


Enhance Your Career With Available Preparation Guide for MCIA-Level-1 Exam [Q55-Q70]



Enhance Your Career With Available Preparation Guide for MCIA-Level-1 Exam [Q55-Q70]

Enhance Your Career With Available Preparation Guide for MCIA-Level-1 Exam Get Special Discount Offer of MCIA-Level-1 Certification Exam Sample Questions and Answers

How to Prepare For MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam **Preparation Guide for MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam Introduction**

MuleSoft has a unique community. Individuals can go ahead and appear for certification which they find suitable as per their expertise and skillset. MuleSoft also recognizes that the community is an important way to engage with its customer base. Certification is evidence of your skills, expertise in those areas in which you like to work. If a candidate wants to work on MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 and prove his knowledge, Certification offered by MuleSoft. This MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Individuals Qualification Certification helps a candidate to validates his skills in MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1. In this guide, we will cover **mcia-level-1 practice exams** and all the aspects of **mcia-level-1 exam dumps**.

QUESTION 55

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?

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

QUESTION 56

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?

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

QUESTION 57

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

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

QUESTION 58

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?

* 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 sent directly to the Fulfillment microservices

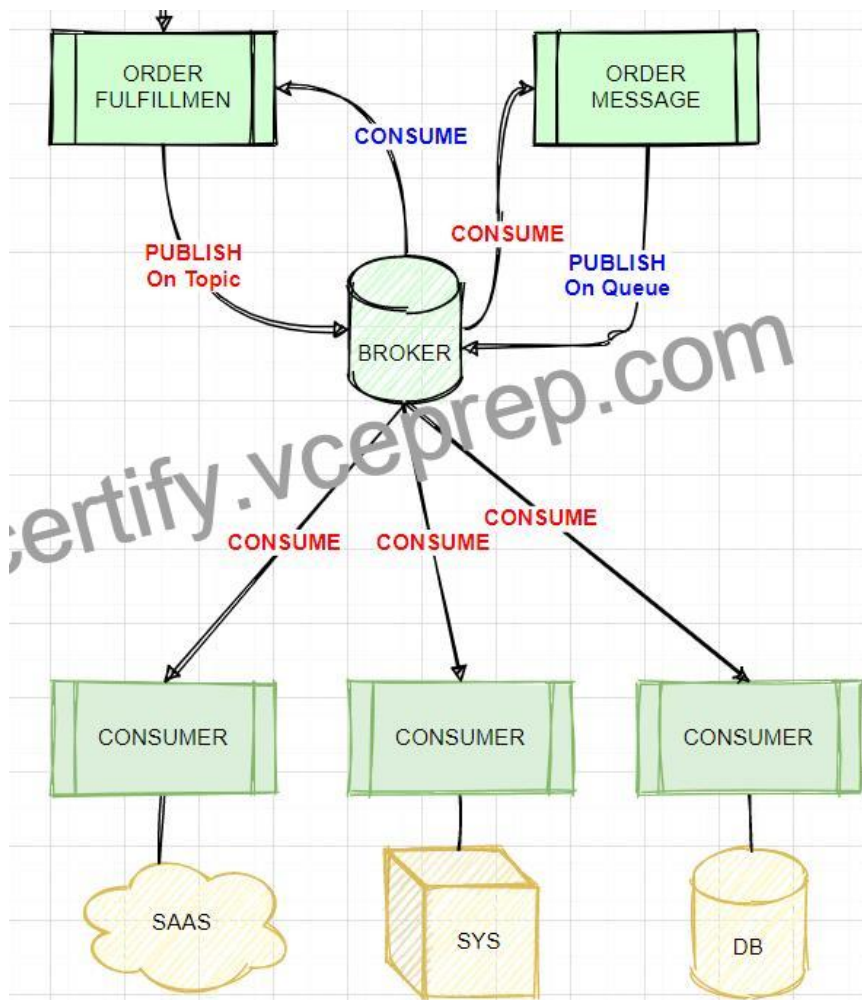
OrderFulfilled messages are sent directly to the Order microservice

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

* 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

* 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

* If you need to scale a JMS provider/ message broker, you can 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 the best choice. * It is 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 a 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 needs to be published on Topic using same broker. * Correct answer:



QUESTION 59

An organization is implementing a Quote of the Day API that caches today's quote. What scenario can use the CloudHub Object Store connector to persist the cache's state?


