MCPA-LEVEL1$^{\text{Q\&As}}$

MuleSoft Certified Platform Architect - Level 1

Pass Mulesoft MCPA-LEVEL1 Exam with 100% Guarantee

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

https://www.pass4itsure.com/mulesoft-certified-platform-architect-level-1.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

Which of the below, when used together, makes the IT Operational Model effective?

A. Create reusable assets, Do marketing on the created assets across organization, Arrange time to time LOB reviews to ensure assets are being consumed or not

B. Create reusable assets, Make them discoverable so that LOB teams can self-serve and browse the APIs, Get active feedback and usage metrics

C. Create resuable assets, make them discoverable so that LOB teams can self-serve and browse the APIs

Correct Answer: C

Create reusable assets, Make them discoverable so that LOB teams can self-serve and browse the APIs, Get active feedback and usage metrics. **

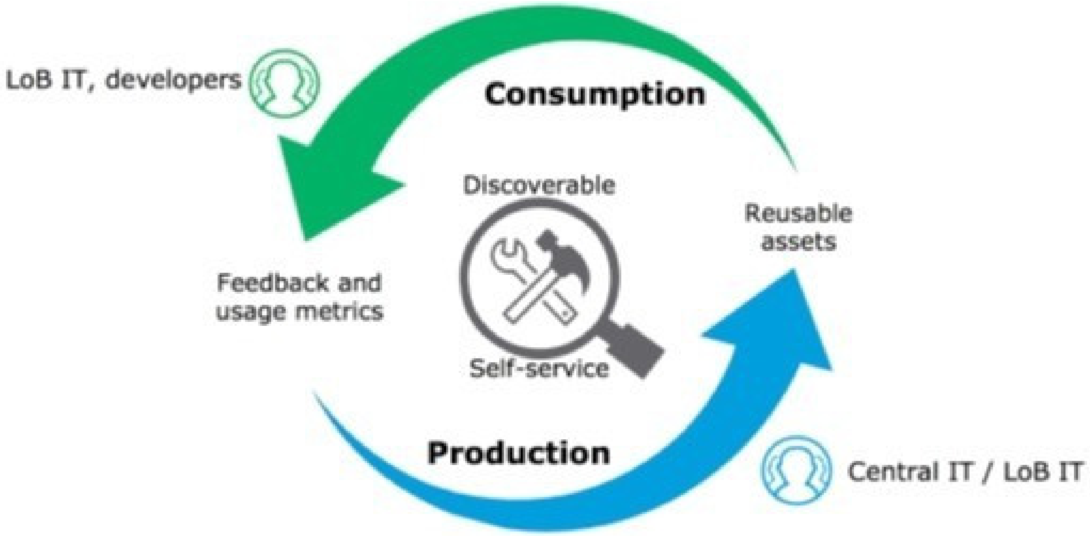

QUESTION 2

A new upstream API Is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity.

The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms.

If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API\\'s desired SLA?

A. Set a timeout of 50 ms; this times out more invocations of that API but gives additional room for retries

B. Set a timeout of 100 ms; that leaves 400 ms for the other two downstream APIs to complete

C. No timeout is possible to meet the upstream API\\'s desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

D. Do not set a timeout; the Invocation of this API Is mandatory and so we must wait until it responds

Correct Answer: B

Set a timeout of 100ms; that leaves 400ms for other two downstream APIs to complete

****

Key details to take from the given scenario:

>> Upstream API\\'s designed SLA is 500ms (median). Lets ignore maximum SLA response times.

>> This API calls 3 downstream APIs sequentially and all these are of similar complexity. >> The first downstream API is offering median SLA of 100ms, 80th percentile: 500ms; 95th percentile: 1000ms.

Based on the above details:

>> We can rule out the option which is suggesting to set 50ms timeout. Because, if the median SLA itself being offered is 100ms then most of the calls are going to timeout and time gets wasted in retried them and eventually gets exhausted

with all retries. Even if some retries gets successful, the remaining time wont leave enough room for 2nd and 3rd downstream APIs to respond within time. >> The option suggesting to NOT set a timeout as the invocation of this API is

mandatory and so we must wait until it responds is silly. As not setting time out would go against the good implementation pattern and moreover if the first API is not responding within its offered median SLA 100ms then most probably it would

either respond in 500ms (80th percentile) or 1000ms (95th percentile). In BOTH cases, getting a successful response from 1st downstream API does NO GOOD because already by this time the Upstream API SLA of 500 ms is breached.

There is no time left to call 2nd and 3rd downstream APIs. >> It is NOT true that no timeout is possible to meet the upstream APIs desired SLA. As 1st downstream API is offering its median SLA of 100ms, it means MOST of the time we

would get the responses within that time. So, setting a timeout of 100ms would be ideal for MOST calls as it leaves enough room of 400ms for remaining 2 downstream API calls.

QUESTION 3

An organization has created an API-led architecture that uses various API layers to integrate mobile clients with a backend system. The backend system consists of a number of specialized components and can be accessed via a REST API. The process and experience APIs share the same bounded-context model that is different from the backend data model. What additional canonical models, bounded-context models, or anti-corruption layers are best added to this architecture to help process data consumed from the backend system?

A. Create a bounded-context model for every layer and overlap them when the boundary contexts overlap, letting API developers know about the differences between upstream and downstream data models

B. Create a canonical model that combines the backend and API-led models to simplify and unify data models, and minimize data transformations.

