and Integration-Platform-as-a-Service (iPaaS) capabilities in a unified platform, meaning that it provides a single control plane to manage both full-lifecycle API management and integration. This allows organizations to easily manage their APIs and integrations, as well as deploy APIs and integrations quickly and

efficiently. According to the MuleSoft Certified Integration Architect - Level 1 Course Book, "Anypoint Platform provides a unified platform for managing, deploying, and monitoring both API and integration solutions, allowing organizations to quickly and easily build and manage their APIs and integrations."

NEW QUESTION: 19

A Mule application is being designed To receive nightly a CSV file containing millions of records from an external vendor over SFTP, The records from the file need to be validated, transformed. And then written to a database. Records can be inserted into the database in any order. In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 20

An integration team uses Anypoint Platform and follows MuleSoft's recommended approach to full lifecycle API development.

Which step should the team's API designer take before the API developers implement the API Specification?

- A.** Generate test cases using MUnit so the API developers can observe the results of running the API
- B.** Publish the API specification to Exchange and solicit feedback from the API's consumers
- C.** Use the scaffolding capability of Anypoint Studio to create an API portal based on the API specification
- D.** Use API Manager to version the API specification

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 21

An external web UI application currently accepts occasional HTTP requests from client web browsers to change (insert, update, or delete) inventory pricing information in an inventory system's database. Each inventory pricing change must be transformed and then synchronized with multiple customer experience systems in near real-time (in under 10 seconds). New customer experience systems are expected to be added in the future.

The database is used heavily and limits the number of SELECT queries that can be made to the database to 10 requests per hour per user.

What is the most scalable, idiomatic (used for its intended purpose), decoupled, reusable, and maintainable integration mechanism available to synchronize each inventory pricing change with the various customer experience systems in near real-time?

- A.** Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the ID attribute set to an appropriate database column. In the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queue. Write other Mule applications to subscribe to the Anypoint MQ queue, transform each received message, and then update the Mule application's corresponding customer experience system(s).
- B.** Add a trigger to the inventory-pricing database table so that for each change to the inventory pricing database, a stored procedure is called that makes a REST call to a Mule application. Write the Mule application to publish each Mule event as a message to an Anypoint MQ exchange. Write other Mule applications to subscribe to the Anypoint MQ exchange, transform each received message, and then update the Mule application's corresponding customer experience system(s).
- C.** Replace the external web UI application with a Mule application to accept HTTP requests from client web browsers. In the same Mule application, use a Batch Job scope to test if the database request will succeed, aggregate pricing changes within a short time window, and then update both the inventory pricing database and each customer experience system using a Parallel For Each scope.
- D.** Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the watermark attribute set to an appropriate database column. In the same flow, use a Scatter-Gather to call each customer experience system's REST API with transformed inventory-pricing records.

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

NEW QUESTION: 22

An insurance company has an existing API which is currently used by customers. API is deployed to customer hosted Mule runtime cluster. The load balancer that is used to access any APIs on the mule cluster is only configured to point to applications hosted on the server at port 443. Mule application team of a company attempted to deploy a second API using port 443 but the application will not start and checking logs shows an error indicating the address is already in use.


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

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

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

NEW QUESTION: 23

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. 100% increase in the throughput of the API
- B. 50% reduction in the response time 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: ([SHOW ANSWER](#))

NEW QUESTION: 24

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 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
- 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.** 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
- D.** Use a Scatter-Gather within the For Each scope to ensure response message order Configure the Scatter-Gather with a persistent object store

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 25

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organization has all of its Mule applications running on its own premises and uses an on-premise 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 straightforward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

- A.** Deploy the Mule applications to Cloudhub
Create CNAME record for `api.acme.com` in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancer 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 dedicated load balancer(DLB) Update the CNAME record for `api.acme.com` in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB) Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications
- C.** Deploy the Mule applications to Cloudhub
Update the CNAME record for `api.acme.com` in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB) Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.
- D.** Deploy the Mule applications to Cloudhub
Update the CNAME record for `api.acme.com` in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB) Apply mapping rules in the DLB to map URLs to their corresponding Mule applications

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

NEW QUESTION: 26

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

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

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

NEW QUESTION: 27

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

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

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

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

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

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

NEW QUESTION: 28

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

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

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

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 29

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

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

NEW QUESTION: 30

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: 31

A company is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an HTTP5 POST and must be acknowledged immediately.

Once acknowledged the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to the rejections from the back-end system will need to be processed manually (outside the banking system).

The mule application will be deployed to a customer hosted runtime and will be able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization's firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.

Which combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

- A.** One or more on-Error scopes to assist calling the back-end system one or more ActiveMQ long-retry queues A persistent dead-letter Object store configuration in the CloudHub object store service
- B.** A batch job scope to call the back in system An Untill successful scope containing Object Store components for long retries. A dead-letter object store configured in the Mule application
- C.** An Until Successful scope to call the back-end system One or more ActiveMQ long-retry queues One or more ActiveMQ dead-letter queues for manual processing
- D.** One or more On Error scopes to assist calling the back-end system An Untill successful scope containing VM components for long retries A persistent dead-letter VM queue configure in Cloud hub

Answer: ([SHOW ANSWER](#))

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here:

https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 32

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: ([SHOW ANSWER](#))

* 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: 33

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 34

A developer is examining the responses from a RESTful web service that is compliant with the Hypertext Transfer Protocol (HTTP/1.1) as defined by the Internet Engineering Task Force (IETF).

In this HTTP/1.1-compliant web service, which class of HTTP response status codes should be specified to indicate when client requests are successfully received, understood, and accepted by the web service?

- A.** 2xx
- B.** 4xx
- C.** 5xx
- D.** 3xx

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 35


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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 36


Refer to the exhibit.



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

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

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



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

Answer: B (LEAVE A REPLY)

* 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: 37

Cloud Hub is an example of which cloud computing service model?

- A. Monitoring as a Service (MaaS)
- B. Infrastructure as a Service (IaaS)
- C. Platform as a Service (PaaS)
- D. Software as a Service (SaaS)

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

NEW QUESTION: 38


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

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

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

NEW QUESTION: 39

Refer to the exhibit.



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

Step 1: The JMS Correlation ID header is read from the received JMS message.

Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.

Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.

Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.

Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.

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

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

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

NEW QUESTION: 40

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

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

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

Explanation

Correct answer is Anypoint Runtime Manager

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

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


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

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

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

You can use the CloudHub custom log appender to integrate with your logging system.


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



Additional Info:

It can be achieved in 3 steps:

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




NEW QUESTION: 41

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 LEAST ONCE by the SAME CloudHub worker where the parent flow was invoked This one CloudHub worker processes ALL 200 item VM messages

B. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub worker This one CloudHub worker processes ALL 200 item VM messages

C. 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

D. 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)

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

NEW QUESTION: 42

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. The Mule runtimes


B. Anypoint Runtime Manager

- C. Anypoint Exchange
- D. Anypoint API Manager

Answer: B (LEAVE A REPLY)

NEW QUESTION: 43

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)

NEW QUESTION: 44

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: 45

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 46

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

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

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

- C. runtimeFabric deployment should be added to POM.xml file in all the mule applications and CI/CD script should be modified as per the RTF configurations
- D. runtimeFabric profile should be added mule configuration files in the mule applications and CI/CD script should be modified as per the RTF configurations

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

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here: https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 47


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

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 48

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 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
- 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 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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 49

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

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

NEW QUESTION: 50

An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub. For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub's default zero-downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers. This is expected to result in several Mule application deployments per hour.

- A.** 10.0.0.0/23(512 IPs)
- B.** 10.0.0.0/21(2048 IPs)
- C.** 10.0.0.0/24(256 IPs)
- D.** 10.0.0.0/22(1024 IPs)

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

NEW QUESTION: 51

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

NEW QUESTION: 52

An organization needs to procure an enterprise software system to increase cross-selling opportunities and better track prospect data.