- * When there is one deployment of the API implementation to CloudHub and another one to customer hosted mule runtime that must share the cache state.
- * When there are two CloudHub deployments of the API implementation by two Anypoint Platform business groups to the same CloudHub region that must share the cache state.
- * When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.
- * When there are three CloudHub deployments of the API implementation to three separate CloudHub regions that must share the cache state.

Object Store Connector is a Mule component that allows for simple key-value storage. Although it can serve a wide variety of use cases, it is mainly design for:
• Storing synchronization information, such as watermarks.
• Storing temporal information such as access tokens.
• Storing user information. Additionally, Mule Runtime uses Object Stores to support some of its own components, for example:
• The Cache module uses an Object Store to maintain all of the cached data.
• The OAuth module (and every OAuth enabled connector) uses Object Stores to store the access and refresh tokens. Object Store data is in the same region as the worker where the app is initially deployed. For example, if you deploy to the Singapore region, the object store persists in the Singapore region. MuleSoft Reference : <https://docs.mulesoft.com/object-store-connector/1.1/> Data can be shared between different instances of the Mule application. This is not recommended for Inter Mule app communication. Coming to the question, object store cannot be used to share cached data if it is deployed as separate Mule

applications or deployed under separate Business Groups. Hence correct answer is When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.

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 Salesforce to a backend system every 5 minutes.

A watermark will be used to only retrieve those Salesforce 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?

- * Persistent Anypoint MQ Queue
- * Persistent Object Store
- * Persistent Cache Scope
- * Persistent VM Queue

QUESTION 61

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?

- * An On Error scope

MuleSoft Object Store

ActiveMQ Dead Letter Queue for manual processing

- * Until Successful component

ActiveMQ long retry Queue

ActiveMQ Dead Letter Queue for manual processing

- * Until Successful component

MuleSoft Object Store

ActiveMQ is NOT needed or used

- * An On Error scope

Non-persistent VM

ActiveMQ Dead Letter Queue for manual processing

QUESTION 62

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

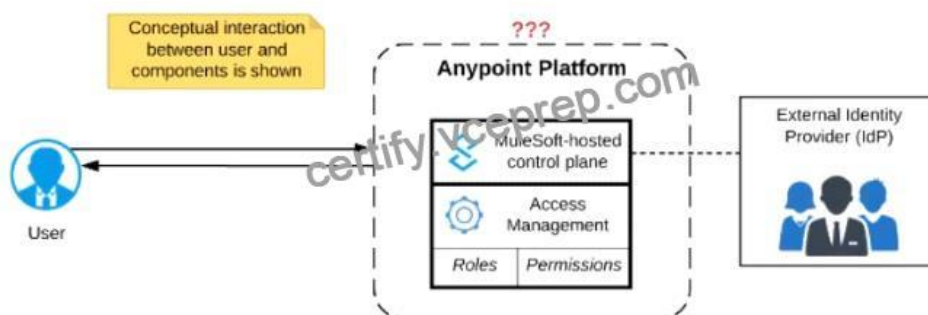
What is the main benefit of using XA transaction over local transactions to synchronize these two database system?

- * Reduce latency
- * Increase throughput
- * Simplifies communication
- * Ensure consistency
- * XA transaction add tremendous latency so Reduce Latency; is incorrect option XA transactions define All or No; commit protocol.
- * Each local XA resource manager supports the A.C.I.D properties (Atomicity, Consistency, Isolation, and Durability).

So correct choice is Ensure consistency;

QUESTION 63

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?

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

QUESTION 64

What API policy would LEAST likely be applied to a Process API?

