


Ultimate Guide to the MCPA-Level-1 - Latest Oct 02, 2024 Edition Available Now [Q47-Q61]



Ultimate Guide to the MCPA-Level-1 - Latest Oct 02, 2024 Edition Available Now [Q47-Q61]

Ultimate Guide to the MCPA-Level-1 - Latest Oct 02, 2024 Edition Available Now
2024 Updated Verified Pass MCPA-Level-1 Exam - Real Questions and Answers

The MCPA-Level-1 certification exam is a valuable credential for architects who want to demonstrate their expertise with MuleSoft's Anypoint Platform. MuleSoft Certified Platform Architect - Level 1 certification validates an individual's ability to design and build integration solutions that meet the complex needs of modern businesses. It also demonstrates a commitment to ongoing education and professional development within the field of enterprise integration.

NEW QUESTION 47

What Mule application deployment scenario requires using Anypoint Platform Private Cloud Edition or Anypoint Platform for Pivotal Cloud Foundry?

- * When it is required to make ALL applications highly available across multiple data centers
- * When it is required that ALL APIs are private and NOT exposed to the public cloud
- * When regulatory requirements mandate on-premises processing of EVERY data item, including meta-data
- * When ALL backend systems in the application network are deployed in the organization's intranet

NEW QUESTION 48

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 APIs public portal.

The API endpoint does NOT change in the new version. How should the developer of an API client respond to this change?

- * The API producer should be requested to run the old version in parallel with the new one
- * The API producer should be contacted to understand the change to existing functionality
- * The API client code only needs to be changed if it needs to take advantage of the new features
- * The API clients need to update the code on their side and need to do full regression

NEW QUESTION 49

In which layer of API-led connectivity, does the business logic orchestration reside?

- * System Layer
- * Experience Layer
- * Process Layer

Process Layer

>> Experience layer is dedicated for enrichment of end user experience. This layer is to meet the needs of different API clients/ consumers.

>> System layer is dedicated to APIs which are modular in nature and implement/ expose various individual functionalities of backend systems

>> Process layer is the place where simple or complex business orchestration logic is written by invoking one or many System layer modular APIs So, Process Layer is the right answer.

NEW QUESTION 50

What CANNOT be effectively enforced using an API policy in Anypoint Platform?

- * Guarding against Denial of Service attacks
- * Maintaining tamper-proof credentials between APIs
- * Logging HTTP requests and responses
- * Backend system overloading

Correct answer: Guarding against Denial of Service attacks

>> Backend system overloading can be handled by enforcing [Spike Control Policy](#);

>> Logging HTTP requests and responses can be done by enforcing [Message Logging Policy](#);

>> Credentials can be tamper-proofed using [Security](#); and [Compliance](#); Policies However, unfortunately, there is no proper way currently on Anypoint Platform to guard against DOS attacks.

NEW QUESTION 51

An organization wants MuleSoft-hosted runtime plane features (such as HTTP load balancing, zero downtime, and horizontal and vertical scaling) in its Azure environment. What runtime plane minimizes the organization's effort to achieve these features?

- * Anypoint Runtime Fabric
- * Anypoint Platform for Pivotal Cloud Foundry
- * CloudHub
- * A hybrid combination of customer-hosted and MuleSoft-hosted Mule runtimes

Correct answer: Anypoint Runtime Fabric

>> When a customer is already having an Azure environment, It is not at all an ideal approach to go with hybrid model having some Mule Runtimes hosted on Azure and some on MuleSoft. This is unnecessary and useless.

>> CloudHub is a Mulesoft-hosted Runtime plane and is on AWS. We cannot customize to point CloudHub to customer's Azure environment.

>> Anypoint Platform for Pivotal Cloud Foundry is specifically for infrastructure provided by Pivotal Cloud Foundry

>> Anypoint Runtime Fabric is right answer as it is a container service that automates the deployment and orchestration of Mule applications and API gateways. Runtime Fabric runs within a customer-managed infrastructure on AWS, Azure, virtual machines (VMs), and bare-metal servers.

-Some of the capabilities of Anypoint Runtime Fabric include:

-Isolation between applications by running a separate Mule runtime per application.

-Ability to run multiple versions of Mule runtime on the same set of resources.

-Scaling applications across multiple replicas.

-Automated application fail-over.

-Application management with Anypoint Runtime Manager.

NEW QUESTION 52

What best explains the use of auto-discovery in API implementations?

- * It makes API Manager aware of API implementations and hence enables it to enforce policies
- * It enables Anypoint Studio to discover API definitions configured in Anypoint Platform
- * It enables Anypoint Exchange to discover assets and makes them available for reuse
- * It enables Anypoint Analytics to gain insight into the usage of APIs

Explanation

<https://docs.mulesoft.com/api-manager/2.x/api-auto-discovery-new-concept>

NEW QUESTION 53

A System API is designed to retrieve data from a backend system that has scalability challenges.

What API policy can best safeguard the backend system?

- * IP whitelist
- * SLA-based rate limiting
- * OAuth 2 token enforcement
- * Client ID enforcement

Explanation/Reference: <https://dzone.com/articles/how-to-secure-apis>

NEW QUESTION 54

Which layer in the API-led connectivity focuses on unlocking key systems, legacy systems, data sources etc and exposes the functionality?