Which category of enterprise software has these core capabilities, when used for its typical and intended purpose?

- A.** Supply Chain Management (SCM)
- B.** Customer Relationship Management (CRM)
- C.** IT Service Management (ITSM)
- D.** Business-to-Business (B2B)

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

NEW QUESTION: 53

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

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 54

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

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

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

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

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

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

NEW QUESTION: 55

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

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

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

* Setting a transaction timeout for the Bitronix transaction manager

* Set the transaction timeout either

- In wrapper.conf

- In CloudHub in the Properties tab of the Mule application deployment

* The default is 60 secs. It is defined as

mule.bitronix.transactiontimeout = 120

* This property defines the timeout for each transaction created for this manager.

If the transaction has not terminated before the timeout expires it will be automatically rolled back.

Additional Info around Transaction Management:

Bitronix is available as the XA transaction manager for Mule applications

* To use Bitronix, declare it as a global configuration element in the Mule application

```
<bt:transaction-manager />
```

* Each Mule runtime can have only one instance of a Bitronix transaction manager, which is shared by all Mule applications

* For customer-hosted deployments, define the XA transaction manager in a Mule domain

- Then share this global element among all Mule applications in the Mule runtime

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

NEW QUESTION: 56

What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

- A. To change log4j2 log levels in Anypoint Runtime Manager without having to restart the Mule application
- B. To log certain messages to a custom log category
- C. To send Mule application log entries to Splunk
- D. To directly reference one shared and customized log4j2.xml file from multiple Mule applications

Answer: B (LEAVE A REPLY)

Explanation/Reference: <https://docs.mulesoft.com/runtime-manager/viewing-log-data>

NEW QUESTION: 57

An organization is migrating all its Mule applications to Runtime Fabric (RTF). None of the Mule applications use Mule domain projects.

Currently, all the Mule applications have been manually deployed to a server group among several customer-hosted Mule runtimes. Port conflicts between these Mule application deployments are currently managed by the DevOps team who carefully manage Mule application properties files.

When the Mule applications are migrated from the current customer-hosted server group to Runtime Fabric (RTF), do the Mule applications need to be rewritten, and what DevOps port configuration responsibilities change or stay the same?

A. NO, the Mule applications do NOT need to be rewritten

DevOps MUST STILL manage port conflicts

B. YES, the Mule applications MUST be rewritten

DevOps MUST STILL manage port conflicts

C. YES, the Mule applications MUST be rewritten

DevOps NO LONGER needs to manage port conflicts between the Mule applications

D. NO, the Mule applications do NOT need to be rewritten

DevOps NO LONGER needs to manage port conflicts between the Mule applications

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

NEW QUESTION: 58

Refer to the exhibit. A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to receive JMS messages from a JMS queue.

To process each received JMS message, the following steps are performed in a flow:

Step 1: The JMS Correlation ID header is read from the received JMS message.


Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.

Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.

Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.

Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.

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



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

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

NEW QUESTION: 59

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

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

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


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

:

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

NEW QUESTION: 60

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 Anypoint MQ Queue
- D. Persistent Object Store

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

NEW QUESTION: 61

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

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

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

- * Setting a transaction timeout for the Bitronix transaction manager
- * Set the transaction timeout either
 - In wrapper.conf
 - In CloudHub in the Properties tab of the Mule application deployment

* The default is 60 secs. It is defined as
mule.bitronix.transactiontimeout = 120

* This property defines the timeout for each transaction created for this manager.

If the transaction has not terminated before the timeout expires it will be automatically rolled back.

Additional Info around Transaction Management:

Bitronix is available as the XA transaction manager for Mule applications

* To use Bitronix, declare it as a global configuration element in the Mule application

```
<bti:transaction-manager />
```

* Each Mule runtime can have only one instance of a Bitronix transaction manager, which is shared by all Mule applications

* For customer-hosted deployments, define the XA transaction manager in a Mule domain

- Then share this global element among all Mule applications in the Mule runtime

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

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here: https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 62

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 Cache scope to filter and store each record from the received file by the order ID
- B. Configure a Database connector to filter and store each record by the order ID
- C. Configure a watermark In an On New or Updated File event source to filter unique records by the order ID
- D. Configure an Idempotent Message Validator component to filter each record by the order ID

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

NEW QUESTION: 63

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 bestrecover from the expected outages?

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

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

NEW QUESTION: 64

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

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

NEW QUESTION: 65

An organization's IT team follows an API-led connectivity approach and must use Anypoint Platform to implement a System API that securely accesses customer data. The organization uses Salesforce as the system of record for all customer data, and its most important objective is to reduce the overall development time to release the System API.

The team's integration architect has identified four different approaches to access the customer data from within the implementation of the System API by using different Anypoint Connectors that all meet the technical requirements of the project.

- A. Use the Anypoint Connector for FTP to download a file containing a recent near-real time extract of the customer data
- B. Use the Anypoint Connector for HTTP to connect to the Salesforce APIs to directly access the customer data
- C. Use the Anypoint Connector for Salesforce to connect to the Salesforce APIs to directly access the customer data
- D. Use the Anypoint Connector for Database to connect to a MySQL database to access a copy of the customer data

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

NEW QUESTION: 66

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

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

Answer: ([SHOW ANSWER](#)**)**

An SLA for the upstream API CANNOT be provided.

NEW QUESTION: 67


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: 68

A REST API is being designed to implement a Mule application.

What standard interface definition language can be used to define REST APIs?

- A. Web Service Definition Language(WSDL)
- B. YAML
- C. OpenAPI Specification (OAS)
- D. AsyncAPI Specification

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

NEW QUESTION: 69

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


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

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

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

NEW QUESTION: 70

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. 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. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 71

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.** Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale
- 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.** 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

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

NEW QUESTION: 72

A mule application must periodically process a large dataset which varies from 6 GB to 8 GB from a back-end database and write transformed data to an FTPS server using a properly configured batch job scope.

The performance requirements of an application are approved to run in the cloud hub 0.2 vCore with 8 GB storage capacity and currency requirements are met.

How can the high rate of records be effectively managed in this application?

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

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

NEW QUESTION: 73

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


What is also required to carry out this change?

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

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

- * Clustering is a group of servers or mule runtime which acts as a single unit.
- * Mulesoft Enterprise Edition supports scalable clustering to provide high availability for the Mulesoft application.
- * In simple terms, virtual servers composed of multiple nodes and they communicate and share information through a distributed shared memory grid.
- * By default, Mulesoft ensures the High availability of applications if clustering implemented.
- * Let's consider the scenario one of the nodes in cluster crashed or goes down and under maintenance. In such cases, Mulesoft will ensure that requests are processed by other nodes in the cluster. Mulesoft clustering also ensures that the request is load balanced between all the nodes in a cluster.
- * Clustering is only supported by on-premise Mule runtime and it is not supported in Cloudhub. Correct answer is External monitoring tools or log aggregators must be configured to recognize the new nodes
- * Rest of the options are automatically taken care of when a new node is added in cluster.

Reference:



NEW QUESTION: 74

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. VM connector first Successful scope On Error Propagate scope
- C. For Each scope On Error Continue scope
- D. Async scope On Error Propagate scope

Answer: C (LEAVE A REPLY)

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


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

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

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

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

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



NEW QUESTION: 75

An XA transaction is being configured that involves a JMS connector listening for incoming JMS messages.

What is the meaning of the timeout attribute of the XA transaction, and what happens after the timeout expires?

