MuleSoft

Exam Questions MCPA-Level-1

MuleSoft Certified Platform Architect - Level 1

NEW QUESTION 1
Refer to the exhibit. An organization is running a Mule standalone runtime and has configured Active Directory as the Anypoint Platform external Identity Provider. The organization does not have budget for other system components.

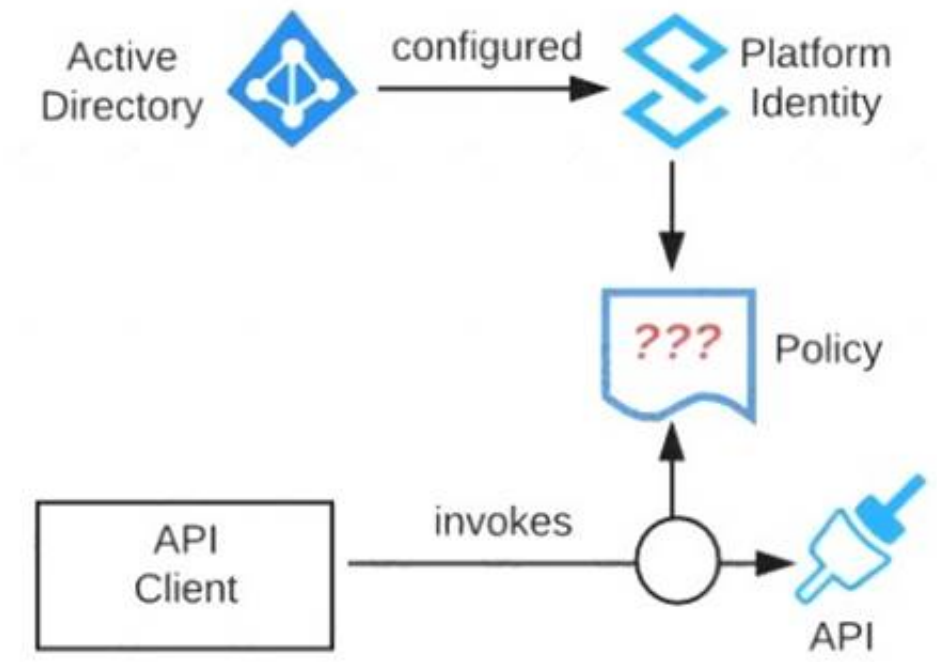

What policy should be applied to all instances of APIs in the organization to most effecuvelyKestrict access to a specific group of internal users?

A. Apply a basic authentication - LDAP policy; the internal Active Directory will be configured as the LDAP source for authenticating users
B. Apply a client ID enforcement policy; the specific group of users will configure their client applications to use their specific client credentials
C. Apply an IP whitelist policy; only the specific users' workstations will be in the whitelist
D. Apply an OAuth 2.0 access token enforcement policy; the internal Active Directory will be configured as the OAuth server

Answer: A

Explanation:
Correct Answer
Apply a basic authentication - LDAP policy; the internal Active Directory will be configured as the LDAP source for authenticating users.
****
>> IP Whitelisting does NOT fit for this purpose. Moreover, the users workstations may not necessarily have static IPs in the network.
>> OAuth 2.0 enforcement requires a client provider which isn't in the organizations system components.
>> It is not an effective approach to let every user create separate client credentials and configure those for their usage.
The effective way it to apply a basic authentication - LDAP policy and the internal Active Directory will be configured as the LDAP source for authenticating users.

NEW QUESTION 2
What Mule application can have API policies applied by
Anypoint Platform to the endpoint exposed by that Mule application?
A) A Mule application that accepts requests over HTTP/1.x

B) A Mule application that accepts JSON requests over TCP but is NOT required to provide a response

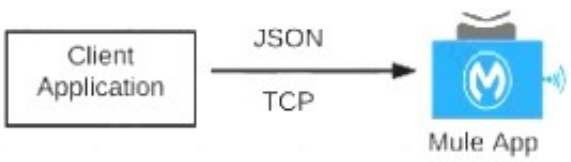

C) A Mute application that accepts JSON requests over WebSocket

D) A Mule application that accepts gRPC requests over HTTP/2

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

Explanation:

Correct Answer
Option A
****
>> Anypoint API Manager and API policies are applicable to all types of HTTP/1.x APIs.
>> They are not applicable to WebSocket APIs, HTTP/2 APIs and gRPC APIs

NEW QUESTION 3
An organization uses various cloud-based SaaS systems and multiple on-premises systems. The on-premises systems are an important part of the organization's application network and can only be accessed from within the organization's intranet.
What is the best way to configure and use Anypoint Platform to support integrations with both the cloud-based SaaS systems and on-premises systems?
A) Use CloudHub-deployed Mule runtimes in an Anypoint VPC managed by Anypoint Platform Private Cloud Edition control plane

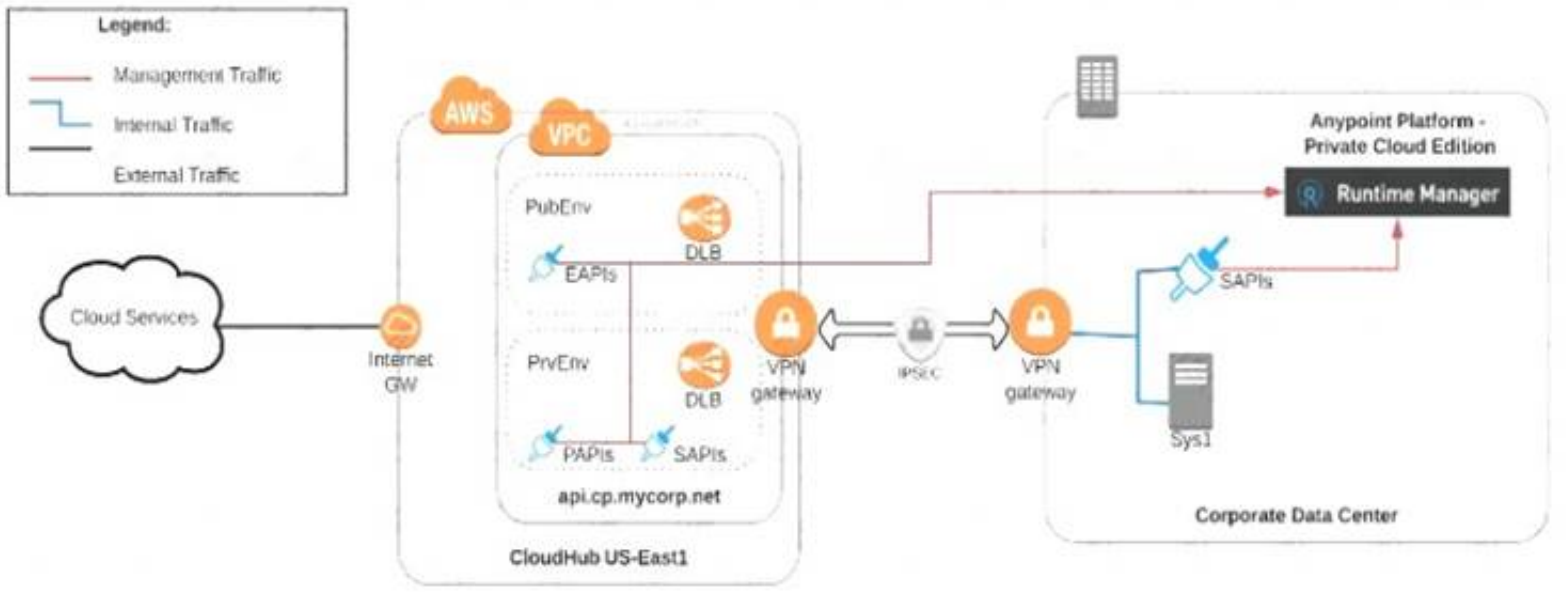

B) Use CloudHub-deployed Mule runtimes in the shared worker cloud managed by the MuleSoft-hosted Anypoint Platform control plane

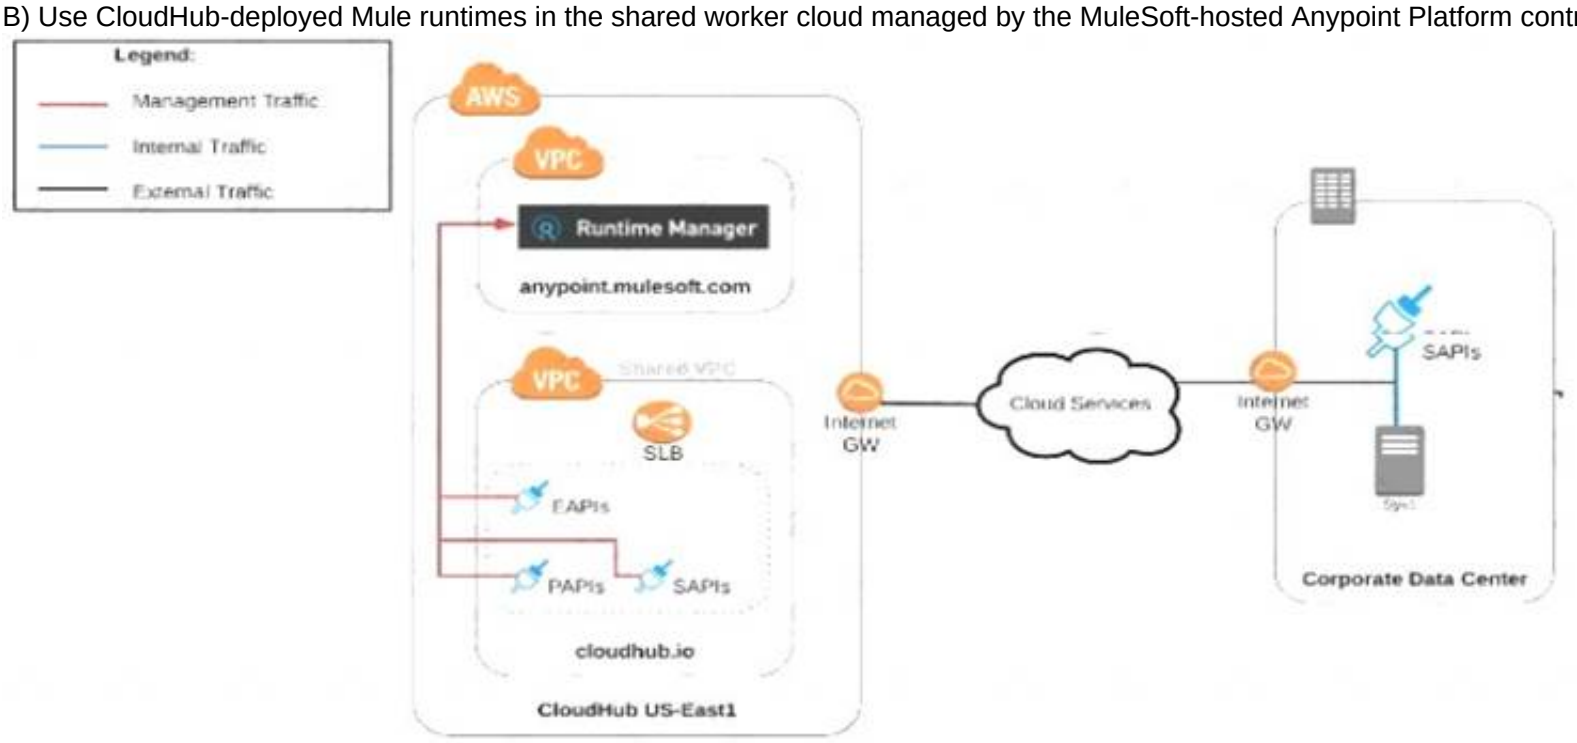

C) Use an on-premises installation of Mule runtimes that are completely isolated with NO external network access, managed by the Anypoint Platform Private Cloud Edition control plane