C. Create a bounded-context model for the system layer to closely match the backend data model, and add an anti-corruption layer to let the different bounded contexts cooperate across the system and process layers

D. Create an anti-corruption layer for every API to perform transformation for every data model to match each other, and let data simply travel between APIs to avoid the complexity and overhead of building canonical models

Correct Answer: C

Create a bounded-context model for the system layer to closely match the backend data model, and add an anti-corruption layer to let the different bounded contexts cooperate across the system and process layers ** >> Canonical models are not an option here as the organization has already put in efforts and created bounded-context models for Experience and Process APIs. >> Anti-corruption layers for ALL APIs is unnecessary and invalid because it is mentioned that experience and process APIs share same bounded-context model. It is just the System layer APIs that need to choose their approach now. >> So, having an anti-corruption layer just between the process and system layers will work well. Also to speed up the approach, system APIs can mimic the backend system data model.

QUESTION 4

Refer to the exhibit.

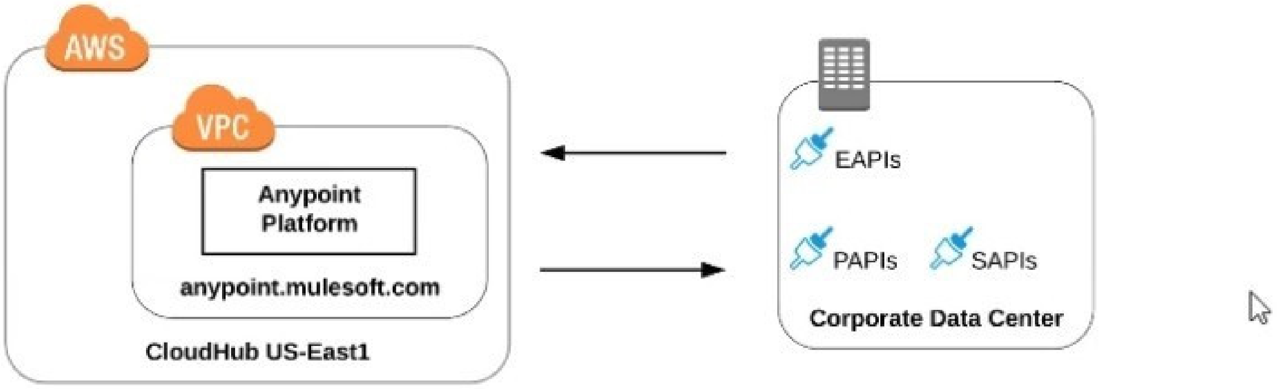

What is true when using customer-hosted Mule runtimes with the MuleSoft-hosted Anypoint Platform control plane (hybrid deployment)?

A. Anypoint Runtime Manager initiates a network connection to a Mule runtime in order to deploy Mule applications

B. The MuleSoft-hosted Shared Load Balancer can be used to load balance API invocations to the Mule runtimes

C. API implementations can run successfully in customer-hosted Mule runtimes, even when they are unable to communicate with the control plane

D. Anypoint Runtime Manager automatically ensures HA in the control plane by creating a new Mule runtime instance in case of a node failure

Correct Answer: C

API implementations can run successfully in customer-hosted Mule runtimes, even when they are unable to communicate with the control plane.

****

>> We CANNOT use Shared Load balancer to load balance APIs on customer hosted runtimes

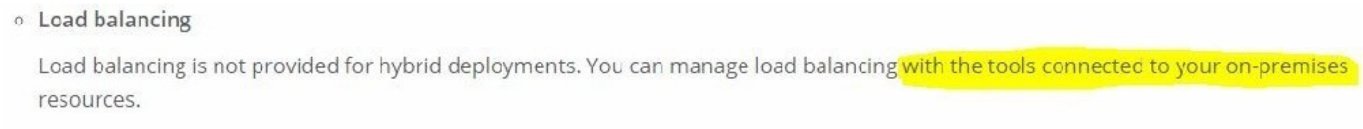

- Load balancing

  Load balancing is not provided for hybrid deployments. You can manage load balancing with the tools connected to your on-premises resources.

>> For Hybrid deployment models, the on-premises are first connected to Runtime Manager using Runtime Manager agent. So, the connection is initiated first from On- premises to Runtime Manager. Then all control can be done from Runtime Manager. >> Anypoint Runtime Manager CANNOT ensure automatic HA. Clusters/Server Groups etc should be configured before hand.

Only TRUE statement in the given choices is, API implementations can run successfully in customer-hosted Mule runtimes, even when they are unable to communicate with the control plane. There are several references below to justify this statement.

References: https://docs.mulesoft.com/runtime-manager/deployment-strategies#hybrid-deployments https://help.mulesoft.com/s/article/On-Premise-Runtimes-Disconnected-From-US-Control- Plane-June-18th-2018 https://help.mulesoft.com/s/article/Runtime-Manager-cannot-manage-On-Prem-Applications-and-Servers-from-US-Control-Plane-June-25th-2019 https://help.mulesoft.com/s/article/On-premise-Runtimes-Appear-Disconnected-in-RuntimeManager-May-29th-2018