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

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

NEW QUESTION: 76

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

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here:

https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 77

Refer to the exhibit.

```
traits:
  error-responses: traits/error-responses.raml
  jwt-required:
    headers:
      x-jwt:
        type: string
        description: JWT token string
```

What is the type data format shown in the exhibit?

- A. JSON

- B. XML
- C. CSV
- D. YAML

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

NEW QUESTION: 78

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: ([SHOW ANSWER](#)**)**

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

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

Mule Ref Doc : <https://docs.mulesoft.com/api-manager/2.x/getting-started-proxy>

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

What is a key difference between synchronous and asynchronous logging from Mule applications?



- A. Synchronous logging writes log messages in a single logging thread but does not block the Mule event being processed by the next event processor
- B. Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event
- C. Asynchronous logging produces more reliable audit trails with more accurate timestamps
- D. Synchronous logging within an ongoing transaction writes log messages in the same thread that processes the current Mule event

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

Types of logging:

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

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



B) Asynchronous:

- * The logging operation occurs in a separate thread, so the actual processing of your message won't be delayed to wait for the logging to complete

- * Substantial improvement in throughput and latency of message processing
- * Mule runtime engine (Mule) 4 uses Log4j 2 asynchronous logging by default
- * The disadvantage of asynchronous logging is error handling.
- * If the logger fails to write to disk, the thread doing the processing won't be aware of any issues writing to the disk, so you won't be able to rollback anything. Because the actual writing of the log gets deferred, there's a chance that log messages might never make it to disk and get lost, if Mule were to crash before the buffers are flushed.

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

NEW QUESTION: 81

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: D (LEAVE A REPLY)

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

NEW QUESTION: 82

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

The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization's firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.

What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while supporting but minimizing manual order processing?

- A. A Batch Job scope to call the back-end system

An Until Successful scope containing Object Store components for long retries A dead-letter object store configured in the Mule application

- B. An Until Successful scope to call the back-end system

One or more ActiveMQ long-retry queues

One or more ActiveMQ dead-letter queues for manual processing

C. One or more On Error scopes to assist calling the back-end system

An Until Successful scope containing VM components for long retries

A persistent dead-letter VM queue configured in CloudHub

D. One or more On Error scopes to assist calling the back-end system

One or more ActiveMQ long-retry queues

A persistent dead-letter object store configured in the CloudHub Object Store service

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

NEW QUESTION: 83

A Mule application is being designed To receive nightly a CSV file containing millions of records from an external vendor over SFTP, The records from the file need to be validated, transformed. And then written to a database. Records can be inserted into the database in any order. In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

A. Use a Parallel for Each scope to Insert records one by one into the database

B. Use a Scatter-Gather to bulk insert records into the database

C. Use a Batch job scope to bulk insert records into the database.

D. Use a DataWeave map operation and an Async scope to insert records one by one into the database.

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

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

* Batch Job is most efficient way to manage millions of records.

A few points to note here are as follows :

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

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

Memory footprint: Since question said that there are millions of records to process, parallel for each will aggregate all the processed records at the end and can possibly cause Out Of Memory. Batch job instead provides a BatchResult in the on complete phase where you can get the count of failures and success. For huge file processing if order is not a concern definitely go ahead with Batch Job

NEW QUESTION: 84

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

applications connect to various backend systems, with multiple applications typically needing to access the backend systems.

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

A. Create a Mule domain project that maintains the credentials as Mule domain-shared resources. Deploy the Mule applications to the Mule domain, so the credentials are available to the Mule applications.

B. Store the credentials in properties files in a shared folder within the organization's data center. Have the Mule applications load properties files from this shared location at startup.

C. Segregate the credentials for each backend system into environment-specific properties files. Package these properties files in each Mule application, from where they are loaded at startup.

D. Configure or create a credentials service that returns the credentials for each backend system, and that is accessible from customer-hosted and MuleSoft-hosted Mule runtimes. Have the Mule applications load the properties at startup by invoking that credentials service.

Answer: C (LEAVE A REPLY)

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

NEW QUESTION: 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. API implementations using an object store must be adjusted to recognize the new nodes and persist to them.

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

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

NEW QUESTION: 86

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 of 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](#))

* 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:



NEW QUESTION: 87

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

location. To achieve optimal performance, what Anypoint Platform region should be chosen for the CloudHub object store?

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

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

*Explanation:

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

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

WHAT DOES "CACHING" MEAN?

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

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

How does Caching work?

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

Caching in MULE 4

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

What is Object Store

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

Cache Scope is used in below-mentioned cases:

- * Need to store the whole response from the outbound processor
- * Data returned from the outbound processor does not change very frequently
- * As Cache scope internally handle the cache hit and cache miss scenarios it is more readable

Object Store is used in below-mentioned cases:

- * Need to store custom/intermediary data
- * To store watermarks
- * Sharing the data/stage across applications, schedulers, batch.

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

NEW QUESTION: 88

An organization has deployed both Mule and non-Mule API implementations to integrate its customer and order management systems. All the APIs are available to REST clients on the public internet.

The organization wants to monitor these APIs by running health checks: for example, to determine if an API can properly accept and process requests. The organization does not have subscriptions to any external monitoring tools and also does not want to extend its IT footprint. What Anypoint Platform feature provides the most idiomatic (used for its intended purpose) way to monitor the availability of both the Mule and the non-Mule API implementations?

- A. API Manager
- B. Runtime Manager
- C. API Functional Monitoring
- D. Anypoint Visualizer

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 89

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

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

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

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

NEW QUESTION: 90

An organization's IT team must secure all of the internal APIs within an integration solution by using an API proxy to apply required authentication and authorization policies.

Which integration technology, when used for its intended purpose, should the team choose to meet these requirements if all other relevant factors are equal?

- A. API Management (APIM)
- B. Electronic Data Interchange (EDI)
- C. Integration Platform-as-a-service (PaaS)
- D. Robotic Process Automation (RPA)

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 91

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

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

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

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here:

https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 92

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: 93

An organization is choosing between API-led connectivity and other integration approaches.

According to MuleSoft, which business benefits is associated with an API-led connectivity approach using Anypoint Platform?

- A. Higher outcome repeatability through centralized development
- B. Greater project predictability through tight coupling of systems
- C. improved security through adoption of monolithic architectures
- D. Increased developer productivity through self-service of API assets

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 94

A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction. What is the behavior of the Scatter-Gather when running within this local transaction?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 95

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

The organization does not currently use AWS in any way.

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

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

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

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

NEW QUESTION: 96

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

How can these requirements be met?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 97

A trading company handles millions of requests a day. Due to nature of its business, it requires excellent performance and reliability within its application.

For this purpose, company uses a number of event-based API's hosted on various mule clusters that communicate across a shared message queue sitting within its network.

Which method should be used to meet the company's requirement for its system?

- A. JMS manual acknowledgements with a reliability pattern
- B. XA transactions and XA connected components
- C. JMS transactions
- D. VM queues with reliability pattern

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

NEW QUESTION: 98

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

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

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

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

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

NEW QUESTION: 99

49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58. <https://api.acme.com/partnefs/v1>. The API implementation is estimated to require deployment to 5 CloudHub workers.

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

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

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

Answer: C (LEAVE A REPLY)

Create a CNAME for `api.acme.com` pointing to the SLB's A record

Explanation:

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


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

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

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

NEW QUESTION: 100

Refer to the exhibit.