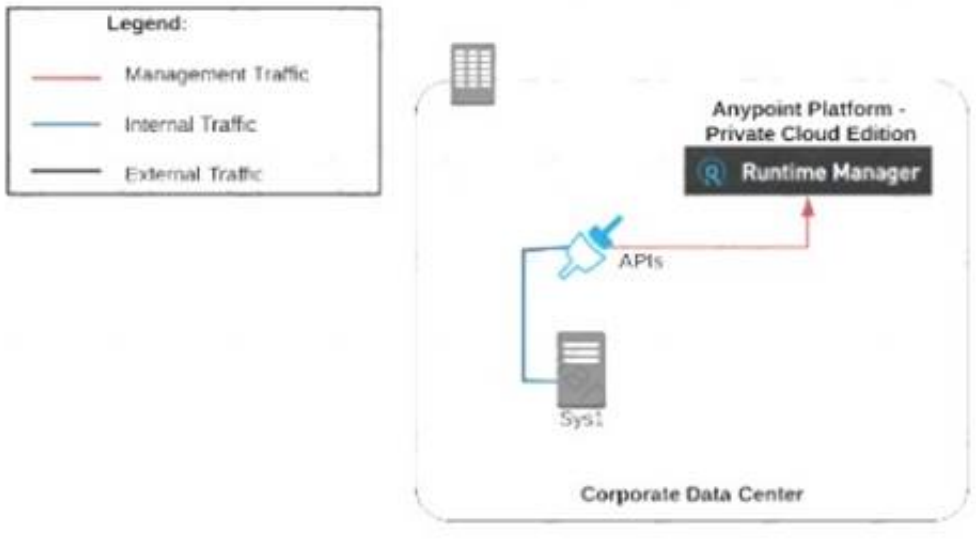

D) Use a combination of Cloud Hub-deployed and manually provisioned on-premises Mule runtimes managed by the MuleSoft-hosted Anypoint Platform control plane

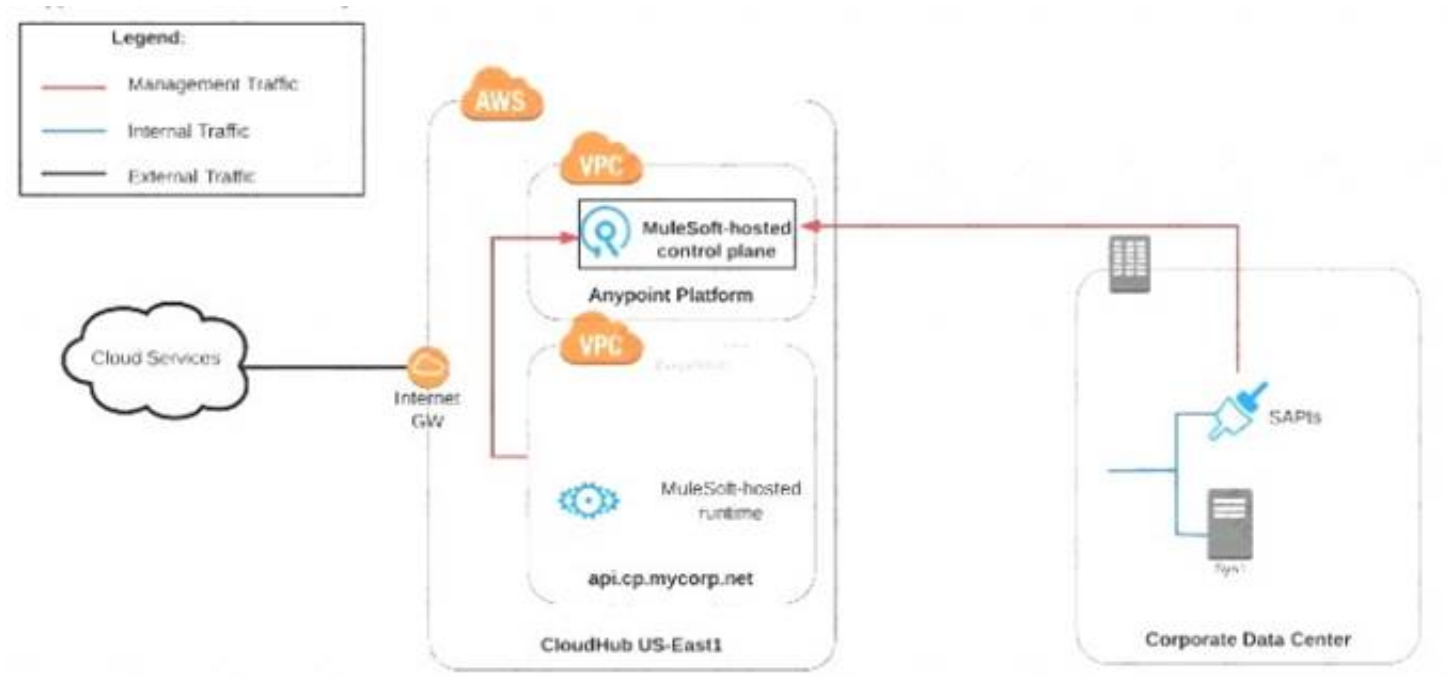

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:
Correct Answer
Use a combination of CloudHub-deployed and manually provisioned on-premises Mule runtimes managed by the MuleSoft-hosted Platform control plane.
** Key details to be taken from the given scenario:
>> Organization uses BOTH cloud-based and on-premises systems
>> On-premises systems can only be accessed from within the organization's intranet Let us evaluate the given choices based on above key details:
>> CloudHub-deployed Mule runtimes can ONLY be controlled using MuleSoft-hosted control plane. We CANNOT use Private Cloud Edition's control plane to control CloudHub Mule Runtimes. So, option suggesting this is INVALID
>> Using CloudHub-deployed Mule runtimes in the shared worker cloud managed by the MuleSoft-hosted Anypoint Platform is completely IRRELEVANT to given scenario and silly choice. So, option suggesting this is INVALID
>> Using 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 would work for On-premises integrations. However, with NO external access, integrations cannot be done to SaaS-based apps. Moreover CloudHub-hosted apps are best-fit for integrating with SaaS-based applications. So, option suggesting this is BEST WAY.
The best way to configure and use Anypoint Platform to support these mixed/hybrid integrations is to use a combination of CloudHub-deployed and manually provisioned on-premises Mule runtimes managed by the MuleSoft-hosted Platform control plane.

NEW QUESTION 4
Refer to the exhibit.

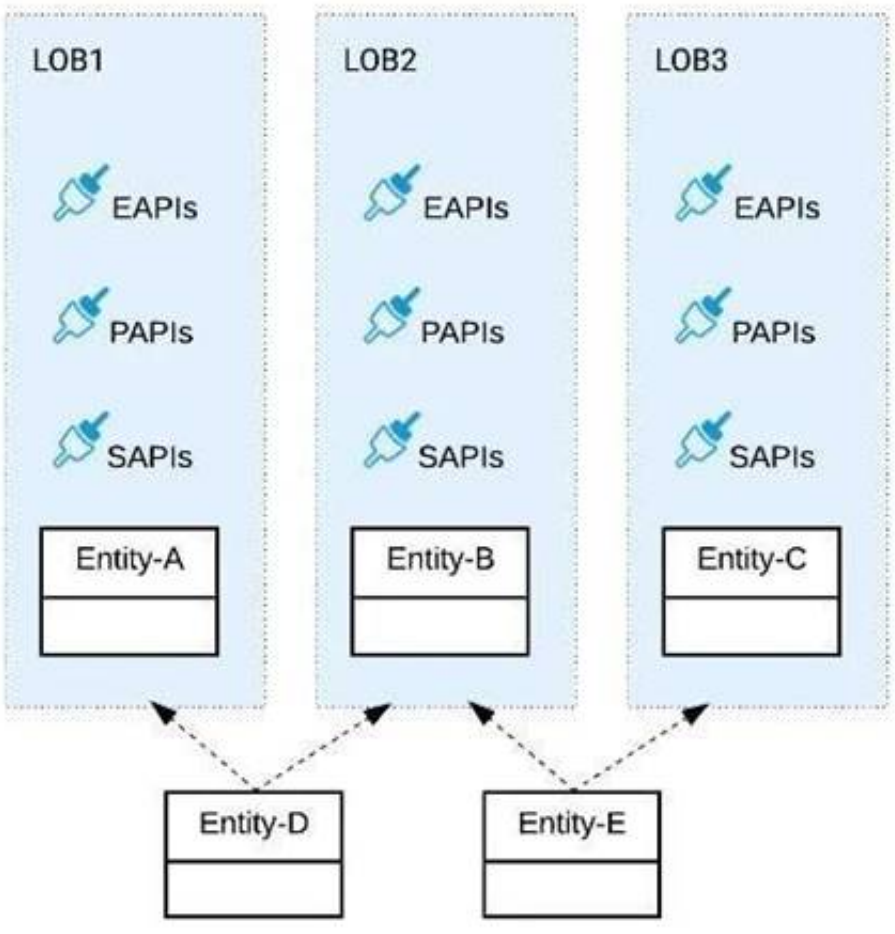

Three business processes need to be implemented, and the implementations need to communicate with several different SaaS applications.
These processes are owned by separate (siloed) LOBs and are mainly independent of each other, but do share a few business entities. Each LOB has one development team and their own budget
In this organizational context, what is the most effective approach to choose the API data models for the APIs that will implement these business processes with minimal redundancy of the data models?
A) Build several Bounded Context Data Models that align with coherent parts of the business processes and the definitions of associated business entities

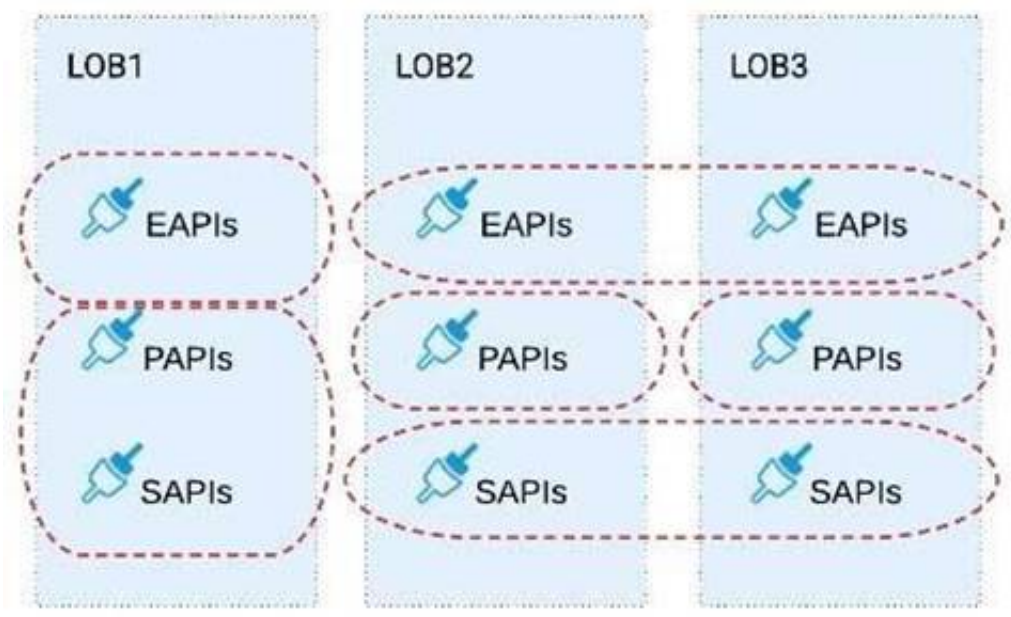

B) Build distinct data models for each API to follow established micro-services and Agile API-centric practices

