MCPA-LEVEL1 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.passapply.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

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

QUESTION 1

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.

QUESTION 2

Refer to the exhibit.

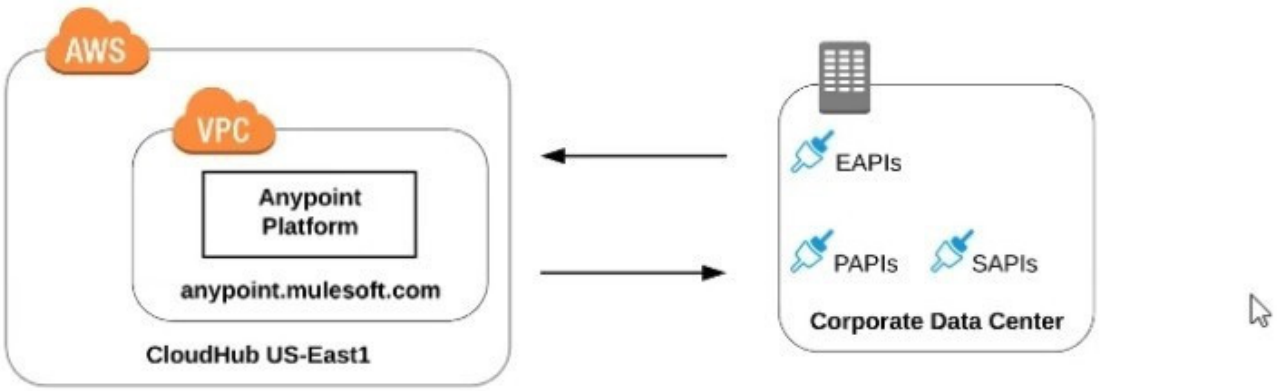

What is true when using customer-hosted Mule runtimes with the MuleSoft-hosted Anypoint Platform control plane