- * Experience Layer
- * Process Layer
- * System Layer

System Layer



The APIs used in an API-led approach to connectivity fall into three categories:

System APIs; these usually access the core systems of record and provide a means of insulating the user from the complexity or any changes to the underlying systems. Once built, many users, can access data without any need to learn the underlying systems and can reuse these APIs in multiple projects.

Process APIs; These APIs interact with and shape data within a single system or across systems (breaking down data silos) and are created here without a dependence on the source systems from which that data originates, as well as the target channels through which that data is delivered.

Experience APIs; Experience APIs are the means by which data can be reconfigured so that it is most easily consumed by its intended audience, all from a common data source, rather than setting up separate point-to-point integrations for each channel. An Experience API is usually created with API-first design principles where the API is designed for the specific user experience in mind.

NEW QUESTION 55

Due to a limitation in the backend system, a system API can only handle up to 500 requests per second. What is the best type of API

policy to apply to the system API to avoid overloading the backend system?

- * Rate limiting
- * HTTP caching
- * Rate limiting & SLA based
- * Spike control

Spike control

>> First things first, HTTP Caching policy is for purposes different than avoiding the backend system from overloading. So this is OUT.

>> Rate Limiting and Throttling/ Spike Control policies are designed to limit API access, but have different intentions.

>> Rate limiting protects an API by applying a hard limit on its access.

>> Throttling/ Spike Control shapes API access by smoothing spikes in traffic.

That is why, Spike Control is the right option.

NEW QUESTION 56

An organization is implementing a Quote of the Day API that caches today's quote.

- * What scenario can use the GoudHub Object Store via the Object Store connector to persist the cache's state?
- * When there are three CloudHub deployments of the API implementation to three separate CloudHub regions 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 deployment of the API implementation to CloudHub and anottV deployment to a customer-hosted Mule runtime that must share the cache state
- * When there is one CloudHub deployment of the API implementation to three CloudHub workers that must share the cache state

NEW QUESTION 57

A System API is designed to retrieve data from a backend system that has scalability challenges. What API policy can best safeguard the backend system?

- * IPwhitelist
- * SLA-based rate limiting
- * Auth 2 token enforcement
- * Client ID enforcement

SLA-based rate limiting

>> Client Id enforcement policy is a Compliance related NFR and does not help in maintaining the Quality of Service (QoS). It CANNOT and NOT meant for protecting the backend systems from scalability challenges.

>> IP Whitelisting and OAuth 2.0 token enforcement are Security related NFRs and again does not help in maintaining the Quality of Service (QoS). They CANNOT and are NOT meant for protecting the backend systems from scalability challenges.

Rate Limiting, Rate Limiting-SLA, Throttling, Spike Control are the policies that are related to Quality of Service (QoS); related NFRs and are meant to help in protecting the backend systems from getting overloaded.

<https://dzone.com/articles/how-to-secure-apis>

NEW QUESTION 58

An organization makes a strategic decision to move towards an IT operating model that emphasizes consumption of reusable IT assets using modern APIs (as defined by MuleSoft).

What best describes each modern API in relation to this new IT operating model?

- * Each modern API has its own software development lifecycle, which reduces the need for documentation and automation.
- * Each modern API must be treated like a product and designed for a particular target audience (for instance, mobile app developers)
- * Each modern API must be easy to consume, so should avoid complex authentication mechanisms such as SAML or JWT.
- * Each modern API must be REST and HTTP based.

NEW QUESTION 59

When designing an upstream API and its implementation, the development team has been advised to NOT set timeouts when invoking a downstream API, because that 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?

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

NEW QUESTION 60

Traffic is routed through an API proxy to an API implementation. The API proxy is managed by API Manager and the API implementation is deployed to a CloudHub VPC using Runtime Manager. API policies have been applied to this API. In this deployment scenario, at what point are the API policies enforced on incoming API client requests?

- * At the API proxy
- * At the API implementation
- * At both the API proxy and the API implementation
- * At a MuleSoft-hosted load balancer

Correct answer: At the API proxy

>> API Policies can be enforced at two places in Mule platform.

>> One; As an Embedded Policy enforcement in the same Mule Runtime where API implementation is running.

>> Two – On an API Proxy sitting in front of the Mule Runtime where API implementation is running.

>> As the deployment scenario in the question has API Proxy involved, the policies will be enforced at the API Proxy.

NEW QUESTION 61

When designing an upstream API and its implementation, the development team has been advised to NOT set timeouts when invoking a downstream API, because that 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?
- * An SLA for the upstream API CANNOT be provided
- * The invocation of the downstream API will run to completion without timing out
- * A default timeout of 500 ms will automatically be applied by the Mule runtime in which the upstream API implementation executes
- * A toad-dependent timeout of less than 1000 ms will be applied by the Mule runtime in which the downstream API implementation executes

MuleSoft MCPA-Level-1 (MuleSoft Certified Platform Architect - Level 1) Certification Exam is a globally recognized certification that validates an individual's expertise in designing, building, and managing MuleSoft APIs and integrations. MuleSoft Certified Platform Architect - Level 1 certification is designed for architects and developers who want to demonstrate their skills in the MuleSoft Anypoint Platform.

Dumps Moneyack Guarantee - MCPA-Level-1 Dumps Approved Dumps:

<https://www.examlabs.com/MuleSoft/MuleSoft-Certified-Platform-Architect/best-MCPA-Level-1-exam-dumps.html>