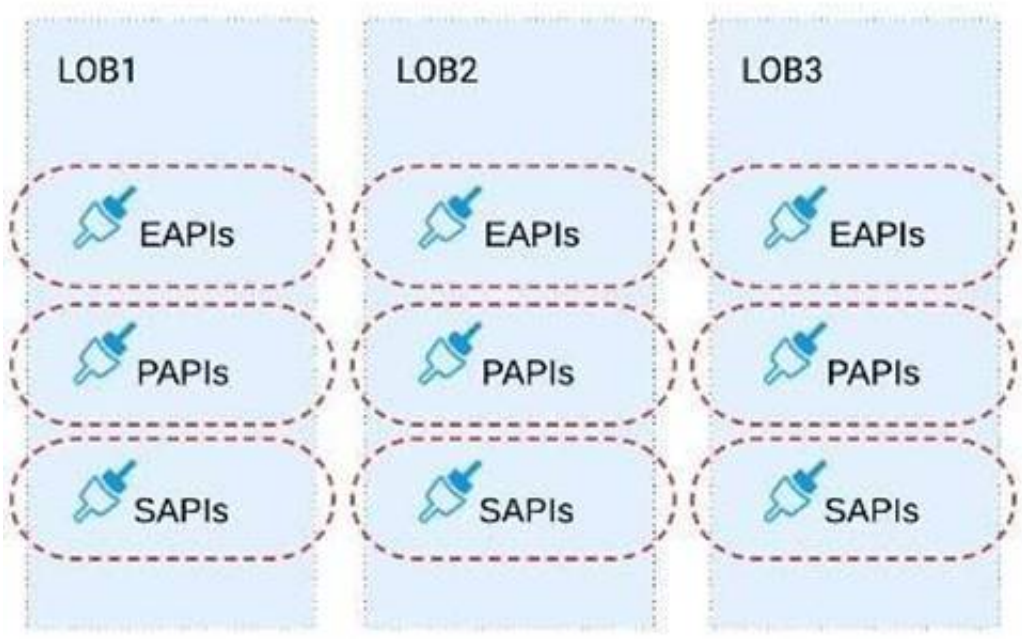

C) Build all API data models using XML schema to drive consistency and reuse across the organization

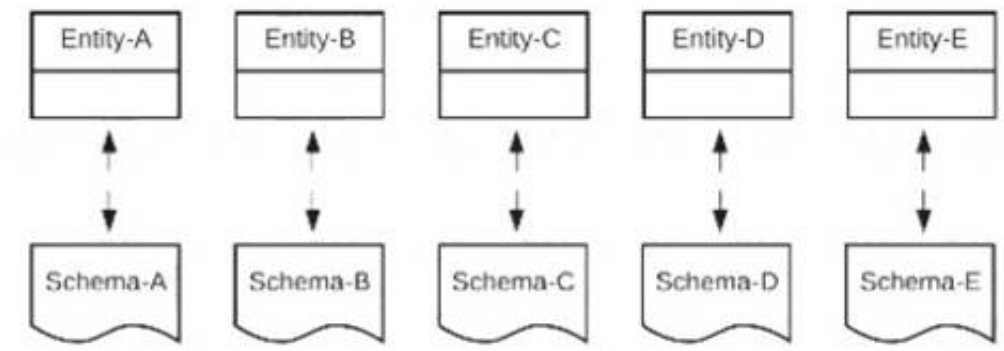

D) Build one centralized Canonical Data Model (Enterprise Data Model) that unifies all the data types from all three business processes, ensuring the data model is consistent and non-redundant

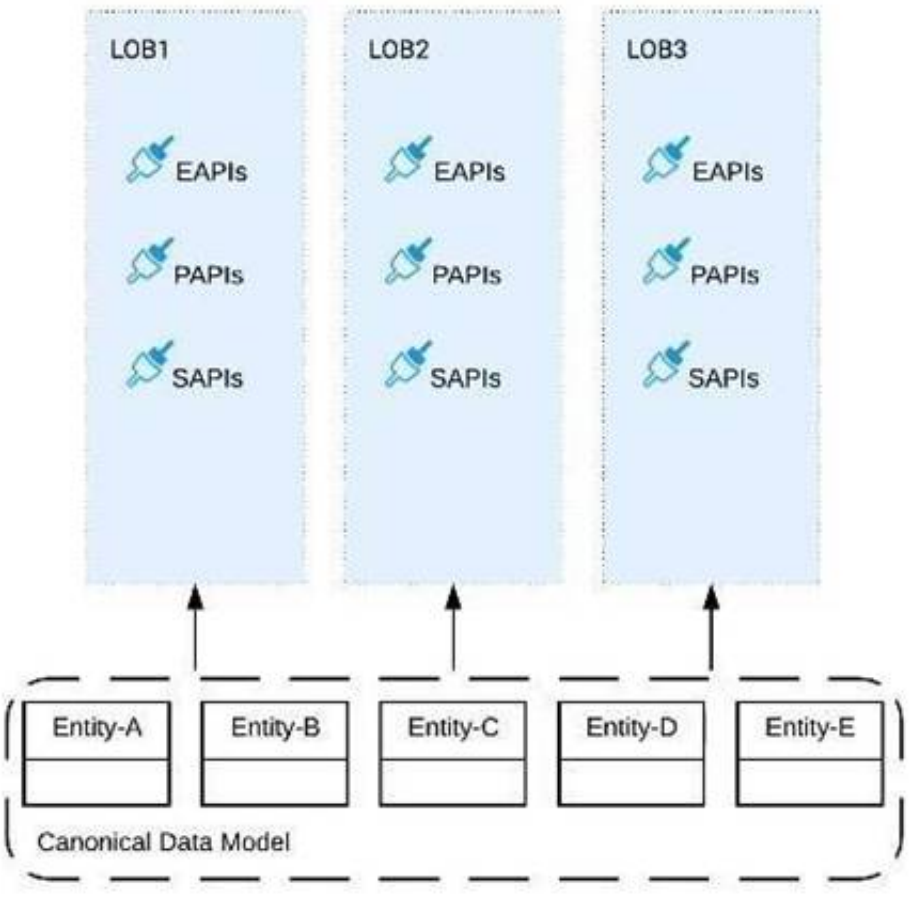

A. Option A

B. Option B
C. Option C
D. Option D

Answer: A

Explanation:

Correct Answer
Build several Bounded Context Data Models that align with coherent parts of the business processes and the definitions of associated business entities.
****
>> The options w.r.t building API data models using XML schema/ Agile API-centric practices are irrelevant to the scenario given in the question. So these two are INVALID.
>> Building EDM (Enterprise Data Model) is not feasible or right fit for this scenario as the teams and LOBs work in silo and they all have different initiatives, budget etc.. Building EDM needs intensive coordination among all the team which evidently seems not possible in this scenario.
So, the right fit for this scenario is to build several Bounded Context Data Models that align with coherent parts of the business processes and the definitions of associated business entities.

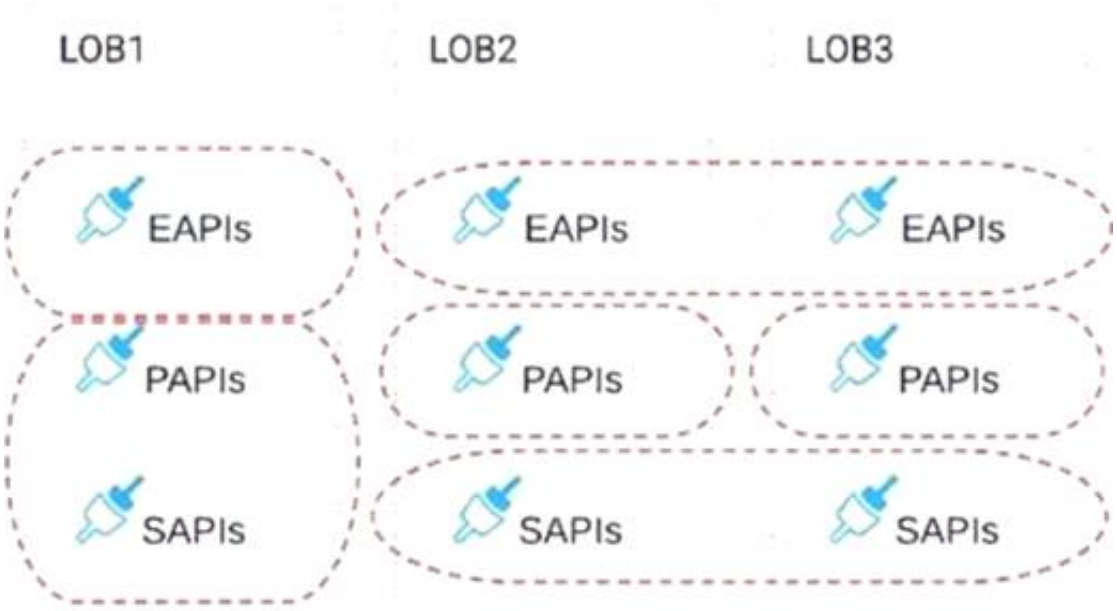

NEW QUESTION 5
An API implementation is deployed to CloudHub.
What conditions can be alerted on using the default Anypoint Platform functionality, where the alert conditions depend on the end-to-end request processing of the API implementation?

A. When the API is invoked by an unrecognized API client
B. When a particular API client invokes the API too often within a given time period
C. When the response time of API invocations exceeds a threshold
D. When the API receives a very high number of API invocations

Answer: C

Explanation:

Correct Answer
When the response time of API invocations exceeds a threshold
****
>> Alerts can be setup for all the given options using the default Anypoint Platform functionality
>> However, the question insists on an alert whose conditions depend on the end-to-end request processing of the API implementation.
>> Alert w.r.t "Response Times" is the only one which requires end-to-end request processing of API implementation in order to determine if the threshold is exceeded or not.

NEW QUESTION 6
In an organization, the InfoSec team is investigating Anypoint Platform related data traffic.
From where does most of the data available to Anypoint Platform for monitoring and alerting originate?

A. From the Mule runtime or the API implementation, depending on the deployment model
B. From various components of Anypoint Platform, such as the Shared Load Balancer, VPC, and Mule runtimes
C. From the Mule runtime or the API Manager, depending on the type of data
D. From the Mule runtime irrespective of the deployment model

Answer: D

Explanation:

Correct Answer
From the Mule runtime irrespective of the deployment model
****
>> Monitoring and Alerting metrics are always originated from Mule Runtimes irrespective of the deployment model.
>> It may seems that some metrics (Runtime Manager) are originated from Mule Runtime and some are (API Invocations/ API Analytics) from API Manager. However, this is realistically NOT TRUE. The reason is, API manager is just a management tool for API instances but all policies upon applying on APIs eventually gets executed on Mule Runtimes only (Either Embedded or API Proxy).
>> Similarly all API Implementations also run on Mule Runtimes.
So, most of the day required for monitoring and alerts are originated fron Mule Runtimes only irrespective of whether the deployment model is MuleSoft-hosted or

Customer-hosted or Hybrid.

NEW QUESTION 7
A Mule application exposes an HTTPS endpoint and is deployed to the CloudHub Shared Worker Cloud. All traffic to that Mule application must stay inside the AWS VPC.
To what TCP port do API invocations to that Mule application need to be sent?

A. 443
B. 8081
C. 8091
D. 8082

Answer: D

Explanation:
Correct Answer 8082
****
>> 8091 and 8092 ports are to be used when keeping your HTTP and HTTPS app private to the LOCAL VPC respectively.
>> Above TWO ports are not for Shared AWS VPC/ Shared Worker Cloud.
>> 8081 is to be used when exposing your HTTP endpoint app to the internet through Shared LB
>> 8082 is to be used when exposing your HTTPS endpoint app to the internet through Shared LB So, API invocations should be sent to port 8082 when calling this HTTPS based app.
References:
https://docs.mulesoft.com/runtime-manager/cloudhub-networking-guide https://help.mulesoft.com/s/article/Configure-Cloudhub-Application-to-Send-a-HTTPS-Request-Directly-to-An
https://help.mulesoft.com/s/question/0D52T00004mXXULSA4/multiple-http-listerners-on-cloudhub-one-with-p

NEW QUESTION 8
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 deployedon-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

Answer: A

Explanation:
Correct Answer
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..

