

MuleSoft MCPA Level 1

**MULESOFT PLATFORM ARCHITECT LEVEL 1 CERTIFICATION
QUESTIONS & ANSWERS**

Exam Summary – Syllabus – Questions

MCPA LEVEL 1

MuleSoft Certified Platform Architect - Level 1 (MCPA)

60 Questions Exam – 70% Cut Score – Duration of 120 minutes

www.CertFun.com

Table of Contents

Know Your MCPA Level 1 Certification Well:	2
MuleSoft MCPA Level 1 Platform Architect Level 1 Certification Details:	2
MCPA Level 1 Syllabus:	3
MuleSoft MCPA Level 1 Sample Questions:	5
Study Guide to Crack MuleSoft Platform Architect Level 1 MCPA Level 1 Exam:	9

Know Your MCPA Level 1 Certification Well:

The MCPA Level 1 is best suitable for candidates who want to gain knowledge in the MuleSoft Architect. Before you start your MCPA Level 1 preparation you may struggle to get all the crucial Platform Architect Level 1 materials like MCPA Level 1 syllabus, sample questions, [study guide](#).

But don't worry the MCPA Level 1 PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the MCPA Level 1 syllabus?
- How many questions are there in the MCPA Level 1 exam?
- Which Practice test would help me to pass the MCPA Level 1 exam at the first attempt?

Passing the MCPA Level 1 exam makes you MuleSoft Certified Platform Architect - Level 1 (MCPA). Having the Platform Architect Level 1 certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

MuleSoft MCPA Level 1 Platform Architect Level 1 Certification Details:

Exam Name	MuleSoft Certified Platform Architect - Level 1 (MCPA)
Exam Code	MCPA Level 1
Exam Price	\$400 (USD)
Duration	120 mins
Number of Questions	60
Passing Score	70%
Schedule Exam	Purchase Exam
Books / Training	Anypoint Platform Architecture: Application Networks
Sample Questions	MuleSoft MCPA Level 1 Sample Questions
Practice Exam	MuleSoft MCPA Level 1 Certification Practice Exam

MCPA Level 1 Syllabus:

Topic	Details
Explaining application network basics	<ul style="list-style-type: none"> - Identify and differentiate between technologies typically used to implement API-led connectivity. - Describe the role and characteristics of web APIs. - Correctly assign APIs to tiers according to ownership, functional focus, and rate of change. - Describe the capabilities and high-level components of Anypoint Platform.
Establishing organizational and platform foundations	<ul style="list-style-type: none"> - Advise on establishing a Center for Enablement (C4E) and identify KPIs to measure its success. - Describe the high-level structure and benefits of MuleSoft Catalyst. - Compare and contrast Identity Management and Client Management options on Anypoint Platform. - Identify data residency of different kinds of data (payload, metrics, and others).
Designing and sharing APIs	<ul style="list-style-type: none"> - Identify dependencies between an API, its API specification, its implementation, and its clients. - Describe the creation and publication of reusable API-related assets using API specifications and Anypoint Platform components. - Identify changes to an API that would require or not require changing the major/minor/patch component of its semantic version. - Given a specific power relationship between two Bounded Contexts, choose the most appropriate strategy for mapping between the API data models of these Bounded Contexts. - Identify idempotent HTTP methods and HTTP-native support for optimistic concurrency. - Recognize the important features and functionality of API Designer to design API specifications.
Designing APIs using System, Process, and Experience Layers	<ul style="list-style-type: none"> - Identify appropriate APIs to implement a business process and assign them to layers of API-led connectivity. - Assign APIs to layers according to ownership, functional focus, and rate of change. - Recommend the most appropriate approach relating the API data model of System APIs to that of their backend system based on specific requirements and organizational characteristics.
Governing APIs on Anypoint Platform	<ul style="list-style-type: none"> - Make appropriate use of API instances and environments in API manager taking into account the nature of the API and the underlying data and

