

MCIA-LEVEL-1^{Q&As}

MuleSoft Certified Integration Architect - Level 1

Pass Mulesoft MCIA-LEVEL-1 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/mcia-level-1.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Mulesoft
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

In Anypoint Platform, a company wants to configure multiple identity providers(Idps) for various lines of business (LOBs) Multiple business groups and environments have been defined for the these LOBs. What Anypoint Platform feature can use multiple Idps access the company\'s business groups and environment?

- A. User management
- B. Roles and permissions
- C. Dedicated load balancers
- D. Client Management

Correct Answer: D

Explanation:

Correct answer is Client Management

*

Anypoint Platform acts as a client provider by default, but you can also configure external client providers to authorize client applications.

*

As an API owner, you can apply an OAuth 2.0 policy to authorize client applications that try to access your API. You need an OAuth 2.0 provider to use an OAuth 2.0 policy.

*

You can configure more than one client provider and associate the client providers with different environments. If you configure multiple client providers after you have already created environments, you can associate the new client providers with the environment.

*

You should review the existing client configuration before reassigning client providers to avoid any downtime with existing assets or APIs.

*

When you delete a client provider from your master organization, the client provider is no longer available in environments that used it.

*

Also, assets or APIs that used the client provider can no longer authorize users who want to access them.

-----MuleSoft

Reference: <https://docs.mulesoft.com/access-management/managing-api-clients>
<https://www.folkstalk.com/2019/11/mulesoft-integration-and-platform.html>

QUESTION 2

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. 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

Correct Answer: D

Explanation: 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>

QUESTION 3

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

- A. A Cloud-based
- B. No-code
- C. Code-first

D. On-premises

Correct Answer: A

QUESTION 4

An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders.

The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS).

At a minimum, what must be stored in each truststore and keystore of these two Mule applications to properly support two-way TLS between the two Mule applications while properly protecting each Mule application's keys?

- A. Orders API truststore: The Orders API public key Process Orders keystore: The Process Orders private key and public key
- B. Orders API truststore: The Orders API private key and public key Process Orders keystore: The Process Orders private key public key
- C. Orders API truststore: The Process Orders public key Orders API keystore: The Orders API private key and public key Process Orders truststore: The Orders API public key Process Orders keystore: The Process Orders private key and public key
- D. Orders API truststore: The Process Orders public key Orders API keystore: The Orders API private key Process Orders truststore: The Orders API public key Process Orders keystore: The Process Orders private key

Correct Answer: C

Reference: <https://www.caeliusconsulting.com/blogs/one-way-and-two-way-tls-and-their-implementation-in-mulesoft/>

QUESTION 5

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

Correct Answer: C

Explanation: * "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 can't 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

QUESTION 6

A team would like to create a project skeleton that developers can use as a starting point when creating API Implementations with Anypoint Studio. This skeleton should help drive consistent use of best practices within the team.

What type of Anypoint Exchange artifact(s) should be added to Anypoint Exchange to publish the project skeleton?

- A. A custom asset with the default API implementation
- B. A RAML archetype and reusable trait definitions to be reused across API implementations
- C. An example of an API implementation following best practices
- D. a Mule application template with the key components and minimal integration logic

Correct Answer: D

Explanation: * Sharing Mule applications as templates is a great way to share your work with other people who are in your organization in Anypoint Platform. When they need to build a similar application they can create the mule application using the template project from Anypoint studio.

*

Anypoint Templates are designed to make it easier and faster to go from a blank canvas to a production application. They're bit for bit Mule applications requiring only Anypoint Studio to build and design, and are deployable both on-premises and in the cloud.

*

Anypoint Templates are based on five common data Integration patterns and can be customized and extended to fit your integration needs. So even if your use case involves different endpoints or connectors than those included in the template, they still offer a great starting point. Some of the best practices while creating the template project: - Define the common error handler as part of template project, either using pom dependency or mule config file - Define common logger/audit framework as part of the template project - Define the env specific properties and secure properties file as per the requirement - Define global.xml for global configuration - Define the config file for connector configuration like Http,Salesforce,File,FTP etc - Create separate folders to create DWL,Properties,SSL certificates etc - Add the dependency and configure the pom.xml as per the business need

-Configure the mule-artifact.json as per the business need

QUESTION 7

According to MuleSoft's recommended REST conventions, which HTTP method should an API use to specify how API clients can request data from a specified resource?

- A. POST
- B. PUT
- C. PATCH D. GET

Correct Answer: D

QUESTION 8

A marketing organization is designing a Mule application to process campaign data. The Mule application will periodically check for a file in a SFTP location and process the records in the file. The size of the file can vary from 10MB to 5GB. Due to the limited availability of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.