NEW QUESTION 9
What are 4 important Platform Capabilities offered by Anypoint Platform?

A. API Versioning, API Runtime Execution and Hosting, API Invocation, API Consumer Engagement
B. API Design and Development, API Runtime Execution and Hosting, API Versioning, API Deprecation
C. API Design and Development, API Runtime Execution and Hosting, API Operations and Management, API Consumer Engagement
D. API Design and Development, API Deprecation, API Versioning, API Consumer Engagement

Answer: C

Explanation:

Correct Answer
API Design and Development, API Runtime Execution and Hosting, API Operations and Management, API Consumer Engagement
****
>> API Design and Development - Anypoint Studio, Anypoint Design Center, Anypoint Connectors
>> API Runtime Execution and Hosting - Mule Runtimes, CloudHub, Runtime Services
>> API Operations and Management - Anypoint API Manager, Anypoint Exchange
>> API Consumer Management - API Contracts, Public Portals, Anypoint Exchange, API Notebooks

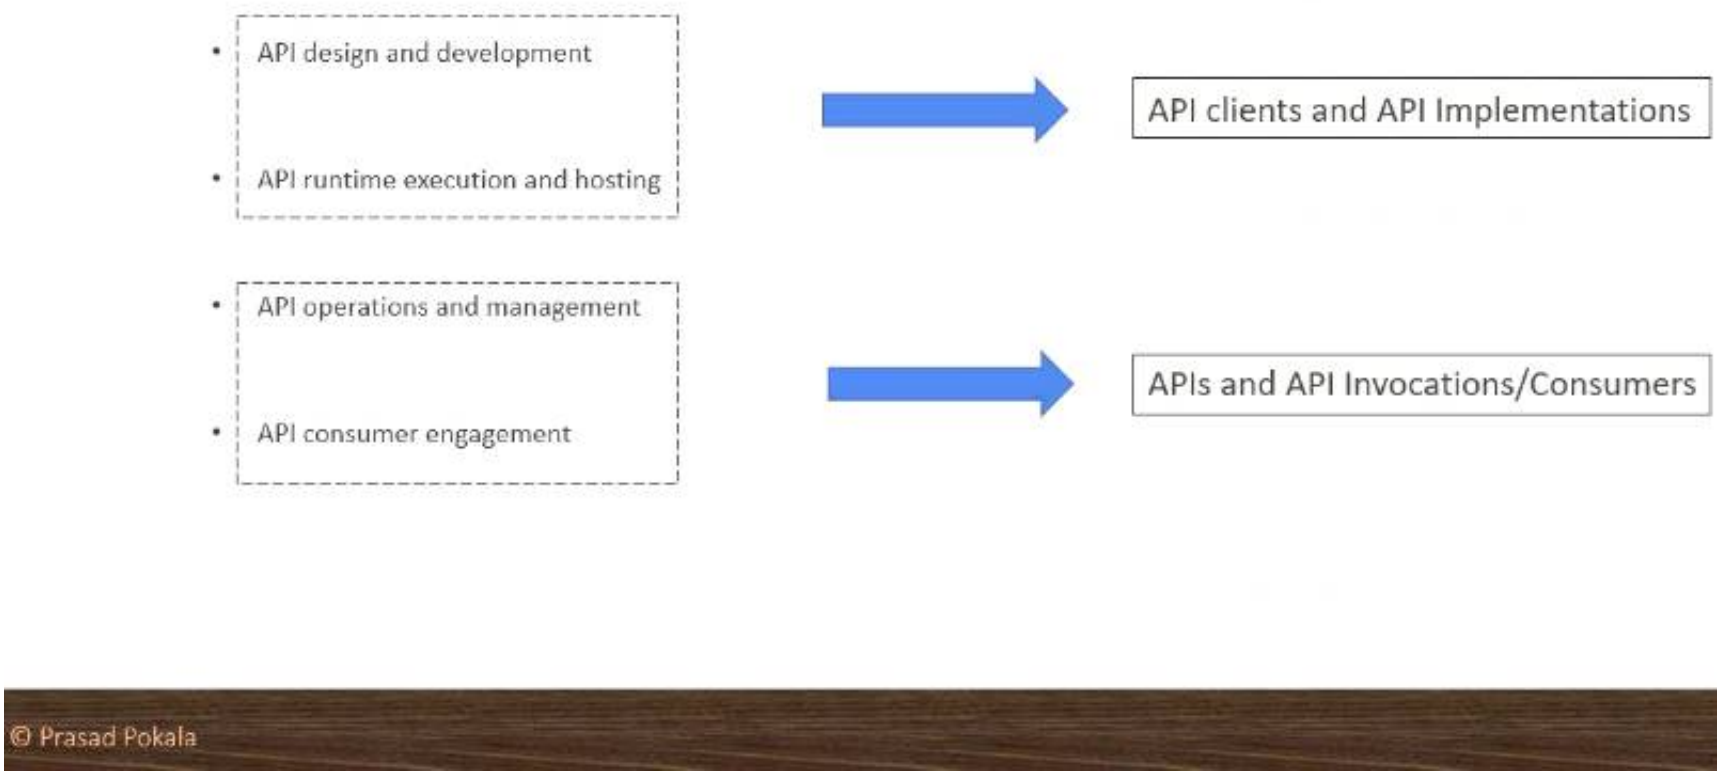

NEW QUESTION 10
When must an API implementation be deployed to an Anypoint VPC?

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

Answer: A

NEW QUESTION 10
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

Answer: A

Explanation:
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 15
Question 10: Skipped
An API implementation returns three X-RateLimit-* HTTP response headers to a requesting API client. What type of information do these response headers indicate to the API client?

A. The error codes that result from throttling
B. A correlation ID that should be sent in the next request
C. The HTTP response size
D. The remaining capacity allowed by the API implementation

Answer: D

Explanation:

Correct Answer
The remaining capacity allowed by the API implementation.
****
>> Reference:

https://docs.mulesoft.com/api-manager/2.x/rate-limiting-and-throttling-sla-based-policies#response-headers

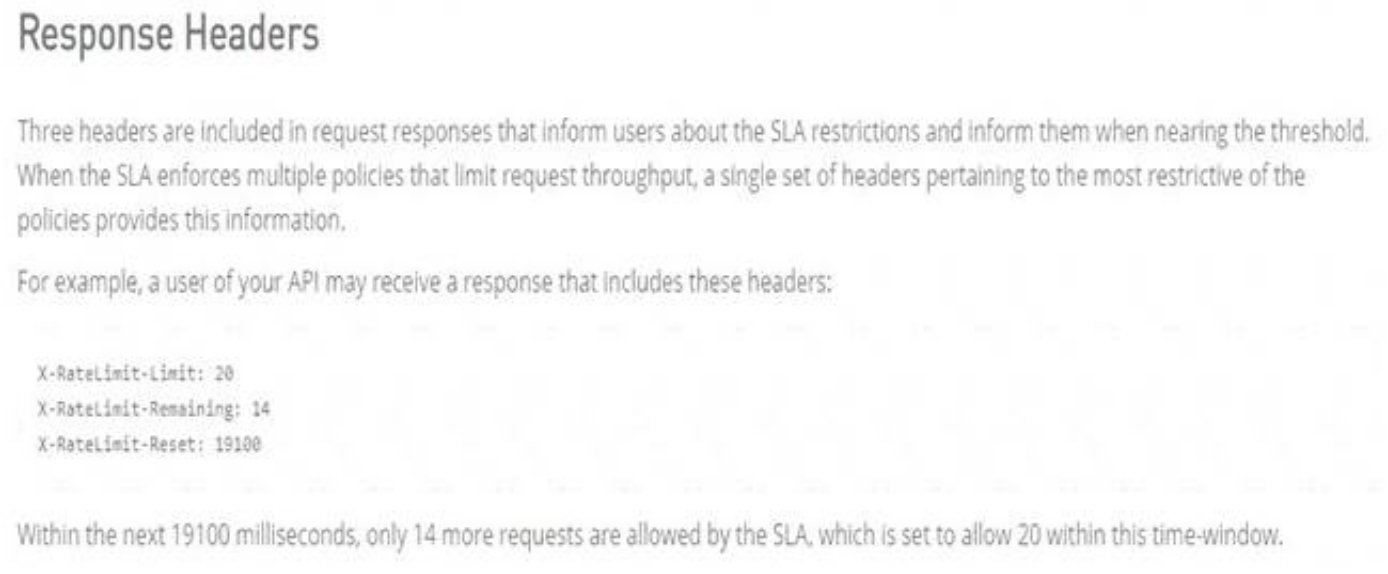

Response Headers

Three headers are included in request responses that inform users about the SLA restrictions and inform them when nearing the threshold. When the SLA enforces multiple policies that limit request throughput, a single set of headers pertaining to the most restrictive of the policies provides this information.

For example, a user of your API may receive a response that includes these headers:

```
X-RateLimit-Limit: 20
X-RateLimit-Remaining: 14
X-RateLimit-Reset: 19100
```

Within the next 19100 milliseconds, only 14 more requests are allowed by the SLA, which is set to allow 20 within this time-window.

NEW QUESTION 17
An Order API must be designed that contains significant amounts of integration logic and involves the invocation of the Product API.
The power relationship between Order API and Product API is one of "Customer/Supplier", because the Product API is used heavily throughout the organization and is developed by a dedicated development team located in the office of the CTO.
What strategy should be used to deal with the API data model of the Product API within the Order API?

A. Convince the development team of the Product API to adopt the API data model of the Order API such that the integration logic of the Order API can work with one consistent internal data model
B. Work with the API data types of the Product API directly when implementing the integration logic of the Order API such that the Order API uses the same (unchanged) data types as the Product API
C. Implement an anti-corruption layer in the Order API that transforms the Product API data model into internal data types of the Order API
D. Start an organization-wide data modeling initiative that will result in an Enterprise Data Model that will then be used in both the Product API and the Order API

Answer: C

Explanation:

Correct Answer
Convince the development team of the product API to adopt the API data model of the Order API such that integration logic of the Order API can work with one consistent internal data model
********** Key details to note from the given scenario:
>> Power relationship between Order API and Product API is customer/supplier
So, as per below rules of "Power Relationships", the caller (in this case Order API) would request for features to the called (Product API team) and the Product API team would need to accomodate those requests.

NEW QUESTION 19
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?

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

Answer: A

Explanation:

Correct Answer
An SLA for the upstream API CANNOT be provided.
****
>> First thing first, the default HTTP response timeout for HTTP connector is 10000 ms (10 seconds). NOT 500 ms.
>> Mule runtime does NOT apply any such "load-dependent" timeouts. There is no such behavior currently in Mule.
>> As there is default 10000 ms time out for HTTP connector, we CANNOT always guarantee that the invocation of the downstream API will run to completion without timing out due to its unreliable SLA times. If the response time crosses 10 seconds then the request may time out.
The main impact due to this is that a proper SLA for the upstream API CANNOT be provided.

NEW QUESTION 23
What do the API invocation metrics provided by Anypoint Platform provide?

A. ROI metrics from APIs that can be directly shared with business users
B. Measurements of the effectiveness of the application network based on the level of reuse
C. Data on past API invocations to help identify anomalies and usage patterns across various APIs
D. Proactive identification of likely future policy violations that exceed a given threat threshold

Answer: C

Explanation:

Correct Answer
Data on past API invocations to help identify anomalies and usage patterns across various APIs
****
API Invocation metrics provided by Anypoint Platform:
>> Does NOT provide any Return Of Investment (ROI) related information. So the option suggesting it is OUT.
>> Does NOT provide any information w.r.t how APIs are reused, whether there is effective usage of APIs or not etc...
>> Does NOT prodive any prediction information as such to help us proactively identify any future policy violations.
So, the kind of data/information we can get from such metrics is on past API invocations to help identify anomalies and usage patterns across various APIs.