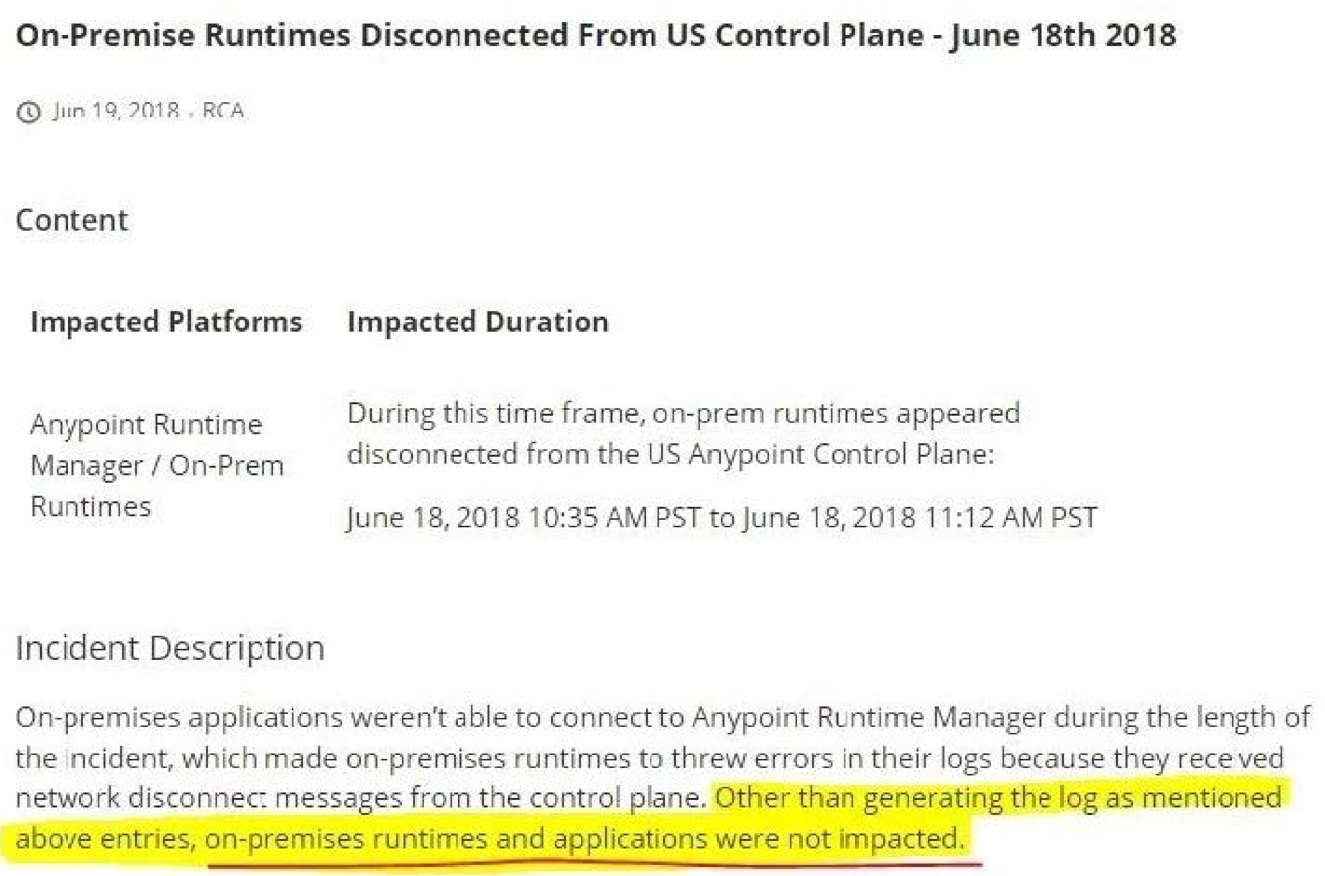

On-Premise Runtimes Disconnected From US Control Plane - June 18th 2018

Jun 19, 2018 . RCA

Content

Impacted Platforms	Impacted Duration
Anypoint Runtime Manager / On-Prem Runtimes	During this time frame, on-prem runtimes appeared disconnected from the US Anypoint Control Plane: June 18, 2018 10:35 AM PST to June 18, 2018 11:12 AM PST

Incident Description

On-premises applications weren't able to connect to Anypoint Runtime Manager during the length of the incident, which made on-premises runtimes to threw errors in their logs because they received network disconnect messages from the control plane. Other than generating the log as mentioned above entries, on-premises runtimes and applications were not impacted.