- * Custom circuit breaker
- * Client ID enforcement
- * Rate limiting
- * JSON threat protection

Key to this question lies in the fact that Process API are not meant to be accessed directly by clients. Lets analyze options one by one. Client ID enforcement : This is applied at process API level generally to ensure that identity of API clients is always known and available for API-based analytics Rate Limiting : This policy is applied on Process Level API to secure API's against degradation of service that can happen in case load received is more than it can handle Custom circuit breaker : This is also quite useful feature on process level API's as it saves the API client the wasted time and effort of invoking a failing API. JSON threat protection : This policy is not required at Process API and rather implemented as Experience API's. This policy is used to safeguard application from malicious attacks by injecting malicious code in JSON object. As ideally Process API's are never called from external world , this policy is never used on Process API's Hence correct answer is JSON threat protection MuleSoft Documentation Reference : <https://docs.mulesoft.com/api-manager/2.x/policy-mule3-json-threat>

QUESTION 65

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?

- * Keep the default logging configuration in RuntimeManager

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

- * Disable Cloudhub logging in Runtime Manager

Define the splunk appender in EACH Mule application's log4j2.xml file

- * Disable Cloudhub logging in Runtime Manager

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

- * Keep the default logging configuration in Runtime Manager

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

QUESTION 66

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?

- * 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
- * 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
- * Use Anypoint Exchange to register the various Mule applications and APIs and share the RAML definitions with the stakeholders, so they can be discovered
- * 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

QUESTION 67

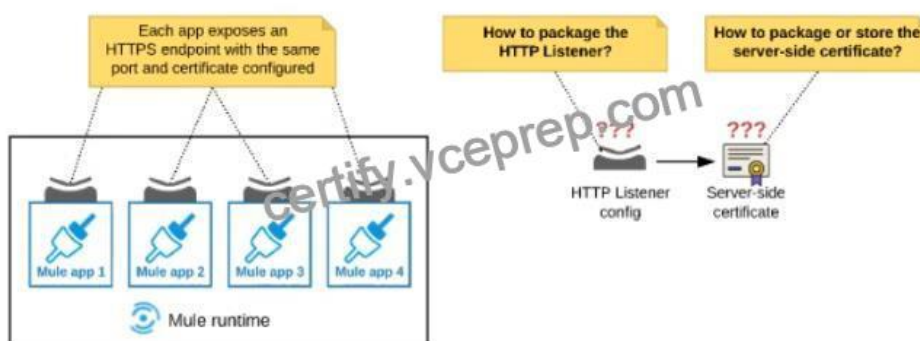
An organization is evaluating using the CloudHub shared Load Balancer (SLB) vs creating a CloudHub dedicated load balancer (DLB). They are evaluating how this choice affects the various types of certificates used by CloudHub deployed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates.

What type of restrictions exist on the types of certificates that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

- * Only MuleSoft-provided certificates are exposed.
- * Only customer-provided wildcard certificates are exposed.
- * Only customer-provided self-signed certificates are exposed.
- * Only underlying Mule application certificates are exposed (pass-through)

QUESTION 68

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?

- * 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
- * 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
- * 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
- * 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

QUESTION 69

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?

- * Keep the default logging configuration in RuntimeManager

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

- * Disable Cloudhub logging in Runtime Manager

Define the splunk appender in EACH Mule application's log4j2.xml file

- * Disable Cloudhub logging in Runtime Manager

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

- * Keep the default logging configuration in Runtime Manager

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

QUESTION 70

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

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

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

- * The request was ACCEPTED by the external API
- * The request was REDIRECTED to a different URL by the external API
- * The request was NOT RECEIVED by the external API
- * The request was NOT ACCEPTED by the external API

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

Who should take the MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam

The MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 exam certification is an internationally recognized certification that helps to have validation for those professionals who are keen to make their career in MuleSoft design, build, test and debug, deploy, and manage basic APIs and integrations.

If a candidate/professional seeks a powerful improvement in career growth needs enhanced knowledge, skills, and talents. The MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 certification provides proof of this advanced knowledge and skill.

MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam Certified Professional salary

The average salary of a MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam Certified Expert in:

- England - 75,000 POUND- India. - 14,00,327 INR- United State - 100,200 USD- Europe - 70,500 EURO **Updated**

MCIA-Level-1 Dumps Questions Are Available For Passing MuleSoft Exam:

<https://www.vceprep.com/MCIA-Level-1-latest-vce-prep.html>]