NEW QUESTION 26
What correctly characterizes unit tests of Mule applications?

A. They test the validity of input and output of source and target systems
B. They must be run in a unit testing environment with dedicated Mule runtimes for the environment
C. They must be triggered by an external client tool or event source
D. They are typically written using MUnit to run in an embedded Mule runtime that does not require external connectivity

Answer: D

Explanation:
Correct Answer
They are typically written using MUnit to run in an embedded Mule runtime that does not require external connectivity.
****
Below TWO are characteristics of Integration Tests but NOT unit tests:
>> They test the validity of input and output of source and target systems.
>> They must be triggered by an external client tool or event source.
It is NOT TRUE that Unit Tests must be run in a unit testing environment with dedicated Mule runtimes for the environment.
MuleSoft offers MUnit for writing Unit Tests and they run in an embedded Mule Runtime without needing any separate/ dedicated Runtimes to execute them. They also do NOT need any external connectivity as MUnit supports mocking via stubs.
https://dzone.com/articles/munit-framework

NEW QUESTION 29
Refer to the exhibit.

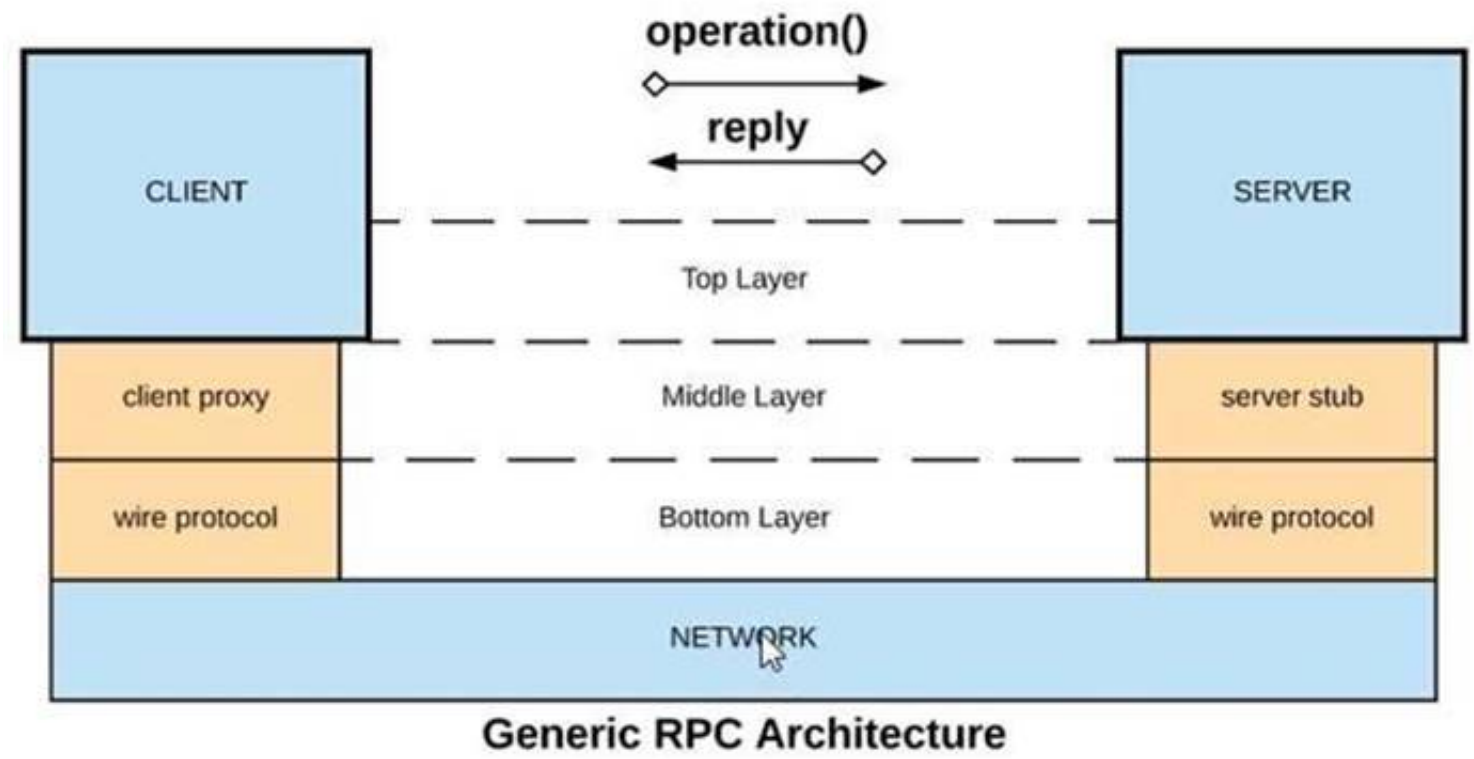

Generic RPC Architecture

What is a valid API in the sense of API-led connectivity and application networks?
A) Java RMI over TCP

