


New Braindumps MuleSoft-Integration-Architect-I Book, Flexible MuleSoft-Integration-Architect-I Testing Engine



2026 Latest TestPDF MuleSoft-Integration-Architect-I PDF Dumps and MuleSoft-Integration-Architect-I Exam Engine Free Share: <https://drive.google.com/open?id=1O50bRKdLwhovIT2pW6URZ0ln0nsFBFrD>

We know that tenet from the bottom of our heart, so all parts of service are made due to your interests. You are entitled to have full money back if you fail the exam even after getting our MuleSoft-Integration-Architect-I test prep. Our staff will help you with genial attitude. We esteem your variant choices so all these versions of MuleSoft-Integration-Architect-I Study Materials are made for your individual preference and inclination. Please get to know our MuleSoft-Integration-Architect-I study materials as follows.

Salesforce MuleSoft-Integration-Architect-I Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">Designing for the Runtime Plane Technology Architecture: It includes analyzing Mule runtime clusters, designing solutions for CloudHub, choosing Mule runtime domains, leveraging Mule 4 class loader isolation, and understanding the reactive event processing model.
Topic 2	<ul style="list-style-type: none">Initiating Integration Solutions on Anypoint Platform: Summarizing MuleSoft Catalyst and Catalyst Knowledge Hub, differentiating between functional and non-functional requirements, selecting features for designing and managing APIs, and choosing deployment options are its sub-topics.
Topic 3	<ul style="list-style-type: none">Designing Integration Solutions to Meet Persistence Requirements: It addresses the usage of VM queues and connectors, object stores and services, and stateful components configured with object stores.
Topic 4	<ul style="list-style-type: none">Designing Automated Tests for Mule Applications: This topic covers unit test suites, and scenarios for integration and performance testing.
Topic 5	<ul style="list-style-type: none">Applying DevOps Practices and Operating Integration Solutions: Its sub-topics are related to designing CI CD pipelines with MuleSoft plugins, automating interactions with Anypoint Platform, designing logging configurations, and identifying Anypoint Monitoring features.

>> New Braindumps MuleSoft-Integration-Architect-I Book <<

100% Pass Quiz 2026 High Pass-Rate MuleSoft-Integration-Architect-I: New Braindumps Salesforce Certified MuleSoft Integration Architect I Book

To do this you just need to pass MuleSoft-Integration-Architect-I exam, which is quite challenging and demands thorough Salesforce Certified MuleSoft Integration Architect I (MuleSoft-Integration-Architect-I) exam preparation. For the complete, comprehensive and quick MuleSoft-Integration-Architect-I Exam Preparation, the TestPDF MuleSoft-Integration-Architect-I

Dumps questions are ideal. You should not ignore it and must try TestPDF MuleSoft-Integration-Architect-I exam questions for preparation today.

Salesforce Certified MuleSoft Integration Architect I Sample Questions (Q245-Q250):

NEW QUESTION # 245

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

Answer: D

NEW QUESTION # 246

A stock trading company handles millions of trades a day and requires excellent performance and reliability within its stock trading system. The company operates a number of event-driven APIs Implemented as Mule applications that are hosted on various customer-hosted Mule clusters and needs to enable message exchanges between the APIs within their internal network using shared message queues.

What is an effective way to meet the cross-cluster messaging requirements of its event-driven APIs?

- A. extended Architecture (XA) transactions and XA connected components with manual acknowledgements
- **B. JMS transactions with automatic acknowledgements**
- C. Non-transactional JMS operations with a reliability pattern and manual acknowledgements
- D. Persistent VM queues with automatic acknowledgements

Answer: B

Explanation:

* JMS (Java Message Service): JMS is a robust messaging standard that supports reliable and asynchronous communication. It allows message producers and consumers to exchange messages via a common message broker.

* Transactions with Automatic Acknowledgements: Utilizing JMS transactions ensures that messages are processed reliably. The automatic acknowledgement mode means that once the consumer receives the message, it acknowledges the broker automatically, ensuring that no messages are lost.

* Performance and Reliability: JMS transactions offer both high performance and reliability. By enabling transactions, each message processing step can be committed or rolled back, ensuring data integrity.

* Cross-Cluster Messaging: For a stock trading company dealing with millions of trades, using JMS transactions allows for consistent and reliable message delivery across different clusters in their network. This approach is more suitable compared to non-transactional or VM queues due to the scale and reliability requirements.

* Event-Driven APIs: The APIs can leverage the transactional nature of JMS to ensure that messages exchanged between different services are reliable and can recover gracefully from failures.

References:

* MuleSoft Documentation on JMS Connector: MuleSoft JMS Connector

* JMS 2.0 Specification: Oracle JMS 2.0

NEW QUESTION # 247

A banking company is developing a new set of APIs for its online business. One of the critical API's is a master lookup API which is a system API. This master lookup API uses persistent object store. This API will be used by all other APIs to provide master lookup data.

