

Get The Most Updated MCIA-Level-1-Maintenance Dumps To MuleSoft Certified Architect Certification [Q14-Q36]



Get The Most Updated MCIA-Level-1-Maintenance Dumps To MuleSoft Certified Architect Certification MuleSoft Certified MCIA-Level-1-Maintenance Dumps Questions Valid MCIA-Level-1-Maintenance Materials

The MCIA-Level-1-Maintenance Exam is a valuable certification for individuals who work with MuleSoft's Anypoint Platform. It is ideal for individuals who are responsible for maintaining and troubleshooting MuleSoft's integration platform, and for those who want to demonstrate their expertise in the field. By obtaining this certification, individuals can increase their value to organizations that use MuleSoft's Anypoint Platform and enhance their career prospects.

Achieving the MuleSoft MCIA-Level-1-Maintenance Certification is a valuable credential for professionals who want to advance their careers in MuleSoft integration architecture. MuleSoft Certified Integration Architect - Level 1 MAINTENANCE certification demonstrates a candidate's expertise in MuleSoft integration and their ability to design, build, and maintain MuleSoft solutions to meet the needs of modern businesses.

QUESTION 14

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

- * After the timeout is exceeded, stale JMS consumer threads are destroyed and new threads are created
- * 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
- * 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
- * The timeout defines the time that is allowed to pass without the transaction ending explicitly and after the timeout expires, the transaction rolls back

QUESTION 15

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

Each System API makes calls to a single database.

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

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

- * Implement eXtended Architecture (XA) transactions between the API implementations Coordinate between the API implementations using a Saga pattern Implement caching in each API implementation to improve performance
- * Implement local transactions within each API implementation

Configure each API implementation to also participate in the same eXtended Architecture (XA) transaction Implement caching in each API implementation to improve performance

- * Implement local transactions in each API implementation

Coordinate between the API implementations using a Saga pattern

Apply various compensating actions depending on where a failure occurs

- * Implement an eXtended Architecture (XA) transaction manager in a Mule application using a Saga pattern Connect each API implementation with the Mule application using XA transactions Apply various compensating actions depending on where a failure occurs

QUESTION 16

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

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

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

- * Anypoint Platform Private Cloud Edition
- * Anypoint Runtime Fabric
- * CloudHub
- * Anypoint Service Mesh

QUESTION 17

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

How can these requirements be met?

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

QUESTION 18

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 message broker
- * An HTTP Load Balancer
- * A database
- * An Object Store

Explanation

Correct answer is An HTTP Load Balancer.

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

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

QUESTION 19

As a part of design , Mule application is required call the Google Maps API to perform a distance computation. The application is deployed to cloudhub.

At the minimum what should be configured in the TLS context of the HTTP request configuration to meet these requirements?

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

QUESTION 20

A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should be published in a way that is guaranteed in normal situations and also minimizes duplicate delivery in less frequent failure scenarios.

The organizational firewall is configured to only allow outbound traffic on ports 80 and 443. Some external event consumers are within the organizational network, while others are located outside the firewall.

What Anypoint Platform service is most idiomatic (used for its intended purpose) for publishing these Mule events to all external consumers while addressing the desired reliability goals?

- * CloudHub VM queues
- * Anypoint MQ
- * Anypoint Exchange
- * CloudHub Shared Load Balancer

Explanation

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

QUESTION 21

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?

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

QUESTION 22

One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make.

Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments.

The production environment however does not have any rate limiting.

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

- * Including logic within the API implementation that bypasses in locations of the back-end system in the staging environment and invoke a Mocking service that replicates typical back-end system responses Then conduct performance test using this API implementation
- * Use MUnit to simulate standard responses from the back-end system.