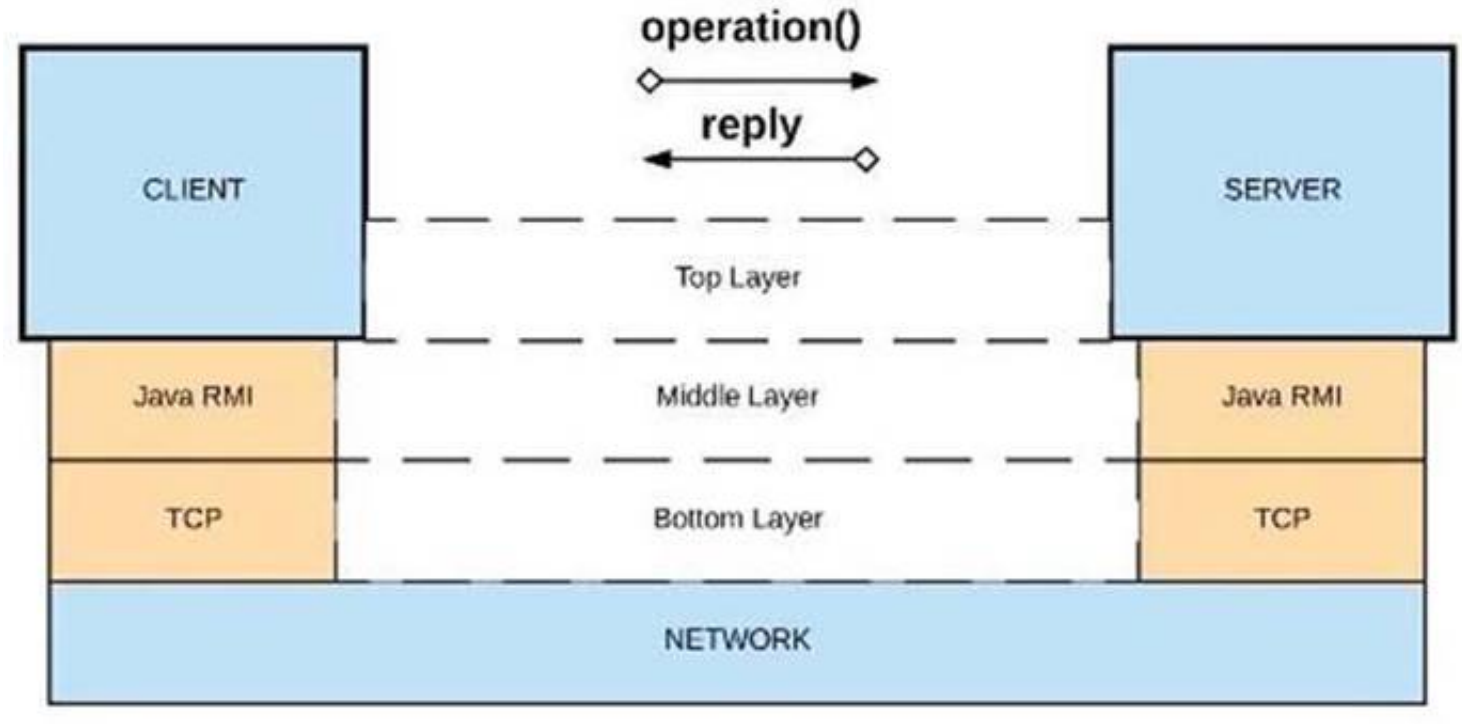

B) Java RMI over TCP

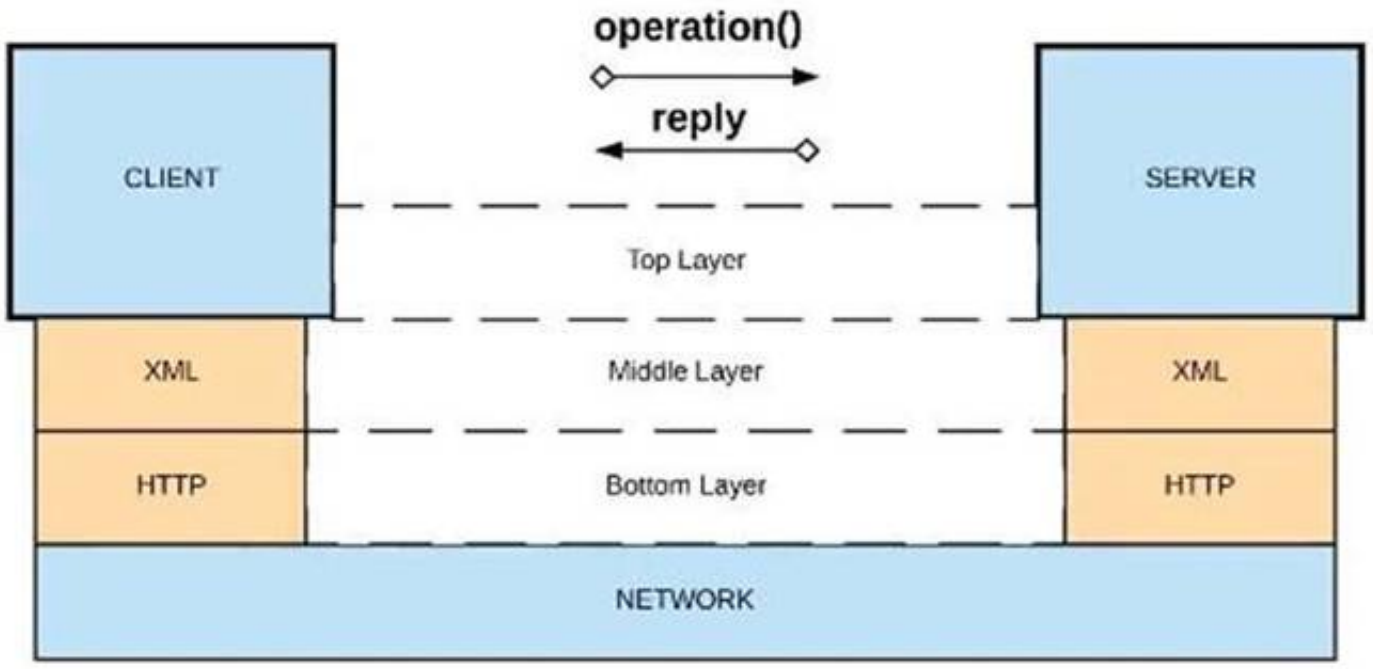

C) CORBA over IIOP

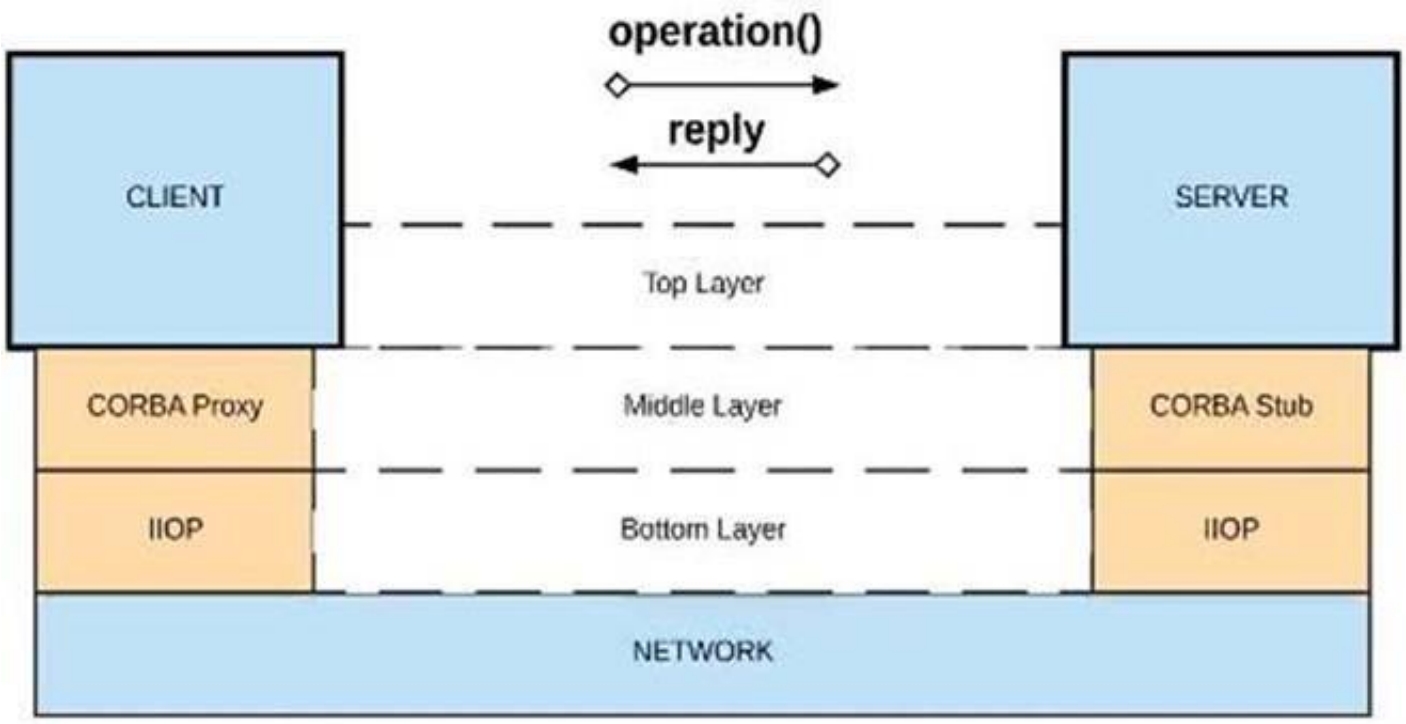

D) XML over UDP

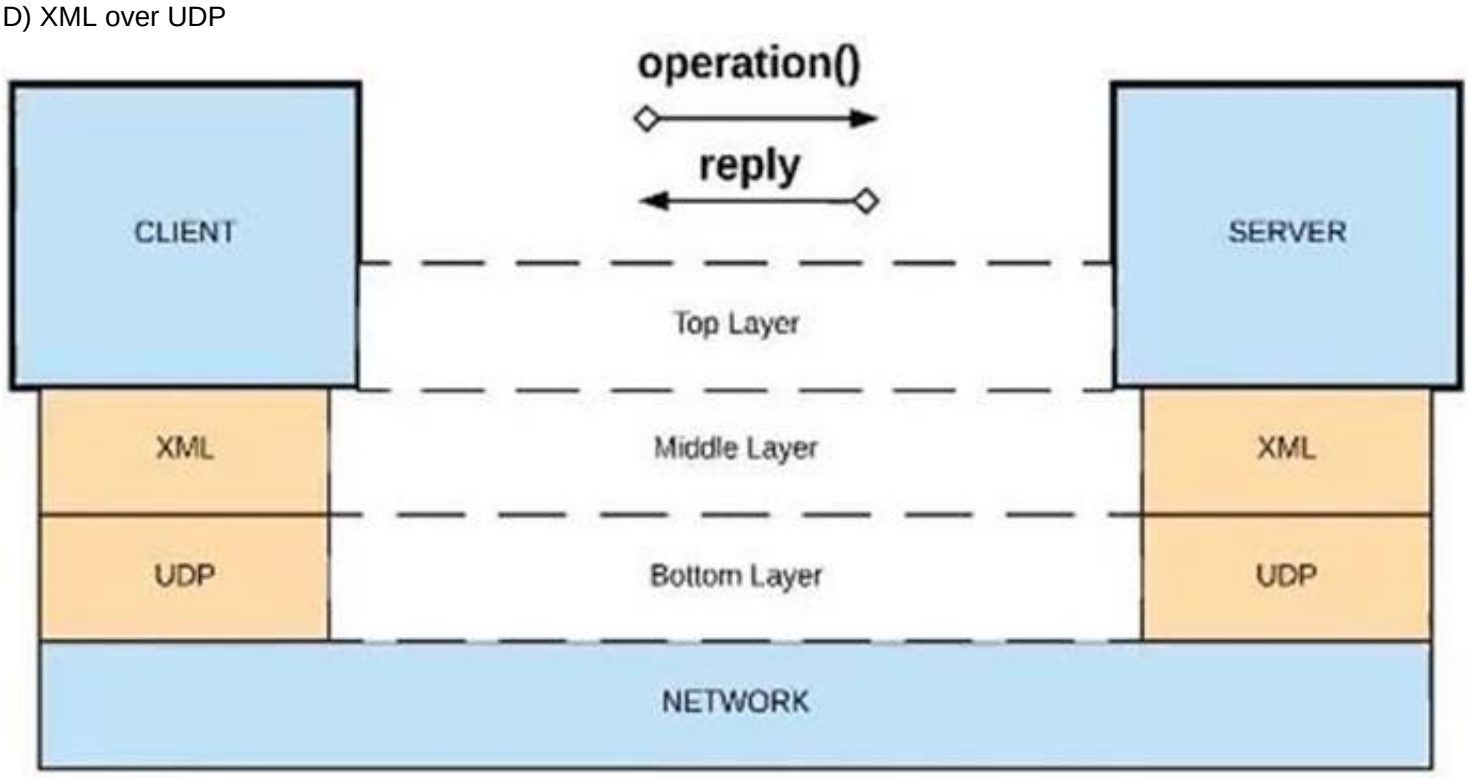

A. Option A
B. Option B
C. Option C
D. Option D

Answer: D

Explanation:

\Correct Answer
XML over HTTP
****
>> API-led connectivity and Application Networks urge to have the APIs on HTTP based protocols for building most effective APIs and networks on top of them.
>> The HTTP based APIs allow the platform to apply various varities of policies to address many NFRs
>> The HTTP based APIs also allow to implement many standard and effective implementation patterns that adhere to HTTP based w3c rules.
Bottom of Form Top of Form

NEW QUESTION 33
Refer to the exhibit.

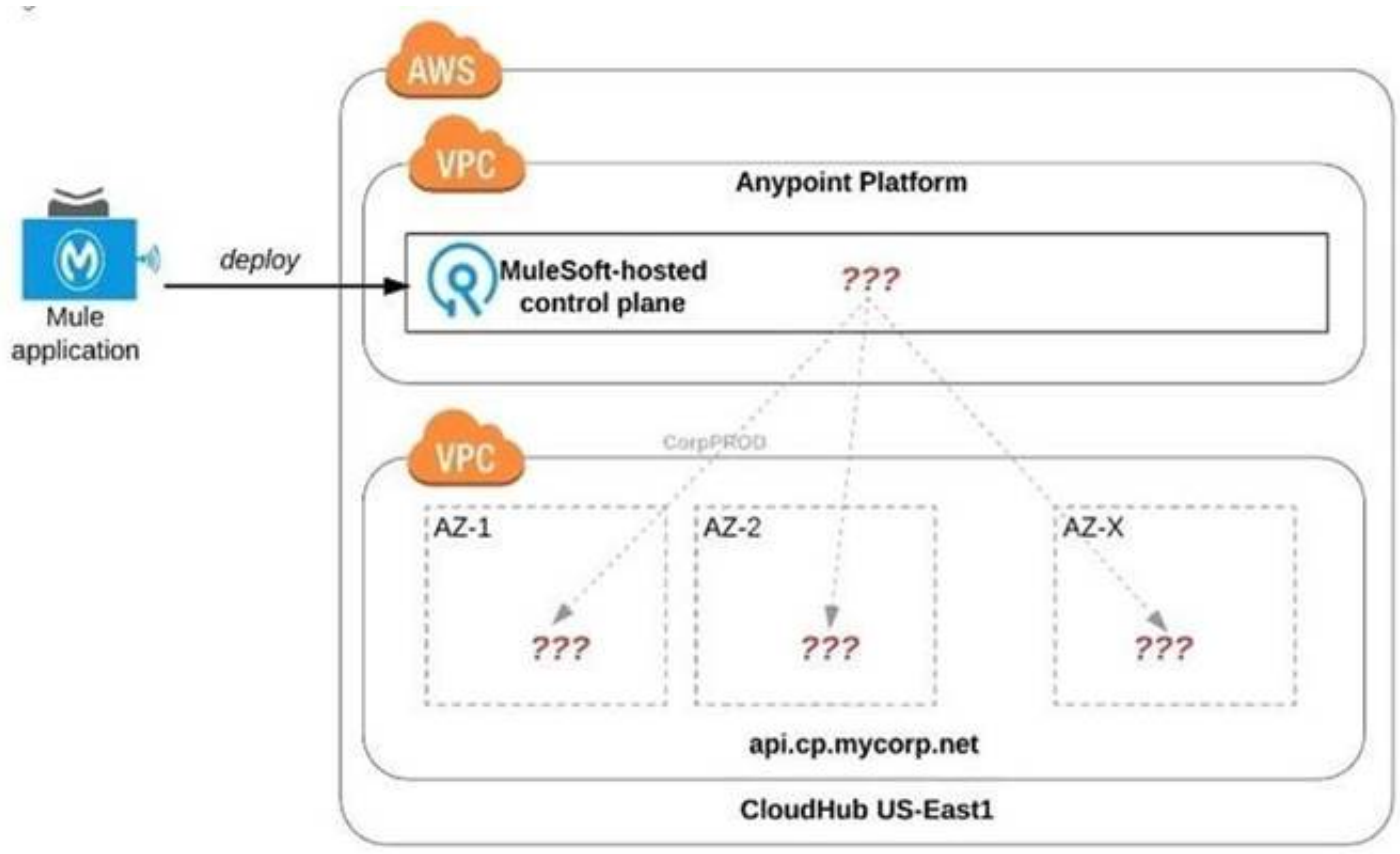

An organization uses one specific CloudHub (AWS) region for all CloudHub deployments.
How are CloudHub workers assigned to availability zones (AZs) when the organization's Mule applications are deployed to CloudHub in that region?

A. Workers belonging to a given environment are assigned to the same AZ within that region
B. AZs are selected as part of the Mule application's deployment configuration
C. Workers are randomly distributed across available AZs within that region
D. An AZ is randomly selected for a Mule application, and all the Mule application's CloudHub workers are assigned to that one AZ

Answer: D

Explanation:
Correct Answer
Workers are randomly distributed across available AZs within that region.
****
>> Currently, we only have control to choose which AWS Region to choose but there is no control at all using any configurations or deployment options to decide what Availability Zone (AZ) to assign to what worker.
>> There arNe O
fixed or implicit rules on platform too w.r.t assignment of AZ to workers based on
environment or application.
>> They are completely assigned inrandom. However, cloudhub definitely ensures that HA is achieved by assigning the workers to more than on AZ so that all workers are not assigned to same AZ for same application.

NEW QUESTION 37
What API policy would LEAST likely be applied to a Process API?