Topic	Details
	<p>system.</p> <ul style="list-style-type: none"> - Select appropriate API policies and other components of the AP to support specific non-functional requirements (NFRs). - Identify any change(s) required to an API specification to reflect the application of an API policy with specific characteristics. - Select an approach to API policy enforcement based on specific preferences and constraints, including, but not limited to, API proxies/API gateways and Anypoint Service Mesh. - Secure web APIs using API policies appropriately chosen for the API's tier (System, Process, Experience). - Describe in what circumstances and how to pass client ID and secret to a web API. - Explain how to request access to an API version for an API client, and how that access is approved and revoked. - Select appropriate API policies to enforce non-functional security constraints on web API invocations. - Explain the relationships of Anypoint Platform, external Identity Providers, AP Business Groups, and API clients in the context of OAuth 2.0. - Identify scenarios needing custom API policies.
Architecting and deploying API implementations	<ul style="list-style-type: none"> - Explain how to use auto-discovery to link a web API implementation to an API instance managed with API Manager. - Identify requirements that call for the use of an Anypoint VPC. - Compare and contrast options for hosting and managing Anypoint Platform runtime planes. - Compare unit and integration tests and specify where MUnit is best employed. - Explain options for automated build, test, and deploy of API implementations and related artifacts in a DevOps setting.
Deploying API implementations to CloudHub	<ul style="list-style-type: none"> - Describe the scenarios for which Object Store should be used with CloudHub. - Select CloudHub worker sizes and configuration as appropriate. - Given an app deployed to the CloudHub shared worker cloud in one or more regions, describe and predict its reliability and performance characteristics.

Topic	Details
	<ul style="list-style-type: none"> - Identify the defining differences between the CloudHub Shared and Dedicated Load Balancers. - Compare and contrast the options for CloudHub networking in the presence of customer-owned Amazon VPCs and on-premises data centers. - Identify and avoid single points of failure in CloudHub deployments of API implementations.
Meeting API quality goals	<ul style="list-style-type: none"> - Design, describe, and differentiate between scenarios that use the an object store or caching. - Select resilience strategies that help web API clients guard against failures when invoking APIs. - Describe when horizontal scaling of an API implementation is or is not likely to benefit response time and throughput as seen by API clients.
Monitoring and analyzing application networks	<ul style="list-style-type: none"> - Identify the components of Anypoint Platform that generate data for monitoring and alerting. - Describe the metrics collected by Anypoint Platform for API invocations. - Specify alerts to define for key metrics of API invocations for all layers of API-led connectivity. - Specify alerts to define for Mule applications.

MuleSoft MCPA Level 1 Sample Questions:

Question: 1

A retail company is using an Order API to accept new orders. The Order API uses a JMS queue to submit orders to a backend order management service. The normal load for orders is being handled using two (2) CloudHub workers, each configured with 0.2 vCore. The CPU load of each CloudHub worker normally runs well below 70%. However, several times during the year the Order API gets four times (4x) the average number of orders.

This causes the CloudHub worker CPU load to exceed 90% and the order submission time to exceed 30 seconds. The cause, however, is NOT the backend order management service, which still responds fast enough to meet the response SLA for the Order API. What is the MOST resource-efficient way to configure the Mule application's CloudHub deployment to help the company cope with this performance challenge?

- Permanently increase the size of each of the two (2) CloudHub workers by at least four times (4x) to one (1) vCore
- Use a horizontal CloudHub autoscaling policy that triggers on CPU utilization greater than 70%
- Use a vertical CloudHub autoscaling policy that triggers on CPU utilization greater than 70%
- Permanently increase the number of CloudHub workers by four times (4x) to eight (8) CloudHub workers

Answer: b

Question: 2

A company wants to move its Mule API implementations into production as quickly as possible. The company's InfoSec group requires that all APIs be accessed using HTTPS and two-way authentication.

What combination of runtime plane and control plane options meets these project lifecycle goals?

- a) Manually provisioned customer-hosted runtime plane and MuleSoft-hosted control plane
- b) iPaaS provisioned customer-hosted runtime plane and MuleSoft-hosted control plane
- c) Manually provisioned customer-hosted runtime plane and customer-hosted control plane
- d) MuleSoft-hosted runtime plane and MuleSoft-hosted control plane

Answer: d

Question: 3

What is the main change to the IT operating model that MuleSoft recommends to organizations to improve innovation and clock speed?