(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

==========

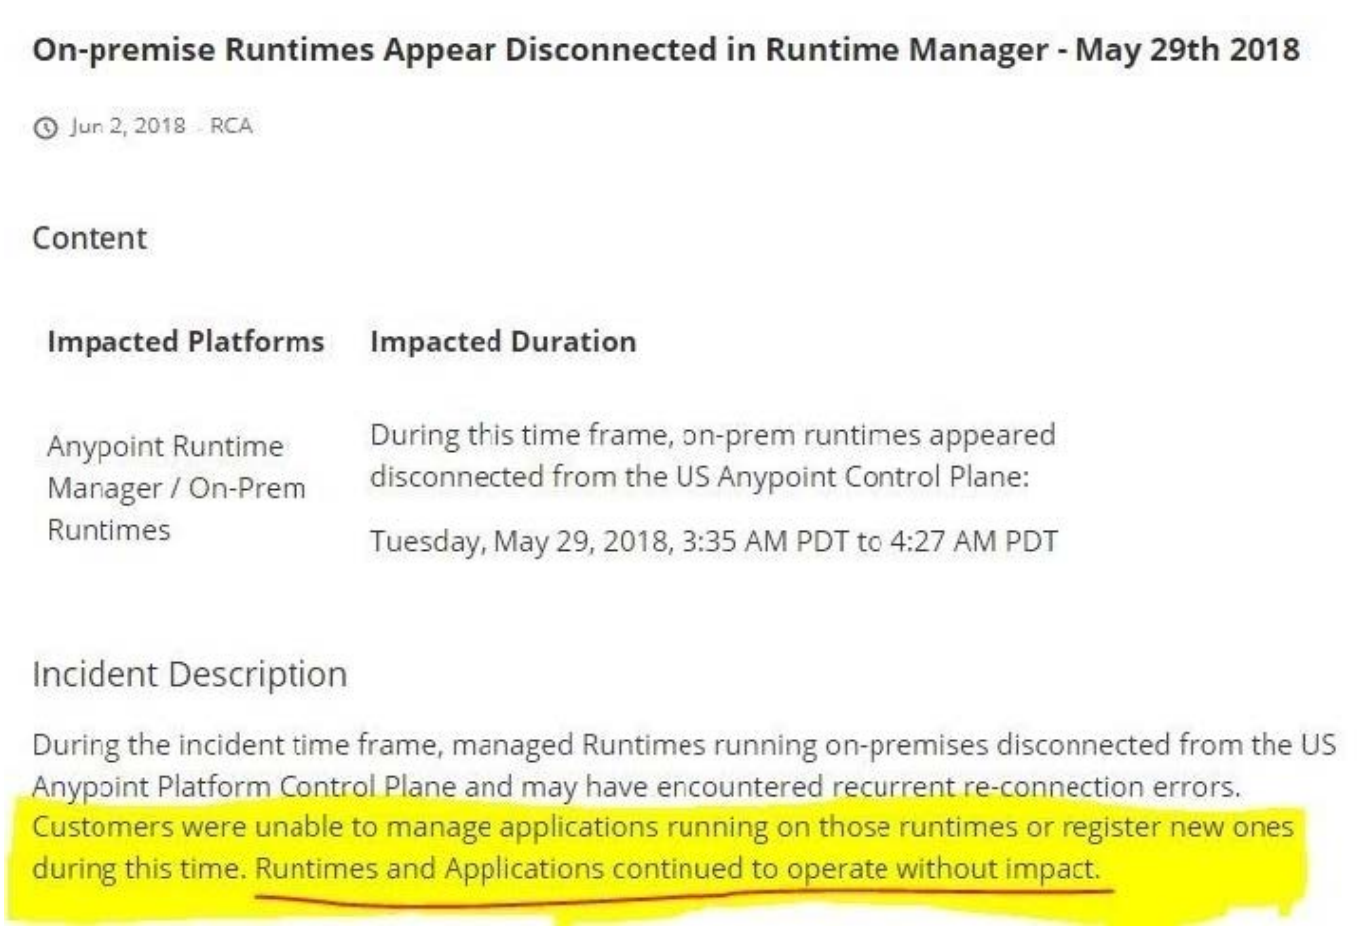

On-premise Runtimes Appear Disconnected in Runtime Manager - May 29th 2018

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 3

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 the API proxy

B. At the API implementation

C. At both the API proxy and the API implementation

D. At a MuleSoft-hosted load balancer

Correct Answer: A

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.

QUESTION 4

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 5

What condition requires using a CloudHub Dedicated Load Balancer?

A. When cross-region load balancing is required between separate deployments of the same Mule application

B. When custom DNS names are required for API implementations deployed to customer- hosted Mule runtimes

C. When API invocations across multiple CloudHub workers must be load balanced

D. When server-side load-balanced TLS mutual authentication is required between API implementations and API clients

Correct Answer: D

When server-side load-balanced TLS mutual authentication is required between API implementations and API clients

****

Fact/ Memory Tip: Although there are many benefits of CloudHub Dedicated Load balancer, TWO important things that should come to ones mind for considering it are:

>> Having URL endpoints with Custom DNS names on CloudHub deployed apps >> Configuring custom certificates for both HTTPS and Two-way (Mutual) authentication.

Coming to the options provided for this question:

>> We CANNOT use DLB to perform cross-region load balancing between separate deployments of the same Mule application.

>> We can have mapping rules to have more than one DLB URL pointing to same Mule app. But vicevera (More than one Mule app having same DLB URL) is NOT POSSIBLE >> It is true that DLB helps to setup custom DNS names for

Cloudhub deployed Mule apps but NOT true for apps deployed to Customer-hosted Mule Runtimes. >> It is true to that we can load balance API invocations across multiple CloudHub workers using DLB but it is NOT A MUST. We can

achieve the same (load balancing) using SLB (Shared Load Balancer) too. We DO NOT necessarily require DLB for achieve it. So the only right option that fits the scenario and requires us to use DLB is when TLS mutual authentication is

required between API implementations and API clients. Reference: https://docs.mulesoft.com/runtime-manager/cloudhub-dedicated-load-balancer