Salesforce

Exam Questions MuleSoft-Platform-Architect-I

Salesforce Certified MuleSoft Platform Architect 1 Exam (SP24)

NEW 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

Answer: C

Explanation:

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

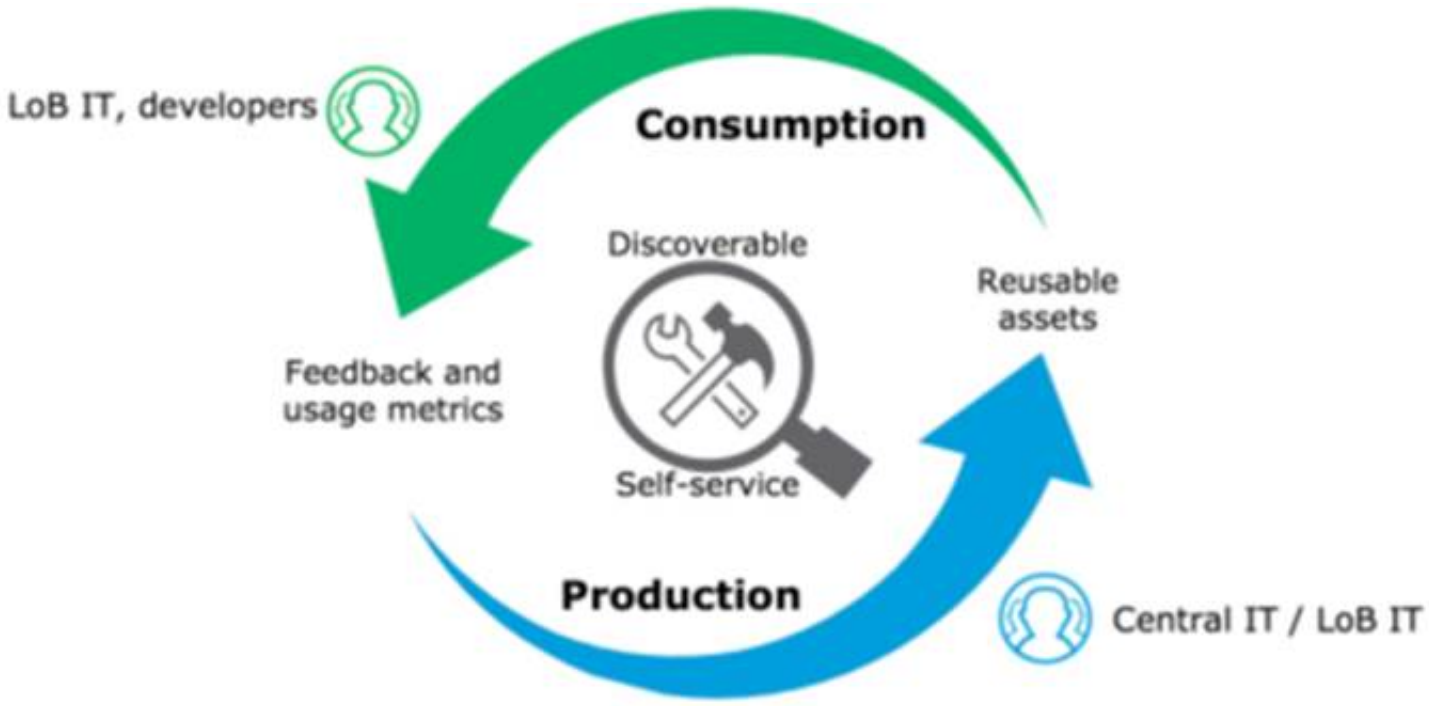

NEW QUESTION 2
An IT Security Compliance Auditor is assessing which nonfunctional requirements (NFRs) are already being implemented to meet security measures.
* The Web API has Rate-Limiting SLA
* Basic Authentication - LDAP
* JSON Threat Protection
* TP Allowlist policies applied Which two NFRs-are enforced?

