Free Questions for MCPA-Level-1 by actualtestdumps

Shared by Vaughan on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

A company wants to move its Mule API implementations into production as quickly as possible. To protect access to all Mule application data and metadata, the company requires that all Mule applications be deployed to the company's customer-hosted infrastructure within the corporate firewall. What combination of runtime plane and control plane options meets these project lifecycle goals?

Options:

A- Manually provisioned customer-hosted runtime plane and customer-hosted control plane

B- MuleSoft-hosted runtime plane and customer-hosted control plane

C- Manually provisioned customer-hosted runtime plane and MuleSoft-hosted control plane

D- iPaaS provisioned customer-hosted runtime plane and MuleSoft-hosted control plane

Answer:

A

Explanation:

Correct Answer: Manually provisioned customer-hosted runtime plane and customer-hosted control plane

****

There are two key factors that are to be taken into consideration from the scenario given in the question.

>> Company requires both data and metadata to be resided within the corporate firewall

>> Company would like to go with customer-hosted infrastructure.

Any deployment model that is to deal with the cloud directly or indirectly (Mulesoft-hosted or Customer's own cloud like Azure, AWS) will have to share atleast the metadata.

Application data can be controlled inside firewall by having Mule Runtimes on customer hosted runtime plane. But if we go with Mulsoft-hosted/ Cloud-based control plane, the control plane required atleast some minimum level of metadata to be sent outside the corporate firewall.

As the customer requirement is pretty clear about the data and metadata both to be within the corporate firewall, even though customer wants to move to production as quickly as possible, unfortunately due to the nature of their security requirements, they have no other option but to go with manually provisioned customer-hosted runtime plane and customer-hosted control plane.

Question 2

Question Type: **MultipleChoice**

Version 3.0.1 of a REST API implementation represents time values in PST time using ISO 8601 hh:mm:ss format. The API implementation needs to be changed to instead represent time values in CEST time using ISO 8601 hh:mm:ss format. When following the semver.org semantic versioning specification, what version should be assigned to the updated API implementation?

Options:

A- 3.0.2

B- 4.0.0

C- 3.1.0

D- 3.0.1

Answer:

B

Explanation:

Correct Answer: 4.0.0

****

As per semver.org semantic versioning specification:

Given a version number MAJOR.MINOR.PATCH, increment the:

- MAJOR version when you make incompatible API changes.

- MINOR version when you add functionality in a backwards compatible manner.

- PATCH version when you make backwards compatible bug fixes.

As per the scenario given in the question, the API implementation is completely changing its behavior. Although the format of the time is still being maintained as hh:mm:ss and there is no change in schema w.r.t format, the API will start functioning different after this change as the times are going to come completely different.

Example: Before the change, say, time is going as 09:00:00 representing the PST. Now on, after the change, the same time will go as 18:00:00 as Central European Summer Time is 9 hours ahead of Pacific Time.

>> This may lead to some uncertain behavior on API clients depending on how they are handling the times in the API response. All the API clients need to be informed that the API functionality is going to change and will return in CEST format. So, this considered as a MAJOR change and the version of API for this new change would be 4.0.0

Question 3

Question Type: MultipleChoice

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

Options:

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

Explanation:

Correct Answer: Drive consumption as much as production of assets; this enables developers to discover and reuse assets from other projects and encourages standardization

****

>> The main motto of the new IT Operating Model that MuleSoft recommends and made popular is to change the way that they are delivered from a production model to a production + consumption model, which is done through an API strategy called API-led connectivity.

>> The assets built should also be discoverable and self-serveable for reusablity across LOBs and organization.

>> MuleSoft's IT operating model does not talk about SDLC model (Agile/ Lean etc) or MDM at all. So, options suggesting these are not valid.

References:

https://blogs.mulesoft.com/biz/connectivity/what-is-a-center-for-enablement-c4e/

https://www.mulesoft.com/resources/api/secret-to-managing-it-projects

Question 4

Question Type: **MultipleChoice**

An Anypoint Platform organization has been configured with an external identity provider (IdP) for identity management and client management. What credentials or token must be provided to Anypoint CLI to execute commands against the Anypoint Platform APIs?

Options:

A- The credentials provided by the IdP for identity management

B- The credentials provided by the IdP for client management

C- An OAuth 2.0 token generated using the credentials provided by the IdP for client management

D- An OAuth 2.0 token generated using the credentials provided by the IdP for identity management

Answer:

A

Explanation:

Correct Answer: The credentials provided by the IdP for identity management

****

>> There is no support for OAuth 2.0 tokens from client/identity providers to authenticate via Anypoint CLI. Only possible tokens are 'bearer tokens' that too only generated using Anypoint Organization/Environment Client Id and Secret from https://anypoint.mulesoft.com/accounts/login. Not the client credentials of client provider. So, OAuth 2.0 is not possible. More over, the token is mainly for API Manager purposes and not associated with a user. You can NOT use it to call most APIs (for example Cloudhub and etc) as per this Mulesoft Knowledge article.

>> The other option allowed by Anypoint CLI is to use client credentials. It is possible to use client credentials of a client provider but requires setting up Connected Apps in client management but such details are not given in the scenario explained in the question.