Then conduct performance test to identify other bottlenecks in the system

- * Create a Mocking service that replicates the back-end system

production performance characteristics

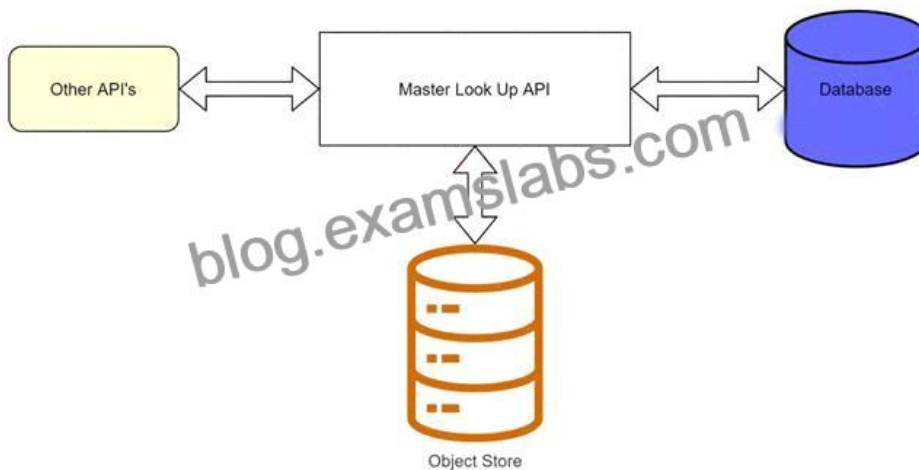
Then configure the API implementation to use the mocking

service and conduct the performance test

* Conduct scaled-down performance tests in the staging environment against rate-limiting back-end system. Then upscale performance results to full production scale

QUESTION 23

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 cloudhub 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 they 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?

- * Implement HTTP caching policy for all GET endpoints for the master lookup API and implement locking to synchronize access to object store
- * Upgrade vCore size from 0.1 vCore to 0,2 vCore
- * Implement HTTP caching policy for all GET endpoints for master lookup API
- * Add an additional Cloudhub worker to provide additional capacity

QUESTION 24

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

- * A Mule application that accepts requests over HTTP/1x

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

Explanation

- * 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

QUESTION 25

An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events What are the two considerations that must be kept in mind while using test recorder (Choose two answers)

- * Tests for flows cannot be created with Mule errors raised inside

the flow or already existing in the incoming event

- * Recorder supports smoking a message before or inside a ForEach processor
- * The recorder support loops where the structure of the data been tested changes inside the iteration
- * A recorded flow execution ends successfully but the result does

not reach its destination because the application is killed

- * Mocking values resulting from parallel processes are possible and will not affect the execution of the processes that follow in the test

QUESTION 26

An organization designing a hybrid, load balanced, single cluster production environment. Due to performance service level agreement goals, it is looking into running the Mule applications in an active-active multi node cluster configuration.

What should be considered when running its Mule applications in this type of environment?

- * All event sources, regardless of time , can be configured as the target source by the primary node in the cluster
- * An external load balancer is required to distribute incoming requests throughout the cluster nodes
- * A Mule application deployed to multiple nodes runs in an isolation from the other nodes in the cluster
- * Although the cluster environment is fully installed configured and running, it will not process any requests until an outage condition is detected by the primary node in the cluster.

QUESTION 27

What condition requires using a CloudHub Dedicated Load Balancer?

- * When cross-region load balancing is required between separate deployments of the same Mule application
- * When custom DNS names are required for API implementations deployed to customer-hosted Mule runtimes
- * When API invocations across multiple CloudHub workers must be load balanced
- * When server-side load-balanced TLS mutual authentication is required between API implementations and API clients

Explanation

Correct answer is When server-side load-balanced TLS mutual authentication is required between API implementations and API clients CloudHub dedicated load balancers (DLBs) are an optional component of Anypoint Platform that enable you to route external HTTP and HTTPS traffic to multiple Mule applications deployed to CloudHub workers in a Virtual Private Cloud (VPC). Dedicated load balancers enable you to: * Handle load balancing among the different CloudHub workers that run your application. * Define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication. * Configure proxy rules that map your applications to custom domains. This enables you to host your applications under a single domain

QUESTION 28

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

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

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

QUESTION 29

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?