============================

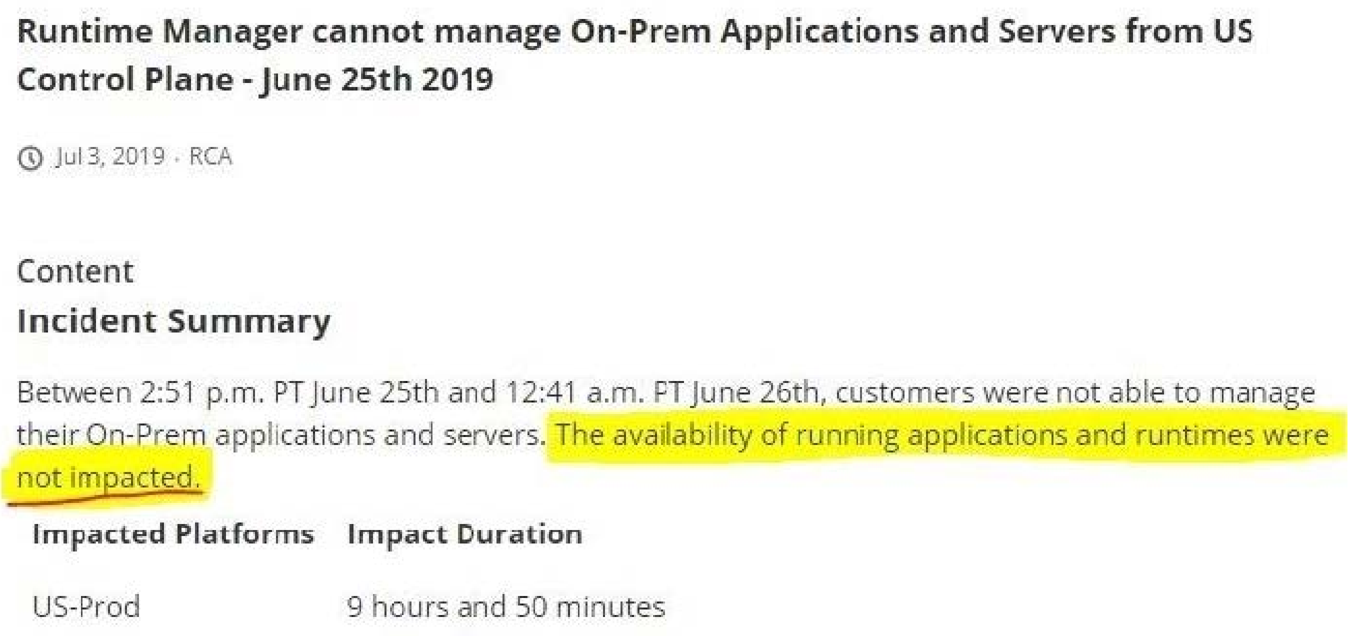

Runtime Manager cannot manage On-Prem Applications and Servers from US Control Plane - June 25th 2019

Jul 3, 2019 · RCA

Content

Incident Summary

Between 2:51 p.m. PT June 25th and 12:41 a.m. PT June 26th, customers were not able to manage their On-Prem applications and servers. The availability of running applications and runtimes were not impacted.