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

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

- A.** Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- B.** Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing
- 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: D ([LEAVE A REPLY](#))

NEW QUESTION: 101

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 and acknowledge the JMS message (NOT in an XA transaction)
2. In a NEW XA transaction, perform BOTH DB inserts
- 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 the JMS message (NOT in an XA transaction)
2. Perform EACH DB insert in a SEPARATE DB transaction
3. Acknowledge the JMS message

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

NEW QUESTION: 102

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

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

NEW QUESTION: 103

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

NEW QUESTION: 104

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)

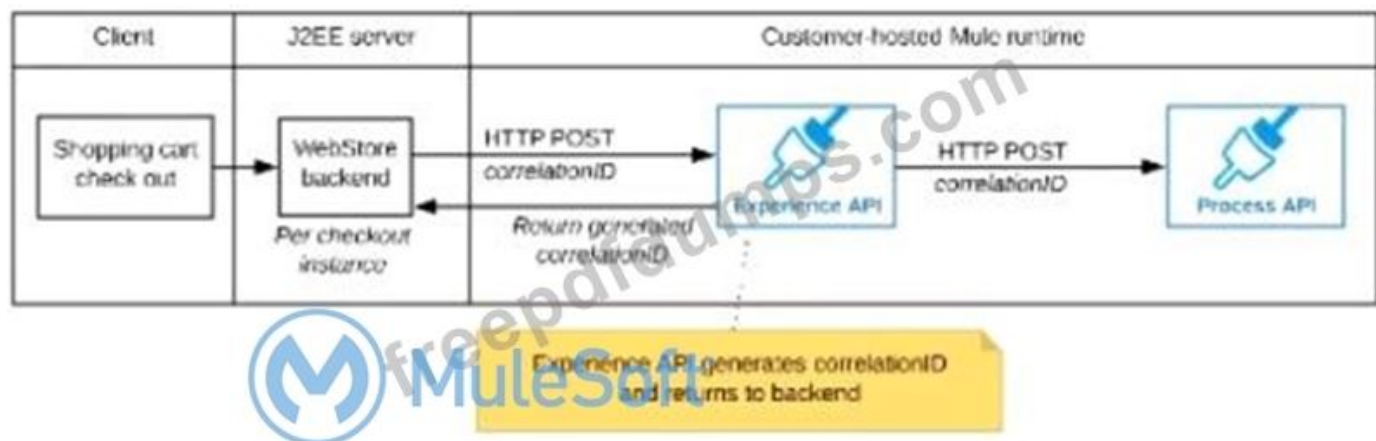
The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header. In each API invocation belonging to that checkout, no special code or configuration is included in the Experience API and Process API implementations to generate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



C)

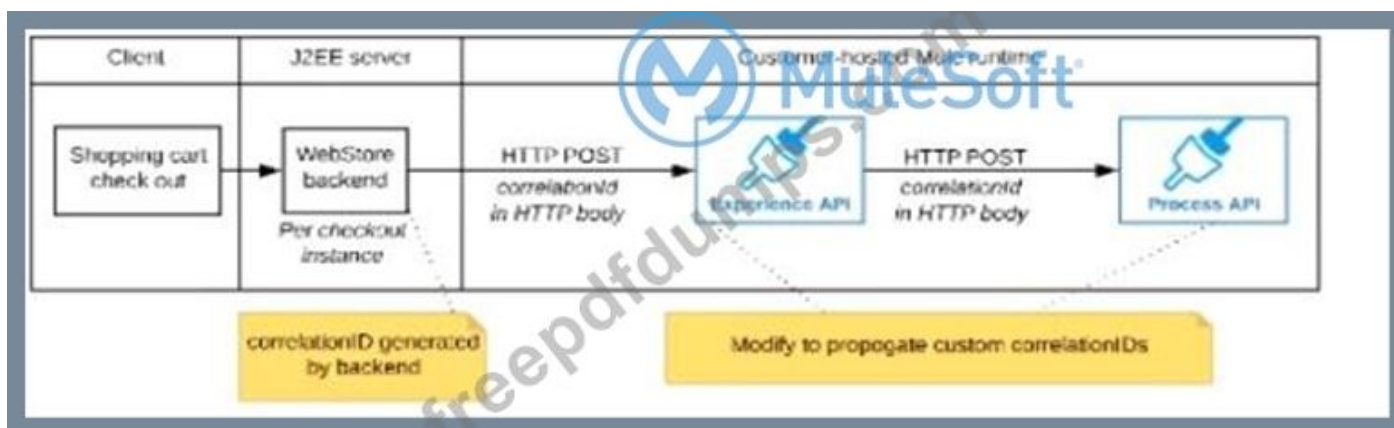
The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

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 A

B. Option B


C. Option C

D. Option D

Answer: B (LEAVE A REPLY)

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

Mulesoft Reference: <https://help.mulesoft.com/s/article/How-to-Set-Custom-Correlation-Id-for-Flows-with-HTTP-Endpoint-in-Mule-4>



NEW QUESTION: 105

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

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

NEW QUESTION: 106

A REST API is being designed to implement a Mule application.

What standard interface definition language can be used to define REST APIs?

- A. YAML
- B. Web Service Definition Language (WSDL)
- C. AsyncAPI Specification
- D. OpenAPI Specification (OAS)

Answer: ([SHOW ANSWER](#)**)**


Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here:

https://www.actual4test.com/MCIA-Level-1_examcollection.html (**275 Q&As Dumps, 30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 107

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.** Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale
- C.** Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
- D.** 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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 108

How does timeout attribute help inform design decisions while using JMS connector listening for incoming messages in an extended architecture (XA) transaction?

- A.** The timeout defines the time that is allowed to pass without the transaction ending explicitly and after the timeout expires, the transaction rolls back
- B.** After the timeout is exceeded, stale JMS consumer threads are destroyed and new threads are created
- C.** The time allowed to pass between committing the transaction and the completion of the mule flow and then after the timeout flow processing triggers an error
- D.** The timeout specifies the time allowed to pass between receiving JMS messages on the same JMS connection and then after the timeout new JMS connection is established

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

NEW QUESTION: 109

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: 110

Which Exchange asset type represents configuration modules that extend the functionality of an API and enforce capabilities such as security?

- A. Policies
- B. Rulesets
- C. REST APIs
- D. Connectors

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 111

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


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

https://aDi.customer.com/ and
https://customer.com/api)

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

NEW QUESTION: 112

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)

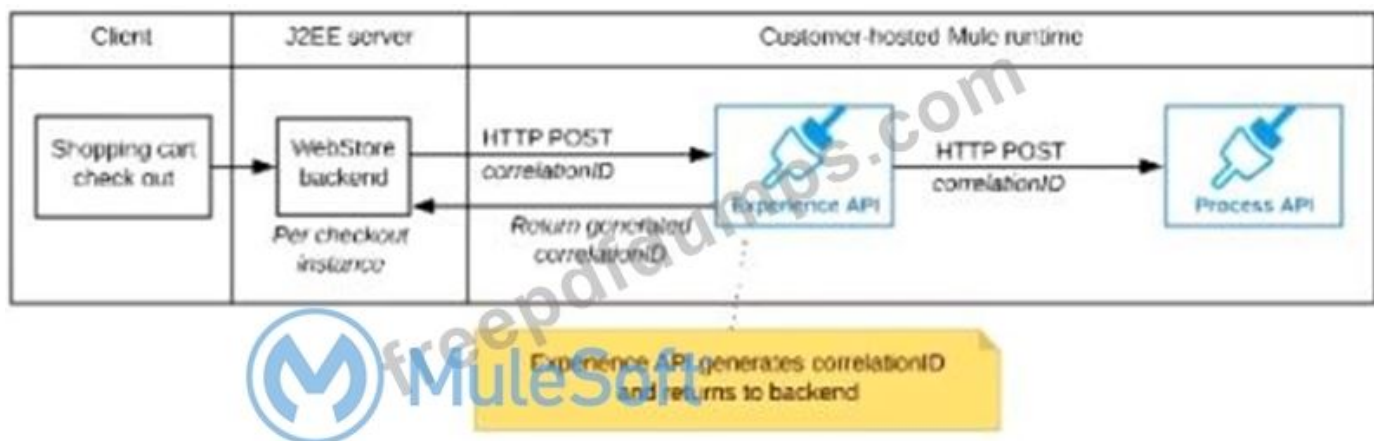
The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-ID HTTP request header. In each API invocation belonging to that checkout. No special code or configuration is included in the Experience API and Process API implementations to generate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