The application must transform and send different formats of this file to three different downstream SFTP locations.

What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

- A. Use an in-memory repeatable stream
- B. Use a file-stored non-repeatable stream
- C. Use an in-memory non-repeatable stream
- D. Use a file-stored repeatable stream

Correct Answer: A

Reference: <https://docs.mulesoft.com/mule-runtime/4.4/streaming-about>

QUESTION 9

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. Use the scaffolding capability of Anypoint Studio to create an API portal based on the API specification
- C. Publish the API specification to Exchange and solicit feedback from the API's consumers
- D. Use API Manager to version the API specification

Correct Answer: C

QUESTION 10

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

Correct Answer: D

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.

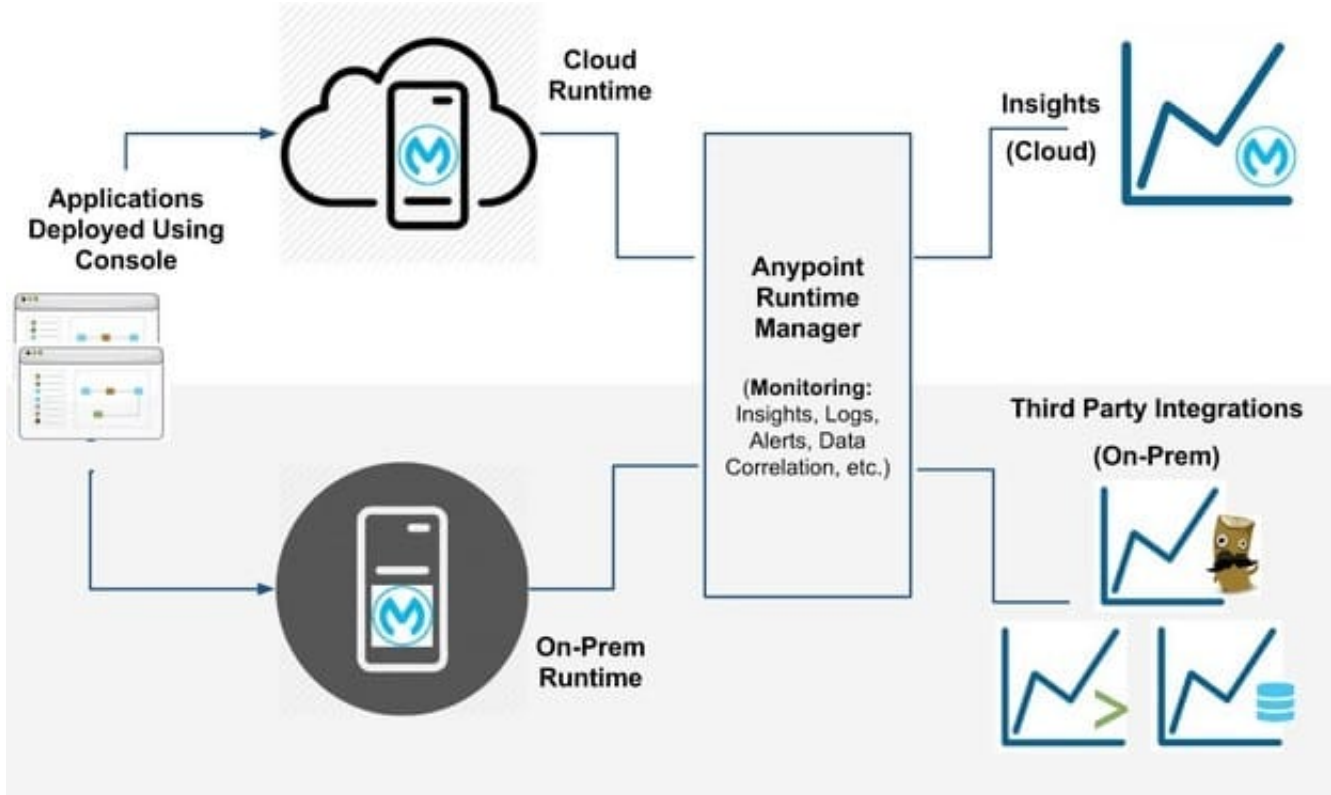
* You can configure the Runtime Manager agent to export data to external analytics tools. Using either the Runtime Manager cloud console or Anypoint Platform Private Cloud Edition, you can:

--> 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:

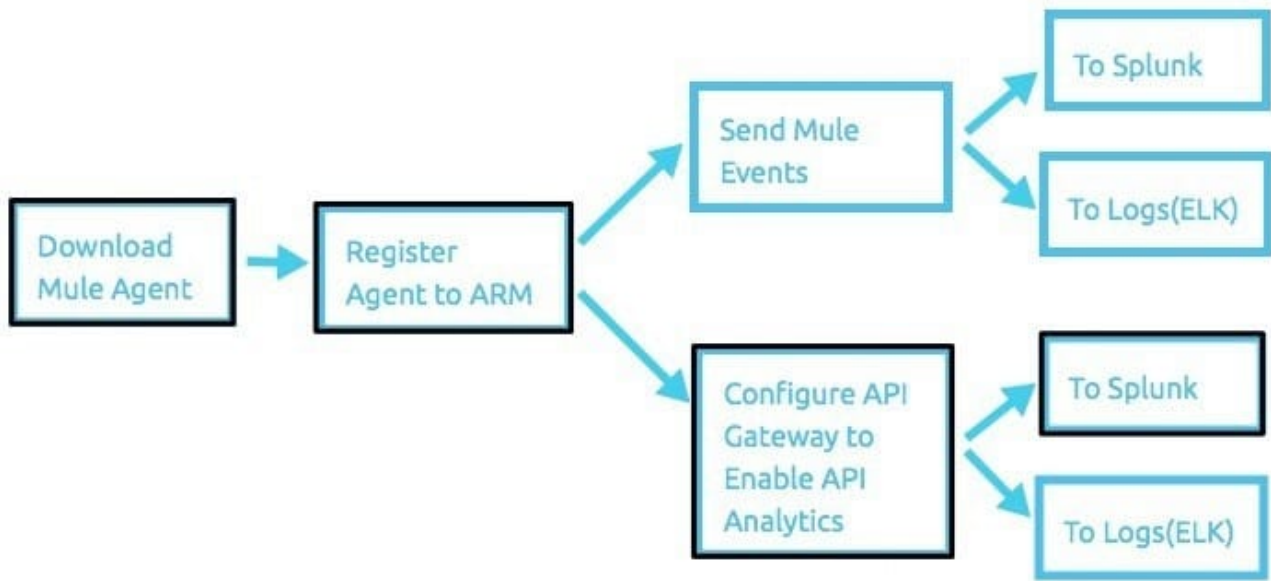
<https://docs.mulesoft.com/runtime-manager/> <https://docs.mulesoft.com/release-notes/runtime-manager-agent/runtime-manager-agentrelease-notes>



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.



Diagram

QUESTION 11

An application load balancer routes requests to a RESTful web API secured by Anypoint Flex Gateway.

Which protocol is involved in the communication between the load balancer and the Gateway?

- A. SFTP
- B. HTTPS
- C. LDAP
- D. SMTP

Correct Answer: B

QUESTION 12

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

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

Correct Answer: C

QUESTION 13

An organization has implemented a continuous integration (CI) lifecycle that promotes Mule applications through code, build, and test stages. To standardize the organization's CI journey, a new dependency control approach is being designed to store artifacts that include information such as dependencies, versioning, and build promotions.

To implement these process improvements, the organization will now require developers to maintain all dependencies related to Mule application code in a shared location.

What is the most idiomatic (used for its intended purpose) type of system the organization should use in a shared location to standardize all dependencies related to Mule application code?

A. A MuleSoft-managed repository at repository.mulesoft.org

B. A binary artifact repository

C. API Community Manager

D. The Anypoint Object Store service at cloudhub.io

Correct Answer: C

QUESTION 14

How are the API implementation , API client, and API consumer combined to invoke and process an API ?

A. The API consumer creates an API implementation , which receives API invocations from an API such that they are processed for an API client

B. The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation

C. An API client creates an API consumer, which receives API invocation from an API such that they are processed for an API implementation

D. The API client creates an API consumer which sends API invocations to an API such that they are processed by API implementation

Correct Answer: C

Explanation: The API consumer creates an API client which sends API invocations to an API such that they are processed by an API implementation This is based on below definitions API client • An application component • that accesses a service • by invoking an API of that service - by definition of the term API over HTTP API consumer • A business role, which is often assigned to an individual • that develops API clients, i.e., performs the activities necessary for enabling an API client to invoke APIs API implementation • An application component • that implements the

functionality

QUESTION 15

An organization is struggling with frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects.

How can these best practices be achieved with the applications having the least amount of responsibility?

- A. Create a Mule plugin project with all the dependencies and add it as a dependency in each application's POM.xml file
- B. Create a mule domain project with all the dependencies defined in its POM.xml file and add each application to the domain Project
- C. Add all dependencies in each application's POM.xml file
- D. Create a parent POM of all the required dependencies and reference each in each application's POM.xml file

Correct Answer: D

[Latest MCIA-LEVEL-1 Dumps](#)

[MCIA-LEVEL-1 PDF Dumps](#) [MCIA-LEVEL-1 Braindumps](#)