Impacted Platforms	Impact Duration
US-Prod	9 hours and 50 minutes

============================

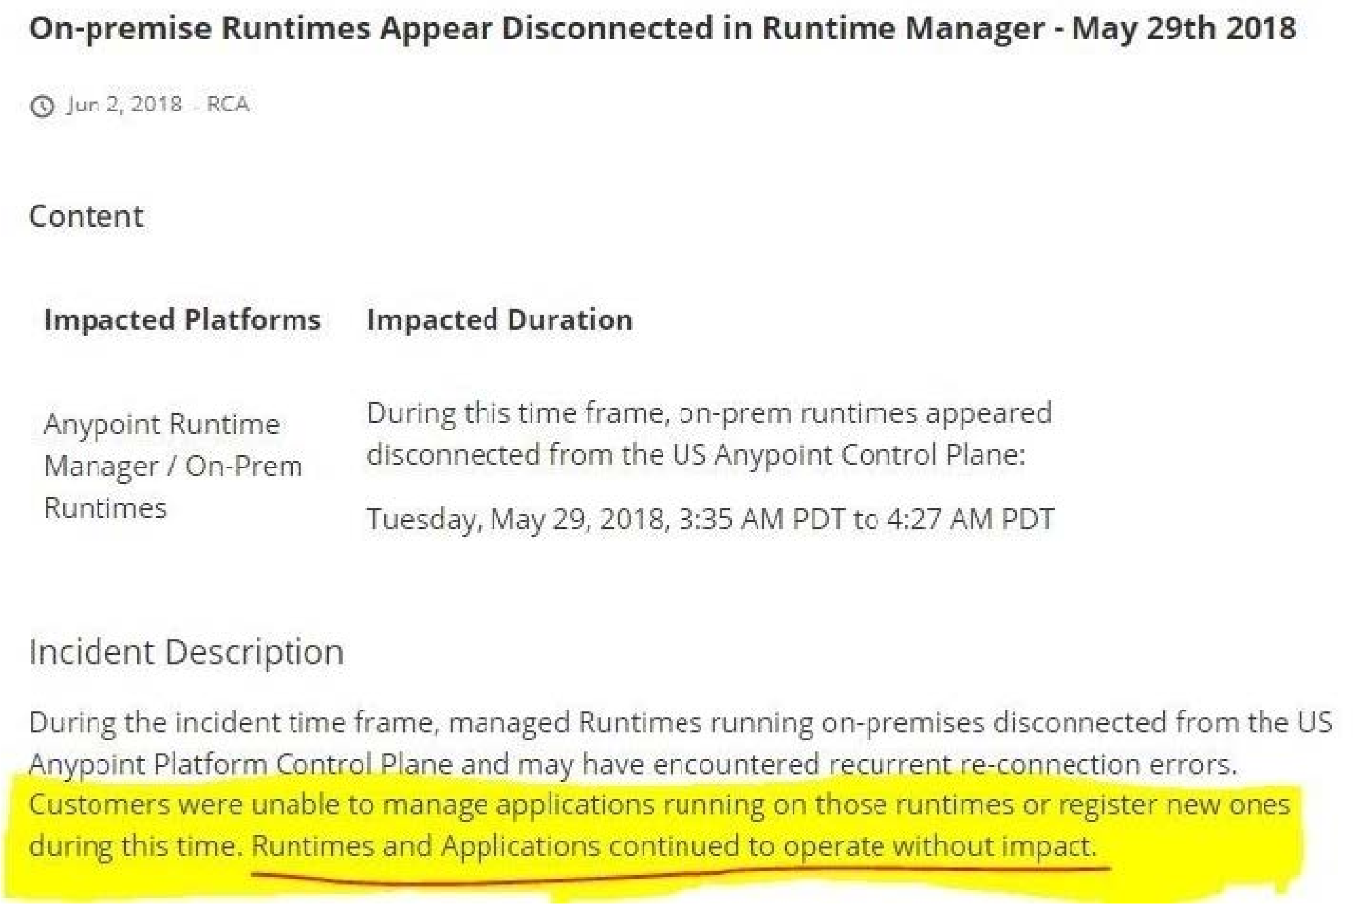

On-premise Runtimes Appear Disconnected in Runtime Manager - May 29th 2018

Jun 2, 2018 · RCA

Content

Impacted Platforms	Impacted Duration
Anypoint Runtime Manager / On-Prem Runtimes	During this time frame, on-prem runtimes appeared disconnected from the US Anypoint Control Plane: Tuesday, May 29, 2018, 3:35 AM PDT to 4:27 AM PDT

Incident Description

During the incident time frame, managed Runtimes running on-premises disconnected from the US Anypoint Platform Control Plane and may have encountered recurrent re-connection errors. Customers were unable to manage applications running on those runtimes or register new ones during this time. Runtimes and Applications continued to operate without impact.

QUESTION 5

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?

A. When there are three CloudHub deployments of the API implementation to three separate CloudHub regions that must share the cache state

B. 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

C. 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

D. When there is one CloudHub deployment of the API implementation to three CloudHub workers that must share the cache state

Correct Answer: D

When there is one CloudHub deployment of the API implementation to three CloudHub workers that must share the cache state.

****

Key details in the scenario:

>> Use the CloudHub Object Store via the Object Store connector Considering above details:

>> CloudHub Object Stores have one-to-one relationship with CloudHub Mule Applications. >> We CANNOT use an application\\'s CloudHub Object Store to be shared among multiple Mule applications running in different Regions or Business

Groups or Customer-hosted Mule Runtimes by using Object Store connector.

>> If it is really necessary and very badly needed, then Anypoint Platform supports a way by allowing access to CloudHub Object Store of another application using Object Store REST API. But NOT using Object Store connector. So, the only

scenario where we can use the CloudHub Object Store via the Object Store connector to persist the cache\\'s state is when there is one CloudHub deployment of the API implementation to multiple CloudHub workers that must share the cache

state.

QUESTION 6

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?

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

Correct Answer: A

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.

Reference: https://docs.mulesoft.com/runtime-fabric/1.7/

QUESTION 7

What is the most performant out-of-the-box solution in Anypoint Platform to track transaction state in an asynchronously executing long-running process implemented as a Mule application deployed to multiple CloudHub workers?

A. Redis distributed cache

B. java.util.WeakHashMap

C. Persistent Object Store

D. File-based storage

Correct Answer: C

Persistent Object Store ***

>> Redis distributed cache is performant but NOT out-of-the-box solution in Anypoint Platform

>> File-storage is neither performant nor out-of-the-box solution in Anypoint Platform >> java.util.WeakHashMap needs

a completely custom implementation of cache from scratch using Java code and is limited to the JVM where it is running.

Which means the state in the cache is not worker aware when running on multiple workers. This type of cache is local to the worker. So, this is neither out-of-the-box nor worker-aware among multiple workers on cloudhub. https://

www.baeldung.com/java-weakhashmap >> Persistent Object Store is an out-of-the-box solution provided by Anypoint Platform which is performant as well as worker aware among multiple workers running on CloudHub. https://

docs.mulesoft.com/object-store/

So, Persistent Object Store is the right answer.

QUESTION 8

What is a typical result of using a fine-grained rather than a coarse-grained API deployment model to implement a given business process?

A. A decrease in the number of connections within the application network supporting the business process

B. A higher number of discoverable API-related assets in the application network

C. A better response time for the end user as a result of the APIs being smaller in scope and complexity

D. An overall tower usage of resources because each fine-grained API consumes less resources

Correct Answer: B

A higher number of discoverable API-related assets in the application network.

****

>> We do NOT get faster response times in fine-grained approach when compared to coarse-grained approach.

>> In fact, we get faster response times from a network having coarse-grained APIs compared to a network having fine-grained APIs model. The reasons are below.

Fine-grained approach:

1.

will have more APIs compared to coarse-grained

2.

So, more orchestration needs to be done to achieve a functionality in business process.

3.

Which means, lots of API calls to be made. So, more connections will needs to be established. So, obviously more hops, more network i/o, more number of integration points compared to coarse-grained approach where fewer APIs with bulk functionality embedded in them.

4.

That is why, because of all these extra hops and added latencies, fine-grained approach will have bit more response times compared to coarse-grained.

5.

Not only added latencies and connections, there will be more resources used up in fine- grained approach due to more number of APIs. That\\'s why, fine-grained APIs are good in a way to expose more number of resuable assets in your network and make them discoverable. However, needs more maintenance, taking care of integration points, connections, resources with a little compromise w.r.t network hops and response times.

QUESTION 9

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?

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

Correct Answer: A

The credentials provided by the IdP for identity management *** Reference: https://docs.mulesoft.com/runtime-manager/anypoint-platform- cli#authentication

>> 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 10

An API implementation is updated. When must the RAML definition of the API also be updated?