C)

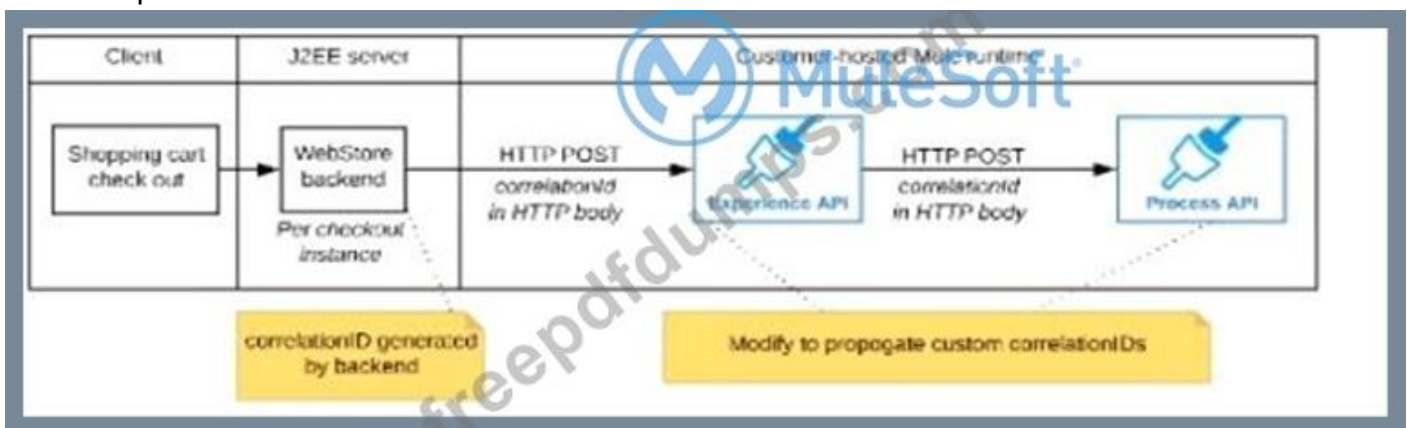
The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API.

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

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

Correct answer is "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" : By design, Correlation Ids cannot be changed within a flow in Mule 4 applications and can be set only at source. This ID is part of the Event Context and is generated as soon as the message is received by the application. When a HTTP Request is received, the request is inspected for "X-Correlation-Id" header. If "X-Correlation-Id" header is present, HTTP connector uses this as the Correlation Id. If "X-Correlation-Id" header is NOT present, a Correlation Id is randomly generated. For Incoming HTTP Requests: In order to set a custom Correlation Id, the client invoking the HTTP request must set "X-Correlation-Id" header. This will ensure that the Mule Flow uses this Correlation Id.

For Outgoing HTTP Requests: You can also propagate the existing Correlation Id to downstream APIs. By default, all outgoing HTTP Requests send "X-Correlation-Id" header. However, you can choose to set a different value to "X-Correlation-Id" header or set "Send Correlation Id" to NEVER.

Mulesoft Reference: <https://help.mulesoft.com/s/article/How-to-Set-Custom-Correlation-Id-for-Flows-with-HTTP-Endpoint-in-Mule-4>



NEW QUESTION: 113

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

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

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

NEW QUESTION: 114

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

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

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

NEW QUESTION: 115

According to MuleSoft, which deployment characteristic applies to a microservices application architecture?

- A.** Services exist as independent deployment artifacts and can be scaled -independently of other services
- B.** Core business capabilities are encapsulated in a single, deployable application
- C.** All services of an application can be deployed together as single Java WAR file
- D.** A deployment to enhance one capability requires a redeployment of all capabilities

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

NEW QUESTION: 116

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

The organization does not currently use AWS in any way.

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

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

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

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

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

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

NEW QUESTION: 117

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the

base URL `https://api.acme.com` As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

A. Deploy the Mule applications to Cloudhub

Update the CNAME record for `api.acme.com` in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB) Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB) Update the CNAME record for `api.acme.com` in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB) Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications

C. Deploy the Mule applications to Cloudhub

Update the CNAME record for an `api.acme.com` in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB) Apply mapping rules in the DLB to map URLs to their corresponding Mule applications

D. Deploy the Mule applications to Cloudhub

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

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

NEW QUESTION: 118

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: ([SHOW ANSWER](#))

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: 119

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 secure the service, requiring all consumers to submit a valid SAML token
- 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 use XML payloads for the service and ensure that it adheres to a specific schema

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

NEW QUESTION: 120

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.** Read and acknowledge the JMS message (NOT in an XA transaction)
2) In a NEW XA transaction, perform BOTH DB inserts
- B.** 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
- C.** Read the JMS message (NOT in an XA transaction)
2) Perform BOTH DB inserts in ONE DB transaction
3) Acknowledge the JMS message

- D. Read the JMS message (NOT in an XA transaction)
- 2) Perform EACH DB insert in a SEPARATE DB transaction
- 3) Acknowledge the JMS message

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 121

Refer to the exhibit.



One of the backend systems invoked by an API implementation enforces rate limits on the number of requests a particular client can make. Both the backend system and the API implementation are deployed to several non-production environments in addition to production. Rate limiting of the backend system applies to all non-production environments. The production environment, however, does NOT have any rate limiting.

What is the most effective approach to conduct performance tests of the API implementation in a staging (non-production) environment?

- A. Create a mocking service that replicates the backend system's production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests
- B. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
- C. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situation. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
- 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: ([SHOW ANSWER](#))

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

* MUnit is for only Unit and integration testing for APIs and Mule apps. Not for performance Testing, even if it has the ability to Mock the backend.

- * Bypassing the backend invocation defeats the whole purpose of performance testing. Hence it is not a valid answer.
- * Scaled down performance tests can't be relied upon as performance of API's is not linear against load.

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here: https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 122

What is a recommended practice when designing an integration Mule 4 application that reads a large XML payload as a stream?

- A. The payload should be dealt with as an XML stream, without converting it to a single Java object (POJO)
- B. The payload size should NOT exceed the maximum available heap memory of the Mule runtime on which the Mule application executes
- C. The payload should be dealt with as a repeatable XML stream, which must only be traversed (iterated-over) once and CANNOT be accessed randomly from DataWeave expressions and scripts
- D. The payload must be cached using a Cache scope if it is to be sent to multiple backend systems

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 123

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

- A. Support for TLS mutual (two-way) authentication with 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 and HTTP response bodies for all HTTP clients
- D. Encryption of both HTTP request header and HTTP request body for all HTTP clients

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

1 way SSL : The server presents its certificate to the client and the client adds it to its list of trusted certificate. And so, the client can talk to the server.

2-way SSL: The same principle but both ways. i.e. both the client and the server has to establish trust between themselves using a trusted certificate. In this way of a digital handshake, the server

needs to present a certificate to authenticate itself to client and client has to present its certificate to server.


* 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. Keystores and Truststores: Truststore and keystore contents differ depending on whether they are used for clients or servers:

For servers: the truststore contains certificates of the trusted clients, the keystore contains the private and public key of the server. For clients: the truststore contains certificates of the trusted servers, the keystore contains the private and public key of the client.

Adding both a keystore and a truststore to the configuration implements two-way TLS authentication also known as mutual authentication.

* in this case, correct answer is Support for TLS mutual (two-way) authentication with HTTP clients.