A. Custom circuit breaker
B. Client ID enforcement
C. Rate limiting
D. JSON threat protection

Answer: D

Explanation:

Correct Answer
JSON threat protection
****
Fact: Technically, there are no restrictions on what policy can be applied in what layer. Any policy can be applied on any layer API. However, context should also be considered properly before blindly applying the policies on APIs.
That is why, this question asked for a policy that would LEAST likely be applied to a Process API. From the given options:
>> All policies except "JSON threat protection" can be applied without hesitation to the APIs in Process tier.
>> JSON threat protection policy ideally fits for experience APIs to prevent suspicious JSON payload coming from external API clients. This covers more of a security aspect by trying to avoid possibly malicious and harmful JSON payloads from external clients calling experience APIs.
As external API clients are NEVER allowed to call Process APIs directly and also these kind of malicious and harmful JSON payloads are always stopped at experience API layer only using this policy, it is LEAST LIKELY that this same policy is again applied on Process Layer API.

NEW QUESTION 39
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

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.

NEW QUESTION 40
Mule applications that implement a number of REST APIs are deployed to their own subnet that is inaccessible from outside the organization.
External business-partners need to access these APIs, which are only allowed to be invoked from a separate subnet dedicated to partners - called Partner-subnet.
This subnet is accessible from the public internet, which allows these external partners to reach it.
Anypoint Platform and Mule runtimes are already deployed in Partner-subnet. These Mule runtimes can already access the APIs.
What is the most resource-efficient solution to comply with these requirements, while having the least impact on other applications that are currently using the APIs?

A. Implement (or generate) an API proxy Mule application for each of the APIs, then deploy the API proxies to the Mule runtimes
B. Redeploy the API implementations to the same servers running the Mule runtimes
C. Add an additional endpoint to each API for partner-enablement consumption
D. Duplicate the APIs as Mule applications, then deploy them to the Mule runtimes

Answer: A

NEW QUESTION 45
What CANNOT be effectively enforced using an API policy in Anypoint Platform?

A. Guarding against Denial of Service attacks
B. Maintaining tamper-proof credentials between APIs
C. Logging HTTP requests and responses
D. Backend system overloading

Answer: A

Explanation:
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 49
The responses to some HTTP requests can be cached depending on the HTTP verb used in the request. According to the HTTP specification, for what HTTP verbs is this safe to do?

A. PUT, POST, DELETE
B. GET, HEAD, POST
C. GET, PUT, OPTIONS
D. GET, OPTIONS, HEAD

Answer: D

Explanation:

Correct Answer
GET, OPTIONS, HEAD

APIs use HTTP-based protocols: cached HTTP responses from previous HTTP requests may potentially be returned if the same HTTP request is seen again.

Safe HTTP methods are ones that do not alter the state of the underlying resource. That is, the *HTTP responses to requests using safe HTTP methods may be cached*.

The HTTP standard requires the following HTTP methods on any resource to be safe:

- GET
- HEAD
- OPTIONS

Safety must be honored by REST APIs (but not by non-REST APIs like SOAP APIs): It is the *responsibility of every API implementation* to implement GET, HEAD or OPTIONS methods such that they never change the state of a resource.

http://restcookbook.com/HTTP%20Methods/idempotency/

NEW QUESTION 53
What is a key requirement when using an external Identity Provider for Client Management in Anypoint Platform?

A. Single sign-on is required to sign in to Anypoint Platform
B. The application network must include System APIs that interact with the Identity Provider
C. To invoke OAuth 2.0-protected APIs managed by Anypoint Platform, API clients must submit access tokens issued by that same Identity Provider
D. APIs managed by Anypoint Platform must be protected by SAML 2.0 policies

Answer: C

Explanation:
https://www.folkstalk.com/2019/11/mulesoft-integration-and-platform.html
Correct Answer
To invoke OAuth 2.0-protected APIs managed by Anypoint Platform, API clients must submit access tokens issued by that same Identity Provider
****
>> It is NOT necessary that single sign-on is required to sign in to Anypoint Platform because we are using an external Identity Provider for Client Management
>> It is NOT necessary that all APIs managed by Anypoint Platform must be protected by SAML 2.0 policies because we are using an external Identity Provider for Client Management
>> Not TRUE that the application network must include System APIs that interact with the Identity Provider because we are using an external Identity Provider for Client Management
Only TRUE statement in the given options is - "To invoke OAuth 2.0-protected APIs managed by Anypoint Platform, API clients must submit access tokens issued by that same Identity Provider"
References:
https://docs.mulesoft.com/api-manager/2.x/external-oauth-2.0-token-validation-policy https://blogs.mulesoft.com/dev/api-dev/api-security-ways-to-authenticate-and-authorize/

NEW QUESTION 54
What should be ensured before sharing an API through a public Anypoint Exchange portal?

A. The visibility level of the API instances of that API that need to be publicly accessible should be set to public visibility
B. The users needing access to the API should be added to the appropriate role in Anypoint Platform
C. The API should be functional with at least an initial implementation deployed and accessible for users to interact with
D. The API should be secured using one of the supported authentication/authorization mechanisms to ensure that data is not compromised

Answer: A

Explanation:

Correct Answer
The visibility level of the API instances of that API that need to be publicly accessible should be set to public visibility.
****

NEW QUESTION 58
Which layer in the API-led connectivity focuses on unlocking key systems, legacy systems, data sources etc and exposes the functionality?

A. Experience Layer
B. Process Layer
C. System Layer

Answer: C

Explanation:
Correct Answer
System Layer

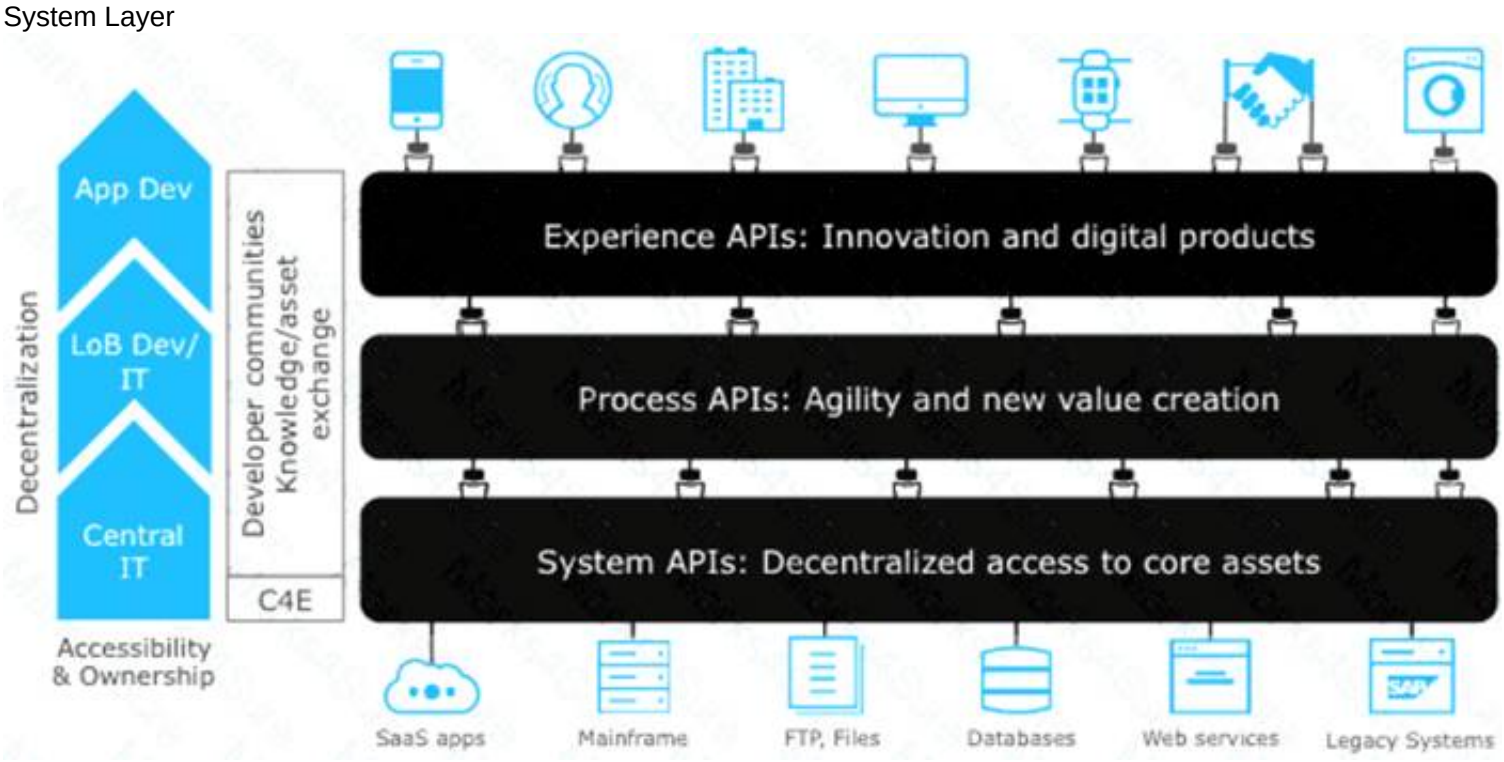

The APIs used in an API-led approach to connectivity fall into three categories:
System APIs – these usually access the core systems of record and provide a means of insulating the user from the complexity or any changes to the underlying systems. Once built, many users, can access data without any need 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 63
An organization has several APIs that accept JSON data over HTTP POST. The APIs are all publicly available and are associated with several mobile applications and web applications.
The organization does NOT want to use any authentication or compliance policies for these APIs, but at the same time, is worried that some bad actor could send payloads that could somehow compromise the applications or servers running the API implementations.
What out-of-the-box Anypoint Platform policy can address exposure to this threat?

A. Shut out bad actors by using HTTPS mutual authentication for all API invocations
B. Apply an IP blacklist policy to all APIs; the blacklist will Include all bad actors
C. Apply a Header injection and removal policy that detects the malicious data before it is used
D. Apply a JSON threat protection policy to all APIs to detect potential threat vectors

Answer: D

Explanation:

Correct Answer
Apply a JSON threat protection policy to all APIs to detect potential threat vectors
****
>> Usually, if the APIs are designed and developed for specific consumers (known consumers/customers) then we would IP Whitelist the same to ensure that traffic only comes from them.
>> However, as this scenario states that the APIs are publicly available and being used by so many mobile and web applications, it is NOT possible to identify and blacklist all possible bad actors.
>> So, JSON threat protection policy is the best chance to prevent any bad JSON payloads from such bad actors.

NEW QUESTION 65
Select the correct Owner-Layer combinations from below options

A. * 1. App Developers owns and focuses on Experience Layer APIs* 2. Central IT owns and focuses on Process Layer APIs* 3. LOB IT owns and focuses on System Layer APIs
B. * 1. Central IT owns and focuses on Experience Layer APIs* 2. LOB IT owns and focuses on Process Layer APIs* 3. App Developers owns and focuses on System Layer APIs
C. * 1. App Developers owns and focuses on Experience Layer APIs* 2. LOB IT owns and focuses on Process Layer APIs* 3. Central IT owns and focuses on System Layer APIs

Answer: C

Explanation:

Correct Answer
* 1. App Developers owns and focuses on Experience Layer APIs
* 2. LOB IT owns and focuses on Process Layer APIs
* 3. Central IT owns and focuses on System Layer APIs
References:
https://blogs.mulesoft.com/biz/api/experience-api-ownership/ https://blogs.mulesoft.com/biz/api/process-api-ownership/ https://blogs.mulesoft.com/biz/api/system-api-ownership/

NEW QUESTION 70
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

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

NEW QUESTION 74
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 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.

NEW QUESTION 75
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

Answer: C

NEW QUESTION 80
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?