A. When the API implementation changes the structure of the request or response messages

B. When the API implementation changes from interacting with a legacy backend system deployed on-premises to a modern, cloud-based (SaaS) system

C. When the API implementation is migrated from an older to a newer version of the Mule runtime

D. When the API implementation is optimized to improve its average response time

Correct Answer: A

When the API implementation changes the structure of the request or response messages

****

>> RAML definition usually needs to be touched only when there are changes in the request/response schemas or in any traits on API. >> It need not be modified for any internal changes in API implementation like performance tuning,

backend system migrations etc..

QUESTION 11

Which of the following best fits the definition of API-led connectivity?

A. API-led connectivity is not just an architecture or technology but also a way to organize people and processes for efficient IT delivery in the organization

B. API-led connectivity is a 3-layered architecture covering Experience, Process and System layers

C. API-led connectivity is a technology which enabled us to implement Experience, Process and System layer based APIs

Correct Answer: A

API-led connectivity is not just an architecture or technology but also a way to organize people and processes for efficient IT delivery in the organization.

****

Reference: https://blogs.mulesoft.com/dev/api-dev/what-is-api-led-connectivity/

QUESTION 12

Refer to the exhibit.

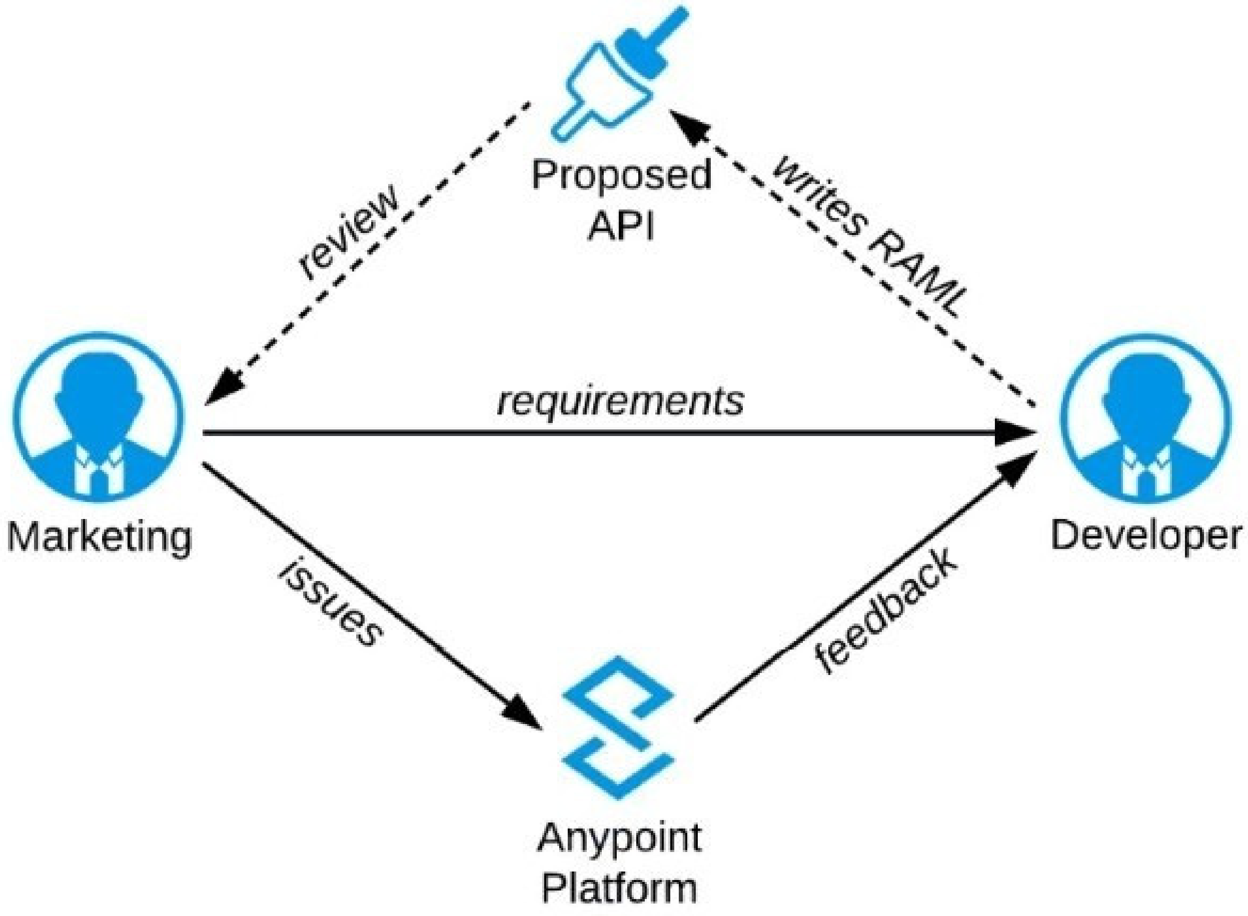

A RAML definition has been proposed for a new Promotions Process API, and has been published to Anypoint Exchange.

The Marketing Department, who will be an important consumer of the Promotions API, has important requirements and expectations that must be met.

What is the most effective way to use Anypoint Platform features to involve the Marketing Department in this early API design phase?