NEW QUESTION: 124

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

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

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

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

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

NEW QUESTION: 125

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

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

Answer: ([SHOW ANSWER](#))

High availability is about up-time of your application


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

Reference:

C) High availability can be achieved by below two models for on-premise MuleSoft implementations.

1) Mule Clustering - Where multiple Mule servers are available within the same cluster environment and the routing of requests will be done by the load balancer. A cluster is a set of up to eight servers that act as a single deployment target and high-availability processing unit. Application instances in a cluster are aware of each other, share common information, and synchronize statuses. If one server fails, another server takes over processing applications. A cluster can run multiple applications. (refer left half of the diagram) In given scenario, it's mentioned that 'data cannot be shared among of different instances'. So this is not a correct choice.

2) Load balanced standalone Mule instances - The high availability can be achieved even without cluster, with the usage of third party load balancer pointing requests to different Mule servers. This approach does not share or synchronize data between Mule runtimes. Also high availability achieved as load balanced algorithms can be implemented using external load balancer. (refer right half of the diagram)



NEW QUESTION: 126

The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure.

Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).

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

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

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

NEW QUESTION: 127

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL <https://api.acme.com> As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

- A. Deploy the Mule applications to Cloudhub

Update the CNAME record for an api.acme.com in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB) Apply mapping rules in the DLB to map URLs to their corresponding Mule applications

- B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB) Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of a cloudhub dedicated

load balancer(DLB) Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications

C. Deploy the Mule applications to Cloudhub

Create CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancer Apply mapping rules in the SLB to map URLs to their corresponding Mule applications

D. Deploy the Mule applications to Cloudhub

Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB) Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

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

Explanation

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

NEW QUESTION: 128

An airline is architecting an API connectivity project to integrate its flight data into an online aggregation website. The interface must allow for secure communication high-performance and asynchronous message exchange.

What are suitable interface technologies for this integration assuming that Mulesoft fully supports these technologies and that Anypoint connectors exist for these interfaces?

A. AsyncAPI over HTTPS

AMQP with RabbitMQ JSON/REST over HTTPS

B. XML over ActiveMQ XML over SFTP XML/REST over HTTPS

C. SOAP over HTTPS HOP over TLS gRPC over HTTPS

D. CSV over FTP YAM L over TLS JSON over HTTPS

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

NEW QUESTION: 129

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 130

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

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

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

A. Add the dbPassword property to the secureProperties section of the pom.xml file

B. Use secure::dbPassword as the property placeholder name and store the cleartext (unencrypted) value in a secure properties placeholder file

C. Add the dbPassword property to the secureProperties section of the mule-artifact.json file

D. Use secure::dbPassword as the property placeholder name and store the property encrypted value in a secure properties placeholder file

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

NEW QUESTION: 131

Which Anypoint Platform component should a MuleSoft developer use to create an API specification prior to building the API implementation?

A. Runtime Manager

B. API Manager


C. MUnit

D. API Designer

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

NEW QUESTION: 132

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

Answer: (SHOW ANSWER)

Mule applications are deployed in CloudHub workers and each worker is assigned with a dedicated IP * For zero downtime deployment, each worker in CloudHub needs additional IP addresses * A few IPs in a VPC are reserved for infrastructure (generally 2 IPs) * The IP addresses are usually in a private range with a subnet block specifier, such as 10.0.0.1/24 * The smallest CIDR network subnet block you can assign for your VPC is /24 (256 IP addresses) (60*3

env * 2 worker per application) + 50% of (total) for zero downtime = 540 In this case correct answer is 10.0.0.0/22 as this provided 1024 IP's .
Other IP's are insufficient.

NEW QUESTION: 133

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: 134

Which Anypoint Platform component helps integration developers discover and share reusable APIs, connectors, and templates?

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

Answer: (SHOW ANSWER)


NEW QUESTION: 135

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

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

NEW QUESTION: 136

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

The organization has a policy that all project teams should leverage existing investments by reusing existing APIs and related resources and documentation that other project teams have already developed and deployed.

To support reuse, where on Anypoint Platform should the project team go to discover and read existing APIs, discover related resources and documentation, and interact with mocked versions of those APIs?

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

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

:

The mocking service is a feature of Anypoint Platform and runs continuously. You can run the mocking service from the text editor, the visual editor, and from Anypoint Exchange. You can

simulate calls to the API in API Designer before publishing the API specification to Exchange or in Exchange after publishing the API specification.

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here: https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 137

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

A. Implement a For-Each scope to invoke the two vendor applications

Use until successful scope to implement the retry mechanism for the timeout errors

B. Implement a Choice scope to invoke the two vendor applications on two different route Use the try-catch scope to implement the retry mechanism for timeout errors on each route

C. Implement Round-Robin scope to invoke the two vendor applications on two different routes Use the Try-Catch scope to implement retry mechanism for timeout errors on each route

D. Implement a scatter-gather scope to invoke the two vendor applications on two different route

Use the Until-Successful scope to implement the retry mechanism for timeout errors on each route

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 138

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

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

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 139

A leading e-commerce giant will use Mulesoft API's on runtime fabric (RTF) to process customer orders. Some customer's sensitive information such as credit card information is also there as a part of a API payload.

What approach minimizes the risk of matching sensitive data to the original and can convert back to the original value whenever and wherever required?

- A.** Apply masking to hide the sensitive information and then use API
- B.** create a tokenization format and apply a tokenization policy to the API Gateway
- C.** manager to detokenize the masking format to return the original value
- D.** Apply a field level encryption policy in the API Gateway
- E.** Used both masking and tokenization

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

NEW QUESTION: 140

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.** Parallel For Each scope 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.** Until Successful router 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: ([SHOW ANSWER](#))

NEW QUESTION: 141

A platform architect includes both an API gateway and a service mesh in the architect of a distributed application for communication management.

Which type of communication management does a service mesh typically perform in this architecture?

- A. Between application services and the firewall
- B. Between services within the application
- C. Between the application and external API implementations.
- D. Between the application and external API clients

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

NEW QUESTION: 142

Mule application muleA deployed in cloudhub uses Object Store v2 to share data across instances. As a part of new requirement , application muleB which is deployed in same region wants to access this Object Store.

Which of the following option you would suggest which will have minimum latency in this scenario?

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

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

V2 Rest API is recommended for on premise applications to access Object Store. It also comes with overhead of encryption and security of using rest api. With Object Store v2, the API call is localized to the same data center as the Runtime Manager app.

But in this case requirement is to access the OS of other mule application and not the same mule application.

You can configure a Mule app to use the Object Store REST API to store and retrieve values from an object store in another Mule app.

However, Object Store v2 is not designed for app-to-app communication.

NEW QUESTION: 143

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

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

- A. Configure or create a credentials service that returns the credentials for each backend system, and that is accessible from customer-hosted and MuleSoft-hosted Mule runtimes Have the Mule applications load the properties at startup by invoking that credentials service
- B. Store the credentials in properties files in a shared folder within the organization's data center Have the Mule applications load properties files from this shared location at startup

C. Create a Mule domain project that maintains the credentials as Mule domain-shared resources
Deploy the Mule applications to the Mule domain, so the credentials are available to the Mule applications

D. Segregate the credentials for each backend system into environment-specific properties files
Package these properties files in each Mule application, from where they are loaded at startup

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

NEW QUESTION: 144

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: 145

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: (SHOW ANSWER)

* "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" is not a feasible option

* "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" - It is risk for sensitive data. - Even if you whitelist the database IP on your app, your app wont be able to connect to the database so this is also not a feasible option

* "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" Adding one VPC with a DLB for each backend system also makes no sense, is way too much work. Why would you add a LB for one system.