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

QUESTION 30

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.

- * VM
- * Anypoint MQ
- * SFTP
- * ObjectStore

Explanation

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

Here is additional information about Transaction management:

Table Description automatically generated

	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}
Fasten	Round Robin	Round Robin
Supports HTTPS Protocol	Yes	Yes
Worker Assignment	No	Yes
IP Blacklisting/Whitelisting	No https://docs.mulesoft.com/runtime-manager/lb-whitelists	Yes
Configure Custom Domain	No	Yes
Custom Certificate	No	Yes
Rate Limit	Lower Rate Limit and applied According to Region	Higher Rate Limit Threshold
VPC	Anypoint VPC optional	Can't Use DLB without Anypoint VPC

QUESTION 31

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?

- * 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
- * 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

- * 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
- * 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

QUESTION 32

A company is designing a mule application to consume batch data from a partner's ftps server The data files have been compressed and then digitally signed using PGP.

What inputs are required for the application to securely consumed these files?

- * ATLS context Key Store requiring the private key and certificate for the company PGP public key of partner PGP private key for the company
- * ATLS context first store containing a public certificate for partner ftps server and the PGP public key of the partner TLS contact Key Store containing the FTP credentials
- * TLS context trust or containing a public certificate for the ftps server The FTP username and password The PGP public key of the partner
- * The PGP public key of the partner

The PGP private key for the company The FTP username and password

QUESTION 33

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.

- * No changes need to be made to POM.xml file and CI/CD script should be modified as per the RTF configurations
- * 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
- * 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

- * runtimeFabric profile should be added mule configuration files in the mule applications and CI/CD script should be modified as per the RTF configurations

QUESTION 34

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?

- * It is not possible to store logs for 2 years in CloudHub deployment. External log management system is required.
- * When deploying an application to CloudHub , logs retention period should be selected as 2 years

- * When deploying an application to CloudHub, worker size should be sufficient to store 2 years data
- * Logging strategy should be configured accordingly in log4j file deployed with the application.

Explanation

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

QUESTION 35

An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.

What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

- * Client management configured in access management
- * Identity management configured in access management
- * Enterprise Security module coded in Mule applications
- * External access configured in API Manager

QUESTION 36

An organization is evaluating using the CloudHub shared Load Balancer (SLB) vs creating a CloudHub dedicated load balancer (DLB). They are evaluating how this choice affects the various types of certificates used by CloudHub deployed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates. What type of restrictions exist on the types of certificates for the service that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

- * Underlying Mule applications need to implement own certificates
- * Only MuleSoft provided certificates can be used for server side certificate
- * Only self signed certificates can be used
- * All certificates which can be used in shared load balancer need to get approved by raising support ticket

Explanation

Correct answer is Only MuleSoft provided certificates can be used for server side certificate

- * The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
- * You would need to use dedicated load balancer which can enable you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
- * 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.

Additional Info on SLB Vs DLB:

Table Description automatically generated

	Shared Load Balancer	Dedicated Load Balancer
VPC	Shared VPC (Mulesoft)	VPC (Customer)
Default Load Balancer	Cloudhub provides Deault Shared Load Balancer available in All Environment	Need to Purchase
Organization Use	Multiple Oragnization	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)
Fastlane	Round Robin	Round Robin
Supports HTTPS Protocol	Yes	Yes
Worker Assignment	No	Yes
IP Blacklisting/Whitelisting	No	Yes
	https://docs.mulesoft.com/runtime-manager/ib-whitelists	
Configure Custom Domain	No	Yes
Custom Certificate	No	Yes
Rate Limit	Lower Rate Limit and applied According to Region	Higher Rate Limit Threshold
VPC	Anypoint VPC optional	Can't Use DLB without Anypoint VPC

MCIA-Level-1-Maintenance Premium PDF & Test Engine Files with 118 Questions & Answers:

<https://www.examslabs.com/MuleSoft/MuleSoft-Certified-Architect/best-MCIA-Level-1-Maintenance-exam-dumps.html>