□

Master lookup API is deployed on two cloudbus workers of 0.1 vCore each because there is a lot of master data to be cached. Master lookup data is stored as a key value pair. The cache gets refreshed if the key is not found in the cache. Doing performance testing it was observed that the Master lookup API has a higher response time due to database queries execution to fetch the master lookup data. Due to this performance issue, go-live of the online business is on hold which could cause potential financial loss to Bank. As an integration architect, which of the below option you would suggest to resolve performance issue?

- A. Implement HTTP caching policy for all GET endpoints for master lookup API
- B. **Implement HTTP caching policy for all GET endpoints for the master lookup API and implement locking to synchronize access to object store**
- C. Upgrade vCore size from 0.1 vCore to 0.2 vCore
- D. Add an additional Cloudbus worker to provide additional capacity

Answer: B

Explanation:

A: Implement HTTP caching policy for all GET endpoints for the master lookup API and implement locking to synchronize access to object store Comprehensive Detailed Step by Step Explanation To resolve the performance issue observed during the performance testing of the Master lookup API, the best approach involves caching and synchronization:

* HTTP Caching Policy:

* Purpose: Reduces the load on the database by caching the responses for GET requests. Once a response is cached, subsequent requests for the same resource can be served from the cache instead of querying the database.

* Implementation: Apply the HTTP caching policy to all GET endpoints of the Master lookup API. This policy ensures that the frequently accessed master lookup data is stored in the cache and served quickly.

* Benefits: This significantly reduces the number of database queries, thus lowering the response time and improving performance.

* Object Store Locking:

* Purpose: Ensures data consistency and prevents race conditions when multiple workers attempt to access or update the cache concurrently.

* Implementation: Use a locking mechanism to synchronize access to the object store. This can be achieved by using MuleSoft's Object Store Connector with lock/unlock operations.

* Benefits: Prevents multiple database queries for the same key and ensures that only one worker updates the cache at a time, thus avoiding redundant operations and maintaining data integrity.

By implementing these two strategies, the response time of the Master lookup API will be significantly improved, and the performance issue will be resolved.

References

* MuleSoft HTTP Caching Policy

* MuleSoft Object Store Connector

NEW QUESTION # 248

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

Explanation:

Before we answer this question, we need to understand what median (50th percentile) and 80th percentile means. If the 50th percentile (median) of a response time is 500ms that means that 50% of my transactions are either as fast or faster than 500ms. If the 90th percentile of the same transaction is at 1000ms it means that 90% are as fast or faster and only 10% are slower. Now as per upstream SLA, 99th percentile is 800 ms which means 99% of the incoming requests should have response time less than or equal to 800 ms. But as per one of the backend API, their 95th percentile is 1000 ms which means that backend API will take 1000 ms or less than that for 95% of requests. As there are three API invocation from upstream API, we can not conclude a

timeout that can be set to meet the desired SLA as backend SLA's do not support it.

Let see why other answers are not correct.

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

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

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

- First API is fetching customer number based on login credentials

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

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

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

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

NEW QUESTION # 249

What is not true about Mule Domain Project?

- A. Expose multiple services within the Mule domain on the same port
- B. This allows Mule applications to share resources
- C. Send events (messages) to other Mule applications using VM queues
- D. Only available Anypoint Runtime Fabric

Answer: D

Explanation:

* Mule Domain Project is ONLY available for customer-hosted Mule runtimes, but not for Anypoint Runtime Fabric

* Mule domain project is available for Hybrid and Private Cloud (PCE). Rest all provide application isolation and can't support domain project.

What is Mule Domain Project?

* A Mule Domain Project is implemented to configure the resources that are shared among different projects.

These resources can be used by all the projects associated with this domain. Mule applications can be associated with only one domain, but a domain can be associated with multiple projects. Shared resources allow multiple development teams to work in parallel using the same set of reusable connectors. Defining these connectors as shared resources at the domain level allows the team to: - Expose multiple services within the domain through the same port. - Share the connection to persistent storage. - Share services between apps through a well-defined interface. - Ensure consistency between apps upon any changes because the configuration is only set in one place.

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

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

You can start using domains in only three steps:

1) Create a Mule Domain project

2) Create the global connector configurations which needs to be shared across the applications inside the Mule Domain project

3) Modify the value of domain in mule-deploy.properties file of the applications Graphical user interface Description automatically generated

NEW QUESTION # 250

.....

The Salesforce MuleSoft-Integration-Architect-I exam questions on the platform have been gathered by subject matter experts to ensure that they accurately reflect the format and difficulty level of the actual Salesforce MuleSoft-Integration-Architect-I exam. This

makes these Salesforce Certified MuleSoft Integration Architect I PDF Questions ideal for individuals looking to pass the Salesforce MuleSoft-Integration-Architect-I Exam on their first try. You can evaluate the product with a free MuleSoft-Integration-Architect-I demo.

Flexible MuleSoft-Integration-Architect-I Testing Engine: <https://www.testpdf.com/MuleSoft-Integration-Architect-I-exam-braindumps.html>

P.S. Free & New MuleSoft-Integration-Architect-I dumps are available on Google Drive shared by TestPDF: <https://drive.google.com/open?id=1O50bRKdLwhovIT2pW6URZ0ln0nsFBFrD>