* Correct answer: "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" IPsec Tunnel You can use an IPsec tunnel with network-to-network configuration to connect your on-premises data centers to your Anypoint VPC. An IPsec VPN tunnel is generally the recommended solution for VPC to on-premises connectivity, as it provides a standardized, secure way to connect. This method also integrates well with existing IT infrastructure such as routers and appliances. Reference: <https://docs.mulesoft.com/runtime-manager/vpc-connectivity-methods-concept>



NEW QUESTION: 146

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: (SHOW ANSWER)

* 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: 147

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

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

A. IBM: ActiveMQ

CRM: REST

B. IBM: REST

CRM: REST

C. IBM: DB access

CRM: gRPC


D. IBM: CICS

CRM: SOAP

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 148

Refer to the exhibit.



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

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

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

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

Answer: C (LEAVE A REPLY)

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

NEW QUESTION: 149

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: 150

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: 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: 151

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 EACH Mule application's log4j2.xml file

B. Disable Cloudhub logging in Runtime Manager

Define the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support all Mule application deployments.

C. Keep the default logging configuration in Runtime Manager

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

D. Keep the default logging configuration in Runtime Manager

Define the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to support all Mule application deployments.

Answer: ([SHOW ANSWER](#))

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here:

NEW QUESTION: 152

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: D (LEAVE A REPLY)

* 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.




NEW QUESTION: 153

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. WSDL
- C. RAMI
- D. JSON Schema

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

NEW QUESTION: 154

An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub. For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub's default zero-downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers. This is expected to result in several Mule application deployments per hour.

- A. 10.0.0.0/24(256 IPs)
- B. 10.0.0.0/23(512 IPs)
- C. 10.0.0.0/21(2048 IPs)
- D. 10.0.0.0/22(1024 IPs)

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

NEW QUESTION: 155

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL <https://api.acme.com> As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

- A. Deploy the Mule applications to Cloudhub

Update the CNAME record for an api.acme.com in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB) Apply mapping rules in the DLB to map URLs to their corresponding Mule applications

B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB) Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB) Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications

C. Deploy the Mule applications to Cloudhub

Create CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancer Apply mapping rules in the SLB to map URLs to their corresponding Mule applications

D. Deploy the Mule applications to Cloudhub

Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB) Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

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

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

NEW QUESTION: 156

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: 157

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: 158

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

A. Set a timeout of 100 ms; that leaves 400 ms for the other two downstream APIs to complete

B. Do not set a timeout; the Invocation of this API is mandatory and so we must wait until it responds

C. Set a timeout of 50 ms; this times out more invocations of that API but gives additional room for retries

D. No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

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

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

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

Let see why other answers are not correct.

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

2) Set a timeout of 100 ms; ---> This will not work as backend API has 100 ms as median meaning only 50% requests will be answered in this time and we will get timeout for 50% of the requests. Important thing to note here is, All APIs need to be executed sequentially, so if you get timeout in first API, there is no use of going to second and third API. As a service provider you wouldn't want to keep 50% of your consumers dissatisfied. So not the best option to go with.

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

- First API is fetching customer number based on login credentials

- Second API is fetching Info in 1 table and returning unique key

- Third API, using unique key provided in second API as primary key, updating remaining details

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

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

NEW QUESTION: 159

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

A. Support for TLS mutual (two-way) authentication with 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 and HTTP response bodies for all HTTP clients

D. Encryption of both HTTP request header and HTTP request body for all HTTP clients

Answer: A (LEAVE A REPLY)

1 way SSL : The server presents its certificate to the client and the client adds it to its list of trusted certificate. And so, the client can talk to the server.

2-way SSL: The same principle but both ways. i.e. both the client and the server has to establish trust between themselves using a trusted certificate. In this way of a digital handshake, the server needs to present a certificate to authenticate itself to client and client has to present its certificate to server.


* 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 Keystores and Truststores Truststore and keystore contents differ depending on whether they are used for clients or servers:

For servers: the truststore contains certificates of the trusted clients, the keystore contains the private and public key of the server. For clients: the truststore contains certificates of the trusted servers, the keystore contains the private and public key of the client.

Adding both a keystore and a truststore to the configuration implements two-way TLS authentication also known as mutual authentication.

* in this case, correct answer is Support for TLS mutual (two-way) authentication with HTTP clients.



NEW QUESTION: 160

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 supports Transactional Type. When an exception occurs, 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/ip-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: 161

According to MuleSoft, which major benefit does a Center for Enablement (C4E) provide for an enterprise and its lines of business?

- A. Enabling Edge security between the lines of business and public devices
- B. Accelerating self-service by the lines of business
- C. Centralizing project management across the lines of business
- D. Centrally managing return on investment (ROI) reporting from lines of business to leadership

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

NEW QUESTION: 162

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

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

NEW QUESTION: 163

According to MuleSoft's IT delivery and operating model, which approach can an organization adopt in order to reduce the frequency of IT project delivery failures?

- A.** Adopt an enterprise data model
- B.** Stop scope creep by centralizing requirements-gathering
- C.** Decouple central IT projects from the innovation that happens within each line of business
- D.** Prevent technology sprawl by reducing production of API assets

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

NEW QUESTION: 164

A mule application is deployed to a Single Cloudhub worker and the public URL appears in Runtime Manager as the APP URL.

Requests are sent by external web clients over the public internet to the mule application App url. Each of these requests routed to the HTTPS Listener event source of the running Mule application.

Later, the DevOps team edits some properties of this running Mule application in Runtime Manager.

Immediately after the new property values are applied in runtime manager, how is the current Mule application deployment affected and how will future web client requests to the Mule application be handled?

- A.** Cloudhub will redeploy the Mule application to the OLD Cloudhub worker New web client requests will RETURN AN ERROR until the Mule application is redeployed to the OLD Cloudhub worker
- B.** Cloudhub will redeploy the mule application to the OLD Cloudhub worker New web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.

C. CloudHub will redeploy the Mule application to a NEW Cloudhub worker New web client requests will RETURN AN ERROR until the NEW Cloudhub worker is available

D. Cloudhub will redeploy the Mule application to a NEW Cloudhub worker New webclient requests are ROUTED to the OLD Cloudhub worker until the NEW Cloudhub worker is available.

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

NEW QUESTION: 165

An integration Mute application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mute application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system. Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system).

The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed.

The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages.

What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

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

B. An On Error scope MuleSoft Object Store ActiveMQ Dead Letter Queue for manual processing

C. Until Successful component MuleSoft Object Store ActiveMQ is NOT needed or used

D. Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing

Answer: ([SHOW ANSWER](#)**)**

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

can intercept dead letter processing and reschedule the failing messages for redelivery. Rather than being delivered to a DLQ, a failing message is scheduled to go to the tail of the original queue and redelivered to a message consumer. How this will help requirement : If backend application is down for a longer duration where Until Successful Scope wont work, then we can make use of ActiveMQ long retry Queue. The redelivery plug-in can intercept dead letter processing and reschedule the failing messages for redelivery. Mule Reference: <https://docs.mulesoft.com/mule-runtime/4.3/migration-core-until-successful>

NEW QUESTION: 166

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 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 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
- 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: B (LEAVE A REPLY)

* 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:



Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here:

NEW QUESTION: 167

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

How are the messages consumed by the Mule application?

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

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

Correct answer is Depending on the Listener operation configuration, either all messages are consumed by ONLY the primary cluster node or else EACH message is consumed by ANY ONE cluster node For applications running in clusters, you have to keep in mind the concept of primary node and how the connector will behave. When running in a cluster, the JMS listener default behavior will be to receive messages only in the primary node, no matter what kind of destination you are consuming from. In case of consuming messages from a Queue, you'll want to change this configuration to receive messages in all the nodes of the cluster, not just the primary.