C. Use an on-premises installation of Mule runtimes that are completely isolated with NO external network access, managed by the Anypoint Platform Private Cloud Edition control plane

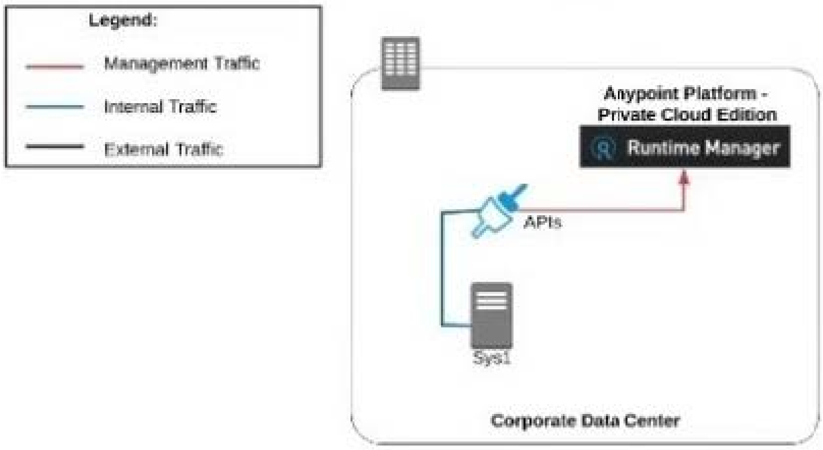

D. Use a combination of Cloud Hub-deployed and manually provisioned on-premises Mule runtimes managed by the MuleSoft-hosted Anypoint Platform control plane

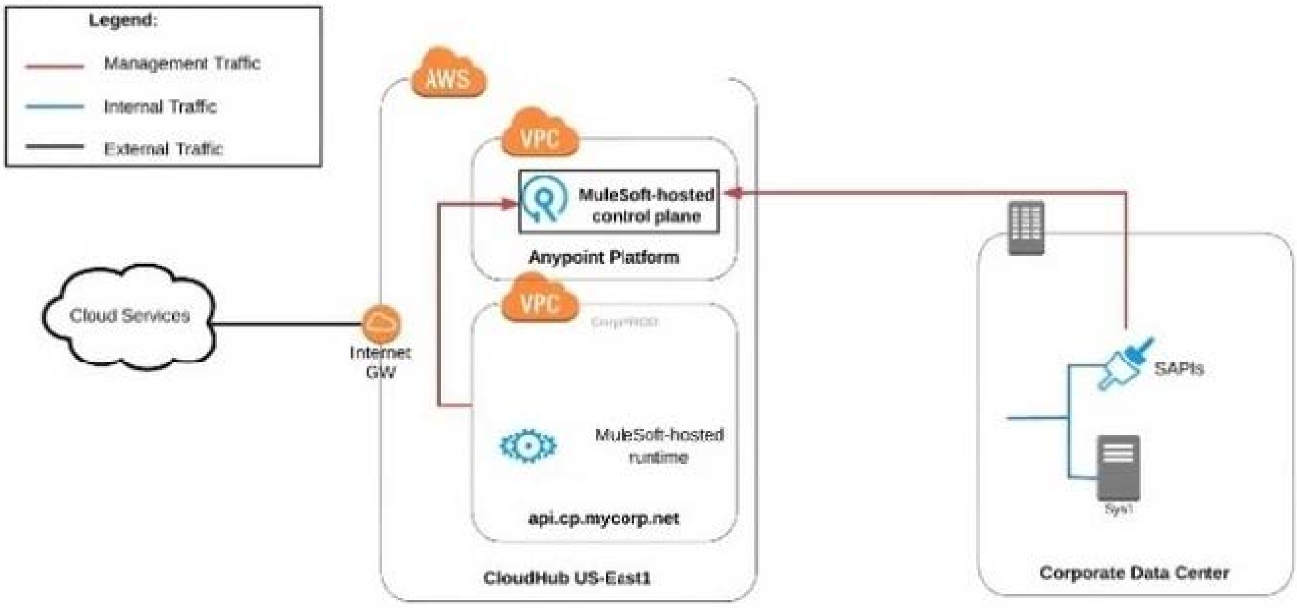

A. Ask the Marketing Department to interact with a mocking implementation of the API using the automatically generated API Console

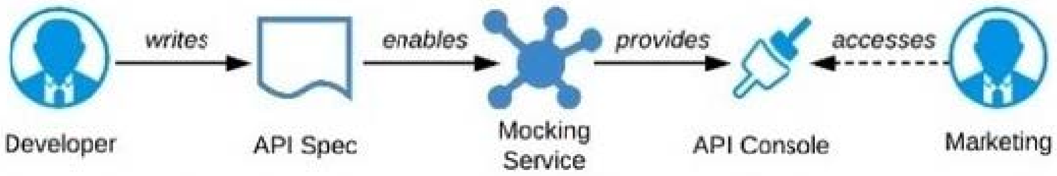

B. Organize a design workshop with the DBAs of the Marketing Department in which the database schema of the Marketing IT systems is translated into RAML

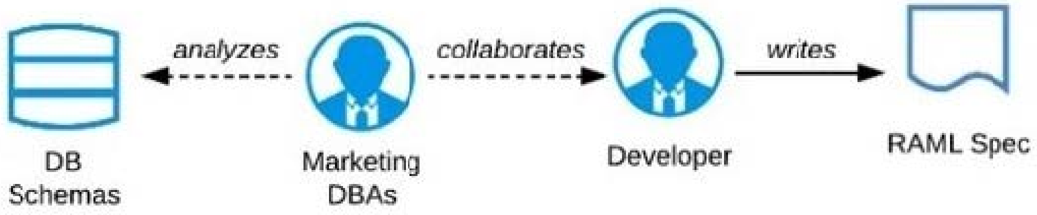

C. Use Anypoint Studio to implement the API as a Mule application, then deploy that API implementation to CloudHub and ask the Marketing Department to interact with it