- a) Drive consumption as much as production of assets; this enables developers to discover and reuse assets from other projects and encourages standardization
- b) Expose assets using a Master Data Management (MDM) system; this standardizes projects and enables developers to quickly discover and reuse assets from other projects
- c) Implement SOA for reusable APIs to focus on production over consumption; this standardizes on XML and WSDL formats to speed up decision making
- d) Create a lean and agile organization that makes many small decisions everyday; this speeds up decision making and enables each line of business to take ownership of its projects

Answer: a

Question: 4

A company has started to create an application network and is now planning to implement a Center for Enablement (C4E) organizational model. What key factor would lead the company to decide upon a federated rather than a centralized C4E?

- a) When various teams responsible for creating APIs are new to integration and hence need extensive training
- b) When the majority of the applications in the application network are cloud based
- c) When development is already organized into several independent initiatives or groups
- d) When there are a large number of existing common assets shared by development teams

Answer: c

Question: 5

A set of tests must be performed prior to deploying API implementations to a staging environment. Due to data security and access restrictions, untested APIs cannot be granted access to the backend systems, so instead mocked data must be used for these tests.

The amount of available mocked data and its contents is sufficient to entirely test the API implementations with no active connections to the backend systems.

What type of tests should be used to incorporate this mocked data?

- a) Integration tests
- b) Performance tests
- c) Functional tests (Blackbox)
- d) Unit tests (Whitebox)

Answer: d

Question: 6

An organization is starting an API-led connectivity journey and is deciding between deployment options. Currently, all their systems are customer-hosted on their premises.

A CloudHub region is available close to their data center and their InfoSec team has approved the usage of an IPsec tunnel. The organization has agreed to use a cloud-first approach and in some Mule application deployments the Mule runtime must be fine-tuned for lower latency.

What is the best runtime plane option for the organization, given their current systems and requirements?

- a) Anypoint Runtime Fabric
- b) CloudHub workers in a CloudHub region nearest to their data center
- c) Customer-hosted Mule runtimes
- d) A hybrid combination of customer-hosted and MuleSoft-hosted Mule runtimes

Answer: d

Question: 7

A Mule application exposes an HTTPS endpoint and is deployed to three CloudHub workers that do not use static IP addresses. The Mule application expects a high volume of client requests in short time periods. What is the most cost-effective infrastructure component that should be used to serve the high volume of client requests?

- a) The CloudHub shared load balancer
- b) Runtime Manager autoscaling
- c) An API proxy
- d) A customer-hosted load balancer

Answer: a

Question: 8

An API client calls one method from an existing API implementation. The API implementation is later updated.

What change to the API implementation would require the API client's invocation logic to also be updated?

- a) When the data type of the response is changed for the method called by the API client
- b) When a new method is added to the resource used by the API client
- c) When a new required field is added to the method called by the API client
- d) When a child method is added to the method called by the API client

Answer: c

Question: 9

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?

- a) At both the API proxy and the API implementation
- b) At the API proxy
- c) At a MuleSoft-hosted load balancer
- d) At the API implementation

Answer: b

Question: 10

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?

- a) Rate limiting
- b) HTTP caching
- c) Rate limiting - SLA based
- d) Spike control

Answer: d

Study Guide to Crack MuleSoft Platform Architect Level 1 MCPA Level 1 Exam:

- Getting details of the MCPA Level 1 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the MCPA Level 1 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the MuleSoft provided training for MCPA Level 1 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the MCPA Level 1 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on MCPA Level 1 practice tests is must. Continuous practice will make you an expert in all syllabus areas.

Reliable Online Practice Test for MCPA Level 1 Certification

Make CertFun.com your best friend during your MuleSoft Platform Architect Level 1 exam preparation. We provide authentic practice tests for the MCPA Level 1 exam. Experts design these online practice tests, so we can offer you an exclusive experience of taking the actual MCPA Level 1 exam. We guarantee you 100% success in your first exam attempt if you continue practicing regularly. Don't bother if you don't get 100% marks in initial practice exam attempts. Just utilize the result section to know your strengths and weaknesses and prepare according to that until you get 100% with our practice tests. Our evaluation makes you confident, and you can score high in the MCPA Level 1 exam.

Start Online Practice of MCPA Level 1 Exam by Visiting URL

<https://www.certfun.com/mulesoft/mcpa-mulesoft-platform-architect-level-1>