This can be done with the `primaryNodeOnly` parameter:

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

NEW QUESTION: 168

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

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

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


Answer: ([SHOW ANSWER](#)**)**

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

You cant use in memory storage because with 0.2 vcore you will get only 1 GB of heap memory.
Hence application will error out for file more than 1 GB.
Hence the correct option is file base repeatable stream

NEW QUESTION: 169

Refer to the exhibit.



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

Step 1: The JMS Correlation ID header is read from the received JMS message.

Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.

Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.

Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.

Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.

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

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

Answer: C (LEAVE A REPLY)

* If we store Correlation id value in step 1 as Mule event variables/attributes, the values will be cleared after server restart and we want system to be fault tolerant.

* The Correlation ID value in Step 1 should be stored in a persistent object store.

* We don't need to store Correlation ID value in Step 3 to persistent object store. We can store it but as we also need to make application performant. We can avoid this step of accessing persistent object store.


* Accessing persistent object stores slow down the performance as persistent object stores are by default stored in shared file systems.

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

Top of Form

Additional Information:

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



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

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

IDEMPOTENCE MuleSoft

WHEN PERFORMING AN OPERATION AGAIN GIVES THE SAME RESULT

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

Bottom of Form

NEW QUESTION: 170

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 171

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. Availability
- D. Response time

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

NEW QUESTION: 172

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

(Choose two answers)

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

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

NEW QUESTION: 173

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

How is access enabled for the API to connect to the database application and Kafka cluster securely?

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

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 174

An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments.

What must the testing team do before they can start integration testing the API in the Staging environment?

- A. They must access the API portal and create an API notebook using the Client ID and Client Secret supplied by the API portal in the Staging environment

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


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

- * It's mentioned that the API is governed by a Client ID Enforcement policy in all environments.
- * Client ID Enforcement policy allows only authorized applications to access the deployed API implementation.
- * Each authorized application is configured with credentials: client_id and client_secret.
- * At runtime, authorized applications provide the credentials with each request to the API implementation.

MuleSoft Reference: <https://docs.mulesoft.com/api-manager/2.x/policy-mule3-client-id-based-policies>

NEW QUESTION: 175

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

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

NEW QUESTION: 176

Which role is primarily responsible for building API implementation as part of a typical MuleSoft integration project?

- A.** API Designer


- B. API Developer
- C. Integration Architect
- D. Operations

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 177

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 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
- B. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing
- C. Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- 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: 178

An API implementation is being designed that must invoke an Order API which is known to repeatedly experience downtime. For this reason a fallback API is to be called when the Order API is unavailable. What approach to designing invocation of the fallback API provides the best resilience?

- A.** Redirect client requests through an HTTP 303 temporary redirect status code to the fallback API whenever the Order API is unavailable
- B.** Set an option in the HTTP Requester component that invokes the order API to instead invoke a fallback API whenever an HTTP 4XX or 5XX response status code is received from Order API
- C.** Create a separate entry for the order API in API manager and then invoke this API as a fallback API if the primary Order API is unavailable
- D.** Search Anypoint Exchange for a suitable existing fallback API and then implement invocations to their fallback API in addition to the Order API

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

* Resilience testing is a type of software testing that observes how applications act under stress. It's meant to ensure the product's ability to perform in chaotic conditions without a loss of core functions or data; it ensures a quick recovery after unforeseen, uncontrollable events.

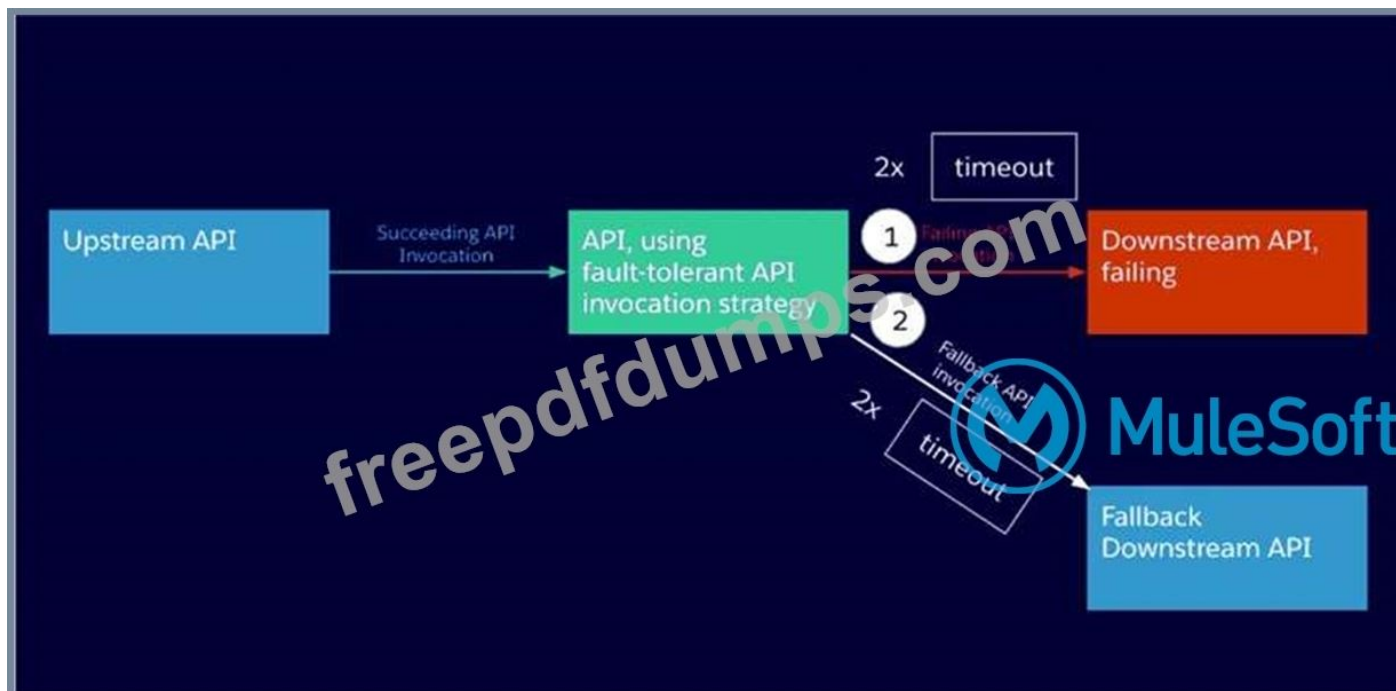
* In case an API invocation fails - even after a certain number of retries - it might be adequate to invoke a different API as a fallback. A fallback API, by definition, will never be ideal for the purpose of the API client, otherwise it would be the primary API.

* Here are some examples for fallback APIs:

- An old, deprecated version of the same API.
- An alternative endpoint of the same API and version (e.g. API in another CloudHub region).
- An API doing more than required, and therefore not as performant as the primary API.
- An API doing less than required and therefore forcing the API Client to offer a degraded service, which is still better than no service at all.

* API clients implemented as Mule applications offer the 'Until Successful Scope and Exception' strategies at their disposal, which together allow configuring fallback actions such as a fallback API invocation.

* All HTTP response status codes within the 3xx category are considered redirection messages. These codes indicate to the user agent (i.e. your web browser) that an additional action is required in order to complete the request and access the desired resource



Hence correct answer is Redirect client requests through an HTTP 303 temporary redirect status code to the fallback API whenever the Order API is unavailable

Valid MCIA-Level-1 Dumps shared by Actual4test.com for Helping Passing MCIA-Level-1 Exam! Actual4test.com now offer the **newest MCIA-Level-1 exam dumps**, the Actual4test.com MCIA-Level-1 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com MCIA-Level-1 dumps with Test Engine here:

https://www.actual4test.com/MCIA-Level-1_examcollection.html (275 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)