>> So only option left is to use user credentials from identify provider

Question 5

Say, there is a legacy CRM system called CRM-Z which is offering below functions:

1. Customer creation

2. Amend details of an existing customer

3. Retrieve details of a customer

4. Suspend a customer

Options:

A- Implement a system API named customerManagement which has all the functionalities wrapped in it as various operations/resources

B- Implement different system APIs named createCustomer, amendCustomer, retrieveCustomer and suspendCustomer as they are modular and has seperation of concerns

C- Implement different system APIs named createCustomerInCRMZ, amendCustomerInCRMZ, retrieveCustomerFromCRMZ and suspendCustomerInCRMZ as they are modular and has seperation of concerns

•Explanation:

Correct Answer: Implement different system APIs named createCustomer, amendCustomer, retrieveCustomer and suspendCustomer as they are modular and has seperation of concerns

****

>> It is quite normal to have a single API and different Verb + Resource combinations. However, this fits well for an Experience API or a Process API but not a best architecture style for System APIs. So, option with just one customerManagement API is not the best choice here.

>> The option with APIs in createCustomerInCRMZ format is next close choice w.r.t modularization and less maintenance but the naming of APIs is directly coupled with the legacy system. A better foreseen approach would be to name your APIs by abstracting the backend system names as it allows seamless replacement/migration of any backend system anytime. So, this is not the correct choice too.

>> createCustomer, amendCustomer, retrieveCustomer and suspendCustomer is the right approach and is the best fit compared to other options as they are both modular and same time got the names decoupled from backend system and it has covered all requirements a System API needs.

Answer:

B

Question 6

Question Type: **MultipleChoice**

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?

Options:

A- Anypoint Runtime Fabric

B- Anypoint Platform for Pivotal Cloud Foundry

C- CloudHub

D- A hybrid combination of customer-hosted and MuleSoft-hosted Mule runtimes

Answer:

A

Explanation:

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.

Question 7

Question Type: **MultipleChoice**

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?

Options:

A- When there are a large number of existing common assets shared by development teams

B- When various teams responsible for creating APIs are new to integration and hence need extensive training

C- When development is already organized into several independent initiatives or groups

D- When the majority of the applications in the application network are cloud based

Answer:

C

Explanation:

Correct Answer: When development is already organized into several independent initiatives or groups

****

>> It would require lot of process effort in an organization to have a single C4E team coordinating with multiple already organized development teams which are into several independent initiatives. A single C4E works well with different teams having at least a common initiative. So, in this scenario, federated C4E works well instead of centralized C4E.

Question 8

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?

Options:

A- Permanently increase the size of each of the two (2) CloudHub workers by at least four times (4x) to one (1) vCore

B- Use a vertical CloudHub autoscaling policy that triggers on CPU utilization greater than 70%

C- Permanently increase the number of CloudHub workers by four times (4x) to eight (8) CloudHub workers

D- Use a horizontal CloudHub autoscaling policy that triggers on CPU utilization greater than 70%

Answer:

D

Explanation:

Correct Answer: Use a horizontal CloudHub autoscaling policy that triggers on CPU utilization greater than 70%

****

The scenario in the question is very clearly stating that the usual traffic in the year is pretty well handled by the existing worker configuration with CPU running well below 70%. The problem occurs only 'sometimes' occasionally when there is spike in the number of orders coming in.

So, based on above, We neither need to permanently increase the size of each worker nor need to permanently increase the number of workers. This is unnecessary as other than those 'occasional' times the resources are idle and wasted.

We have two options left now. Either to use horizontal Cloudhub autoscaling policy to automatically increase the number of workers or to use vertical Cloudhub autoscaling policy to automatically increase the vCore size of each worker.

Here, we need to take two things into consideration:

1. CPU

2. Order Submission Rate to JMS Queue

>> From CPU perspective, both the options (horizontal and vertical scaling) solves the issue. Both helps to bring down the usage below 90%.

>> However, If we go with Vertical Scaling, then from Order Submission Rate perspective, as the application is still being load balanced with two workers only, there may not be much improvement in the incoming request processing rate and order submission rate to JMS queue. The throughput would be same as before. Only CPU utilization comes down.

>> But, if we go with Horizontal Scaling, it will spawn new workers and adds extra hand to increase the throughput as more workers are being load balanced now. This way we can address both CPU and Order Submission rate.

Hence, Horizontal CloudHub Autoscaling policy is the right and best answer.

Question 9

Question Type: MultipleChoice

A REST API is being designed to implement a Mule application.

What standard interface definition language can be used to define REST APIs?

Options:

A- Web Service Definition Language(WSDL)

B- OpenAPI Specification (OAS)

C- YAML

D- AsyncAPI Specification

Answer:

B

Question 10

Question Type: MultipleChoice

What Anypoint Connectors support transactions?

Options:

A- Database, JMS, VM

B- Database, 3MS, HTTP

C- Database, JMS, VM, SFTP

D- Database, VM, File

Answer:

A

To Get Premium Files for MCPA-Level-1 Visit

For More Free Questions Visit

20% DISCOUNT