A. The API invocations are coming from a known subnet range
B. Username/password supported to validate login credentials
C. Sensitive data is masked to prevent compromising critical information
D. The API is protected against XML invocation attacks
E. Performance expectations are to be allowed up to 1,000 requests per second

Answer: AB

NEW QUESTION 3
A company uses a hybrid Anypoint Platform deployment model that combines the EU control plane with customer-hosted Mule runtimes. After successfully testing a Mule API implementation in the Staging environment, the Mule API implementation is set with environment-specific properties and must be promoted to the Production environment. What is a way that MuleSoft recommends to configure the Mule API implementation and automate its promotion to the Production environment?

A. Bundle properties files for each environment into the Mule API implementation's deployable archive, then promote the Mule API implementation to the Production environment using Anypoint CLI or the Anypoint Platform REST APIsB.
B. Modify the Mule API implementation's properties in the API Manager Properties tab, then promote the Mule API implementation to the Production environment using API Manager
C. Modify the Mule API implementation's properties in Anypoint Exchange, then promote the Mule API implementation to the Production environment using Runtime Manager
D. Use an API policy to change properties in the Mule API implementation deployed to the Staging environment and another API policy to deploy the Mule API implementation to the Production environment

Answer: A

NEW QUESTION 4
A manufacturing company has deployed an API implementation to CloudHub and has not configured it to be automatically restarted by CloudHub when the worker is not responding.
Which statement is true when no API Client invokes that API implementation?

A. No alert on the API invocations and APT implementation can be raised
B. Alerts on the APT invocation and API implementation can be raised
C. No alert on the API invocations is raised but alerts on the API implementation can be raised
D. Alerts on the API invocations are raised but no alerts on the API implementation can be raised

Answer: C

NEW QUESTION 5
A company deploys Mule applications with default configurations through Runtime Manager to customer-hosted Mule runtimes. Each Mule application is an API implementation that exposes RESTful interfaces to API clients. The Mule runtimes are managed by the MuleSoft-hosted control plane. The payload is never used by any Logger components.
When an API client sends an HTTP request to a customer-hosted Mule application, which metadata or data (payload) is pushed to the MuleSoft-hosted control plane?

A. Only the data
B. No data
C. The data and metadata
D. Only the metadata

Answer: D

NEW QUESTION 6
What is most likely NOT a characteristic of an integration test for a REST API implementation?

A. The test needs all source and/or target systems configured and accessible
B. The test runs immediately after the Mule application has been compiled and packaged
C. The test is triggered by an external HTTP request
D. The test prepares a known request payload and validates the response payload