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

Answer: B

Explanation:

Correct Answers
* 1. Each modern API must be treated like a product and designed for a particular target audience (for instance mobile app developers)
****
Bottom of Form Top of Form

NEW QUESTION 81
A company requires Mule applications deployed to CloudHub to be isolated between non-production and production environments. This is so Mule applications deployed to non-production environments can only access backend systems running in their customer-hosted non-production environment, and so Mule applications deployed to production environments can only access backend systems running in their customer-hosted production environment. How does MuleSoft recommend modifying Mule applications,
configuring environments, or changing infrastructure to support this type of per-environment isolation between Mule applications and backend systems?

A. Modify properties of Mule applications deployed to the production Anypoint Platform environments to prevent access from non-production Mule applications
B. Configure firewall rules in the infrastructure inside each customer-hosted environment so that only IP addresses from the corresponding Anypoint Platform environments are allowed to communicate with corresponding backend systems
C. Create non-production and production environments in different Anypoint Platform business groups
D. Create separate Anypoint VPCs for non-production and production environments, then configure connections to the backend systems in the corresponding customer-hosted environments

Answer: D

Explanation:
Correct Answer
Create separate Anypoint VPCs for non-production and production environments, then configure connections to the backend systems in the corresponding customer-hosted environments.
****
>> Creating different Business Groups does NOT make any difference w.r.t accessing the non-prod and prod customer-hosted environments. Still they will be accessing from both Business Groups unless process network restrictions are put in place.
>> We need to modify or couple the Mule Application Implementations with the environment. In fact, we should never implements application coupled with environments by binding them in the properties. Only basic things like endpoint URL etc should be bundled in properties but not environment level access restrictions.
>> IP addresses on CloudHub are dynamic until unless a special static addresses are assigned. So it is not possible to setup firewall rules in customer-hosted infrastrcture. More over, even if static IP addresses are assigned, there could be 100s of applications running on cloudhub and setting up rules for all of them would be a hectic task, non-maintainable and definitely got a good practice.
>> Thbeest practice recommended
by MulesoftIn(fact any cloud provider), is to have your Anypoint VPCs
seperated for Prod and Non-Prod and perform the VPC peering or VPN tunneling for these Anypoint VPCs to respective Prod and Non-Prod customer-hosted environment networks.

NEW QUESTION 86
The implementation of a Process API must change.
What is a valid approach that minimizes the impact of this change on API clients?

A. Update the RAML definition of the current Process API and notify API client developers by sending them links to the updated RAML definition
B. Postpone changes until API consumers acknowledge they are ready to migrate to a new Process API or API version
C. Implement required changes to the Process API implementation so that whenever possible, the Process API's RAML definition remains unchanged
D. Implement the Process API changes in a new API implementation, and have the old API implementation return an HTTP status code 301 - Moved Permanently to inform API clients they should be calling the new API implementation

Answer: C

Explanation:
Correct Answer
Implement required changes to the Process API implementation so that, whenever possible, the Process API's RAML definition remains unchanged.
********** Key requirement in the question is:
>> Approach that minimizes the impact of this change on API clients Based on above:
>> Updating the RAML definition would possibly impact the API clients if the changes require any thing mandatory from client side. So, one should try to avoid doing that until really necessary.
>> Implementing the changes as a completely different API and then redirectly the clients with 3xx status code is really upsetting design and heavily impacts the

API clients.
>> Organisations and IT cannot simply postpone the changes required until all API consumers acknowledge they are ready to migrate to a new Process API or API version. This is unrealistic and not possible.
The best way to handle the changes always is to implement required changes to the API implementations so that, whenever possible, the API's RAML definition remains unchanged.

NEW QUESTION 90
A retail company with thousands of stores has an API to receive data about purchases and insert it into a single database. Each individual store sends a batch of purchase data to the API about every 30 minutes. The API implementation uses a database bulk insert command to submit all the purchase data to a database using a custom JDBC driver provided by a data analytics solution provider. The API implementation is deployed to a single CloudHub worker. The JDBC driver processes the data into a set of several temporary disk files on the CloudHub worker, and then the data is sent to an analytics engine using a proprietary protocol. This process usually takes less than a few minutes. Sometimes a request fails. In this case, the logs show a message from the JDBC driver indicating an out-of-file-space message. When the request is resubmitted, it is successful. What is the best way to try to resolve this throughput issue?

A. se a CloudHub autoscaling policy to add CloudHub workers
B. Use a CloudHub autoscaling policy to increase the size of the CloudHub worker
C. Increase the size of the CloudHub worker(s)
D. Increase the number of CloudHub workers

Answer: D

Explanation:
Correct Answer
Increase the size of the CloudHub worker(s)
****
The key details that we can take out from the given scenario are:
>> API implementation uses a database bulk insert command to submit all the purchase data to a database
>> JDBC driver processes the data into a set of several temporary disk files on the CloudHub worker
>> Sometimes a request fails and the logs show a message indicating an out-of-file-space message Based on above details:
>> Both auto-scaling options does NOT help because we cannot set auto-scaling rules based on error messages. Auto-scaling rules are kicked-off based on CPU/Memory usages and not due to some given error or disk space issues.
>> Increasing the number of CloudHub workers also does NOT help here because the reason for the failure is not due to performance aspects w.r.t CPU or Memory. It is due to disk-space.
>> Moreover, the API is doing bulk insert to submit the received batch data. Which means, all data is handled by ONE worker only at a time. So, the disk space issue should be tackled on "per worker" basis. Having multiple workers does not help as the batch may still fail on any worker when disk is out of space on that particular worker.
Therefore, the right way to deal this issue and resolve this is to increase the vCore size of the worker so that a new worker with more disk space will be provisioned.

NEW QUESTION 91
What Anypoint Platform Capabilities listed below fall under APIs and API Invocations/Consumers category? Select TWO.

A. API Operations and Management
B. API Runtime Execution and Hosting
C. API Consumer Engagement
D. API Design and Development

Answer: D

Explanation:
Correct Answers: API Operations and Management and API Consumer Engagement
****
>> API Design and Development
-
Anypoint Studio, Anypoint Design Center, Anypoint Connectors
>> API Runtime Execution and Hosting
-
Mule Runtimes, CloudHub, Runtime Services
>> API Operations and Management
-
Anypoint API Manager, Anypoint Exchange
>> API Consumer Management
-
API Contracts, Public Portals, Anypoint Exchange, API Notebooks
Bottom of Form Top of Form

NEW QUESTION 96
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

Answer: B

Explanation:

Correct Answer
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.

NEW QUESTION 97
An organization has implemented a Customer Address API to retrieve customer address information. This API has been deployed to multiple environments and has been configured to enforce client IDs everywhere.
A developer is writing a client application to allow a user to update their address. The developer has found the Customer Address API in Anypoint Exchange and wants to use it in their client application.
What step of gaining access to the API can be performed automatically by Anypoint Platform?

A. Approve the client application request for the chosen SLA tier
B. Request access to the appropriate API Instances deployed to multiple environments using the client application's credentials
C. Modify the client application to call the API using the client application's credentials
D. Create a new application in Anypoint Exchange for requesting access to the API

Answer: A

Explanation:

Correct Answer
Approve the client application request for the chosen SLA tier
****
>> Only approving the client application request for the chosen SLA tier can be automated
>> Rest of the provided options are not valid

NEW QUESTION 100
A system API has a guaranteed SLA of 100 ms per request. The system API is deployed to a primary environment as well as to a disaster recovery (DR) environment, with different DNS names in each environment. An upstream process API invokes the system API and the main goal of this process API is to respond to client requests in the least possible time. In what order should the system APIs be invoked, and what changes should be made in order to speed up the response time for requests from the process API?

A. In parallel, invoke the system API deployed to the primary environment and the system API deployed to the DR environment, and ONLY use the first response
B. In parallel, invoke the system API deployed to the primary environment and the system API deployed to the DR environment using a scatter-gather configured with a timeout, and then merge the responses
C. Invoke the system API deployed to the primary environment, and if it fails, invoke the system API deployed to the DR environment
D. Invoke ONLY the system API deployed to the primary environment, and add timeout and retry logic to avoid intermittent failures

Answer: A

Explanation:
Correct Answer
In parallel, invoke the system API deployed to the primary environment and the system API deployed to the DR environment, and ONLY use the first response.
****
>> The API requirement in the given scenario is to respond in least possible time.
>> The option that is suggesting to first try the API in primary environment and then fallback to API in DR environment would result in successful response but NOT in least possible time. So, this is NOT a right choice of implementation for given requirement.
>> Another option that is suggesting to ONLY invoke API in primary environment and to add timeout and retries may also result in successful response upon retries but NOT in least possible time. So, this is also NOT a right choice of implementation for given requirement.
>> One more option that is suggesting to invoke API in primary environment and API in DR environment in parallel using Scatter-Gather would result in wrong API response as it would return merged results and moreover, Scatter-Gather does things in parallel which is true but still completes its scope only on finishing all routes inside it. So again, NOT a right choice of implementation for given requirement
The Correct choice is to invoke the API in primary environment and the API in DR environment parallelly, and using ONLY the first response received from one of them.

NEW QUESTION 103
Refer to the exhibit.

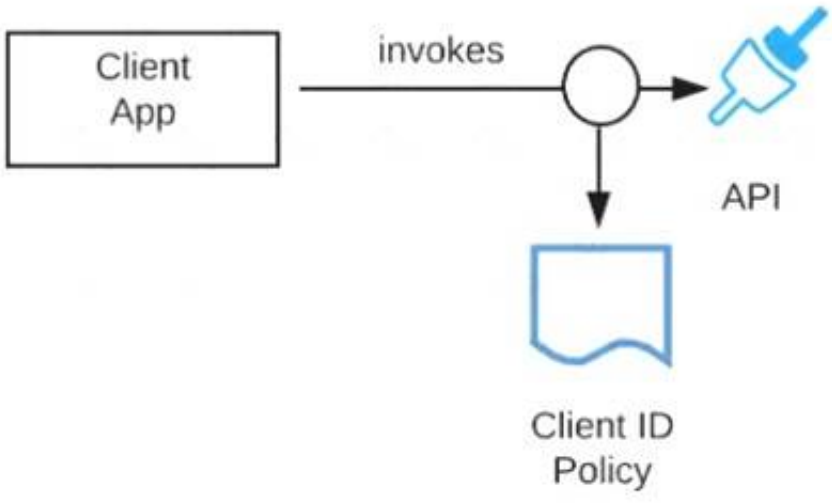

A developer is building a client application to invoke an API deployed to the STAGING environment that is governed by a client ID enforcement policy.
What is required to successfully invoke the API?

A. The client ID and secret for the Anypoint Platform account owning the API in the STAGING environment
B. The client ID and secret for the Anypoint Platform account's STAGING environment
C. The client ID and secret obtained from Anypoint Exchange for the API instance in the STAGING environment
D. A valid OAuth token obtained from Anypoint Platform and its associated client ID and secret

Answer: C

Explanation:
Correct Answer
The client ID and secret obtained from Anypoint Exchange for the API instance in the STAGING environment
****
>> We CANNOT use the client ID and secret of Anypoint Platform account or any individual environments for accessing the APIs
>> As the type of policy that is enforced on the API in question is "Client ID Enforcment Policy", OAuth token based access won't work.
Right way to access the API is to use the client ID and secret obtained from Anypoint Exchange for the API instance in a particular environment we want to work on.
References:
Managing API instance Contracts on API Manager https://docs.mulesoft.com/api-manager/1.x/request-access-to-api-task https://docs.mulesoft.com/exchange/to-request-access https://docs.mulesoft.com/api-manager/2.x/policy-mule3-client-id-based-policies

NEW QUESTION 106
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

Answer: C

Explanation:
Correct Answer
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.

NEW QUESTION 107
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

MCPA-Level-1 Practice Exam Features:

* MCPA-Level-1 Questions and Answers Updated Frequently

* MCPA-Level-1 Practice Questions Verified by Expert Senior Certified Staff

* MCPA-Level-1 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* MCPA-Level-1 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year