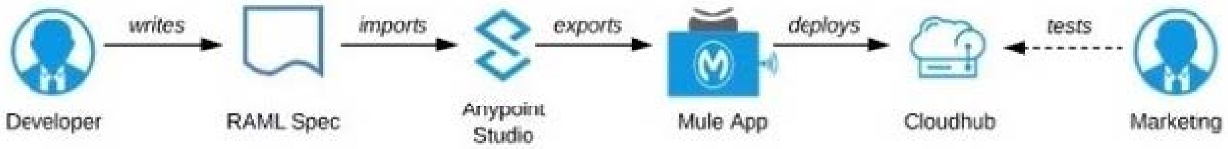

D. Export an integration test suite from API designer and have the Marketing Department execute the tests In that suite to ensure they pass

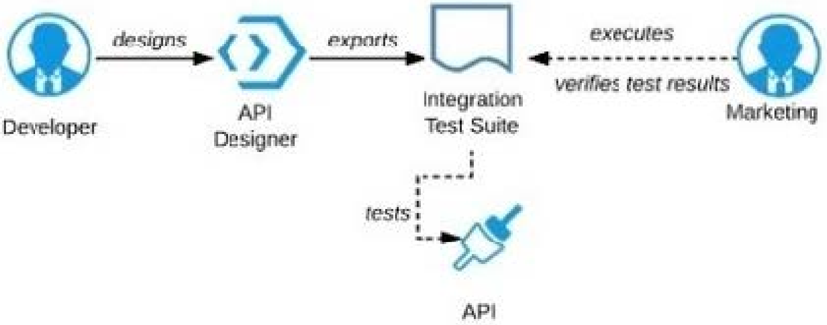

A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: A

Ask the Marketing Department to interact with a mocking implementation of the API using the automatically generated API Console.

****

As per MuleSoft\\'s IT Operating Model:

>> API consumers need NOT wait until the full API implementation is ready. >> NO technical test-suites needs to be shared with end users to interact with APIs. >> Anypoint Platform offers a mocking capability on all the published API

specifications to Anypoint Exchange which also will be rich in documentation covering all details of API functionalities and working nature.

>> No needs of arranging days of workshops with end users for feedback.

API consumers can use Anypoint Exchange features on the platform and interact with the API using its mocking feature. The feedback can be shared quickly on the same to incorporate any changes.

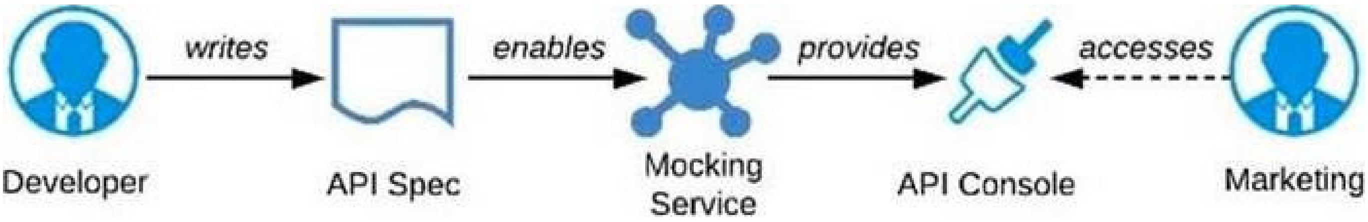

QUESTION 13

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?

A. The API producer should be requested to run the old version in parallel with the new one

B. The API producer should be contacted to understand the change to existing functionality

C. The API client code only needs to be changed if it needs to take advantage of the new features

D. The API clients need to update the code on their side and need to do full regression

Correct Answer: C

QUESTION 14

When must an API implementation be deployed to an Anypoint VPC?

A. When the API Implementation must invoke publicly exposed services that are deployed outside of CloudHub in a customer- managed AWS instance

B. When the API implementation must be accessible within a subnet of a restricted customer-hosted network that does not allow public access

C. When the API implementation must be deployed to a production AWS VPC using the Mule Maven plugin

D. When the API Implementation must write to a persistent Object Store

Correct Answer: A

QUESTION 15

What is a best practice when building System APIs?

A. Document the API using an easily consumable asset like a RAML definition

B. Model all API resources and methods to closely mimic the operations of the backend system

C. Build an Enterprise Data Model (Canonical Data Model) for each backend system and apply it to System APIs

D. Expose to API clients all technical details of the API implementation\\'s interaction wifch the backend system

Correct Answer: B

Model all API resources and methods to closely mimic the operations of the backend system.

****

>> There are NO fixed and straight best practices while opting data models for APIs. They are completly contextual and depends on number of factors. Based upon those factors, an enterprise can choose if they have to go with Enterprise

Canonical Data Model or Bounded Context Model etc.

>> One should NEVER expose the technical details of API implementation to their API clients. Only the API interface/ RAML is exposed to API clients. >> It is true that the RAML definitions of APIs should be as detailed as possible and should

reflect most of the documentation. However, just that is NOT enough to call your API as best documented API. There should be even more documentation on Anypoint Exchange with API Notebooks etc. to make and create a developer

friendly API and repository.. >> The best practice always when creating System APIs is to create their API interfaces by modeling their resources and methods to closely reflect the operations and functionalities of that backend system.