Answer: B

NEW QUESTION 7
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

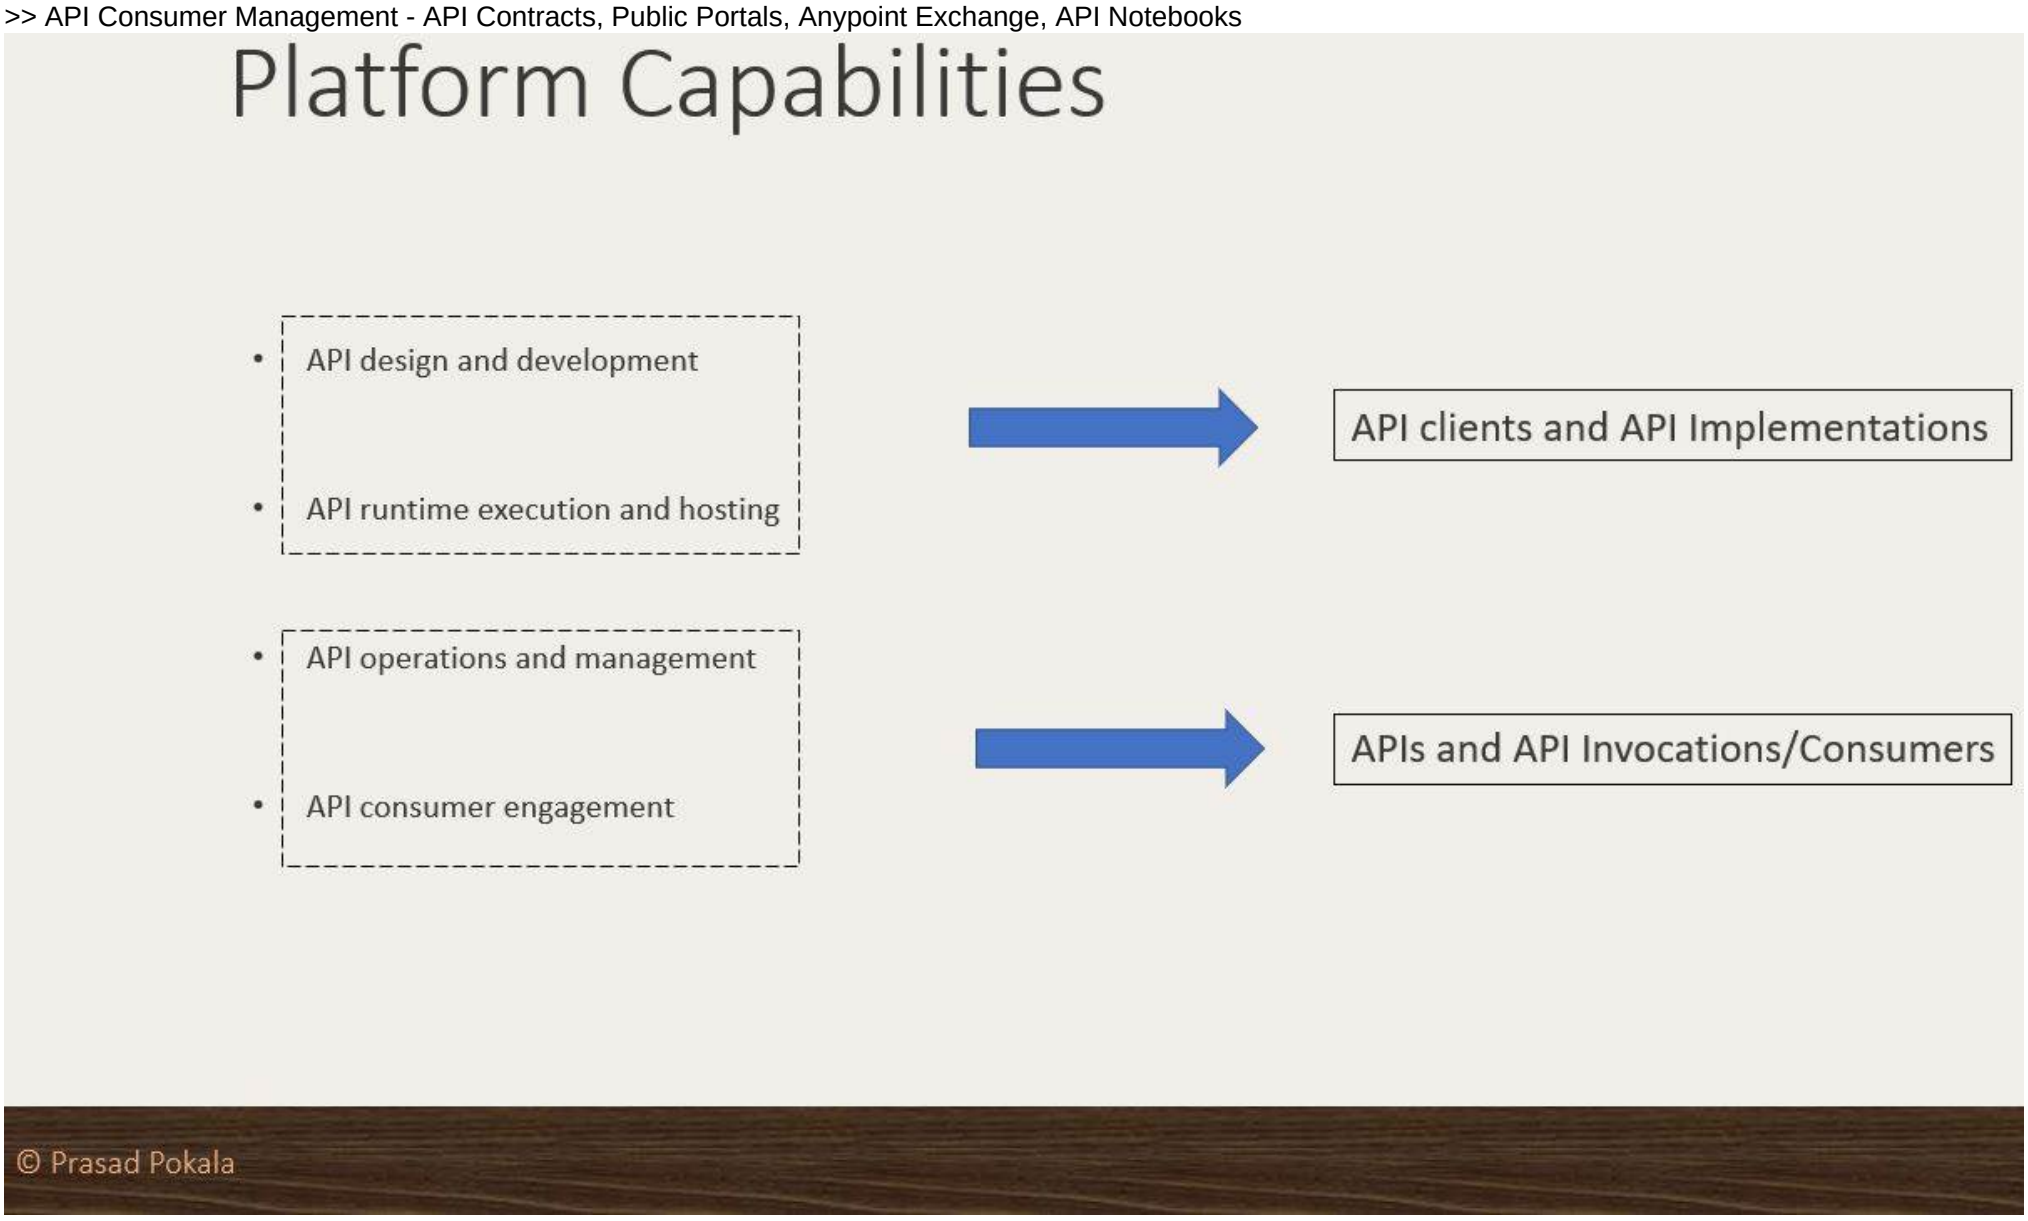

NEW QUESTION 8
An auto manufacturer has a mature CI/CD practice and wants to automate packaging and deployment of any Mule applications to various deployment targets, including CloudHub workers/replicas, customer-hosted Mule runtimes, and Anypoint Runtime Fabric.
Which MuleSoft-provided tool or component facilitates automating the packaging and deployment of Mule applications to various deployment targets as part of the company's CI/CD practice?

A. Anypoint Runtime Manager
B. Mule Maven plugin
C. Anypoint Platform CLI
D. Anypoint Platform REST APIs

Answer: B

NEW QUESTION 9
What Mule application deployment scenario requires using Anypoint Platform Private Cloud Edition or Anypoint Platform for Pivotal Cloud Foundry?

A. When it Is required to make ALL applications highly available across multiple data centers
B. When it is required that ALL APIs are private and NOT exposed to the public cloud
C. When regulatory requirements mandate on-premises processing of EVERY data item, including meta-data
D. When ALL backend systems in the application network are deployed in the organization's intranet

Answer: C

NEW QUESTION 10
What best describes the Fully Qualified Domain Names (FQDNs), also known as DNS entries, created when a Mule application is deployed to the CloudHub Shared Worker Cloud?

A. A fixed number of FQDNs are created, IRRESPECTIVE of the environment and VPC design
B. The FQDNs are determined by the application name chosen, IRRESPECTIVE of the region
C. The FQDNs are determined by the application name, but can be modified by an administrator after deployment
D. The FQDNs are determined by both the application name and the Anypoint Platform organization

Answer: B

NEW QUESTION 10
A customer wants to monitor and gain insights about the number of requests coming in a given time period as well as to measure key performance indicators (response times, CPU utilization, number of active APIs). Which tool provides these data insights?

A. Anypoint Monitoring
B. APT Manager
C. Runtime Alerts
D. Functional Monitoring

Answer: A

NEW QUESTION 11
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

NEW QUESTION 13
Which APIs can be used with DataGraph to create a unified schema?

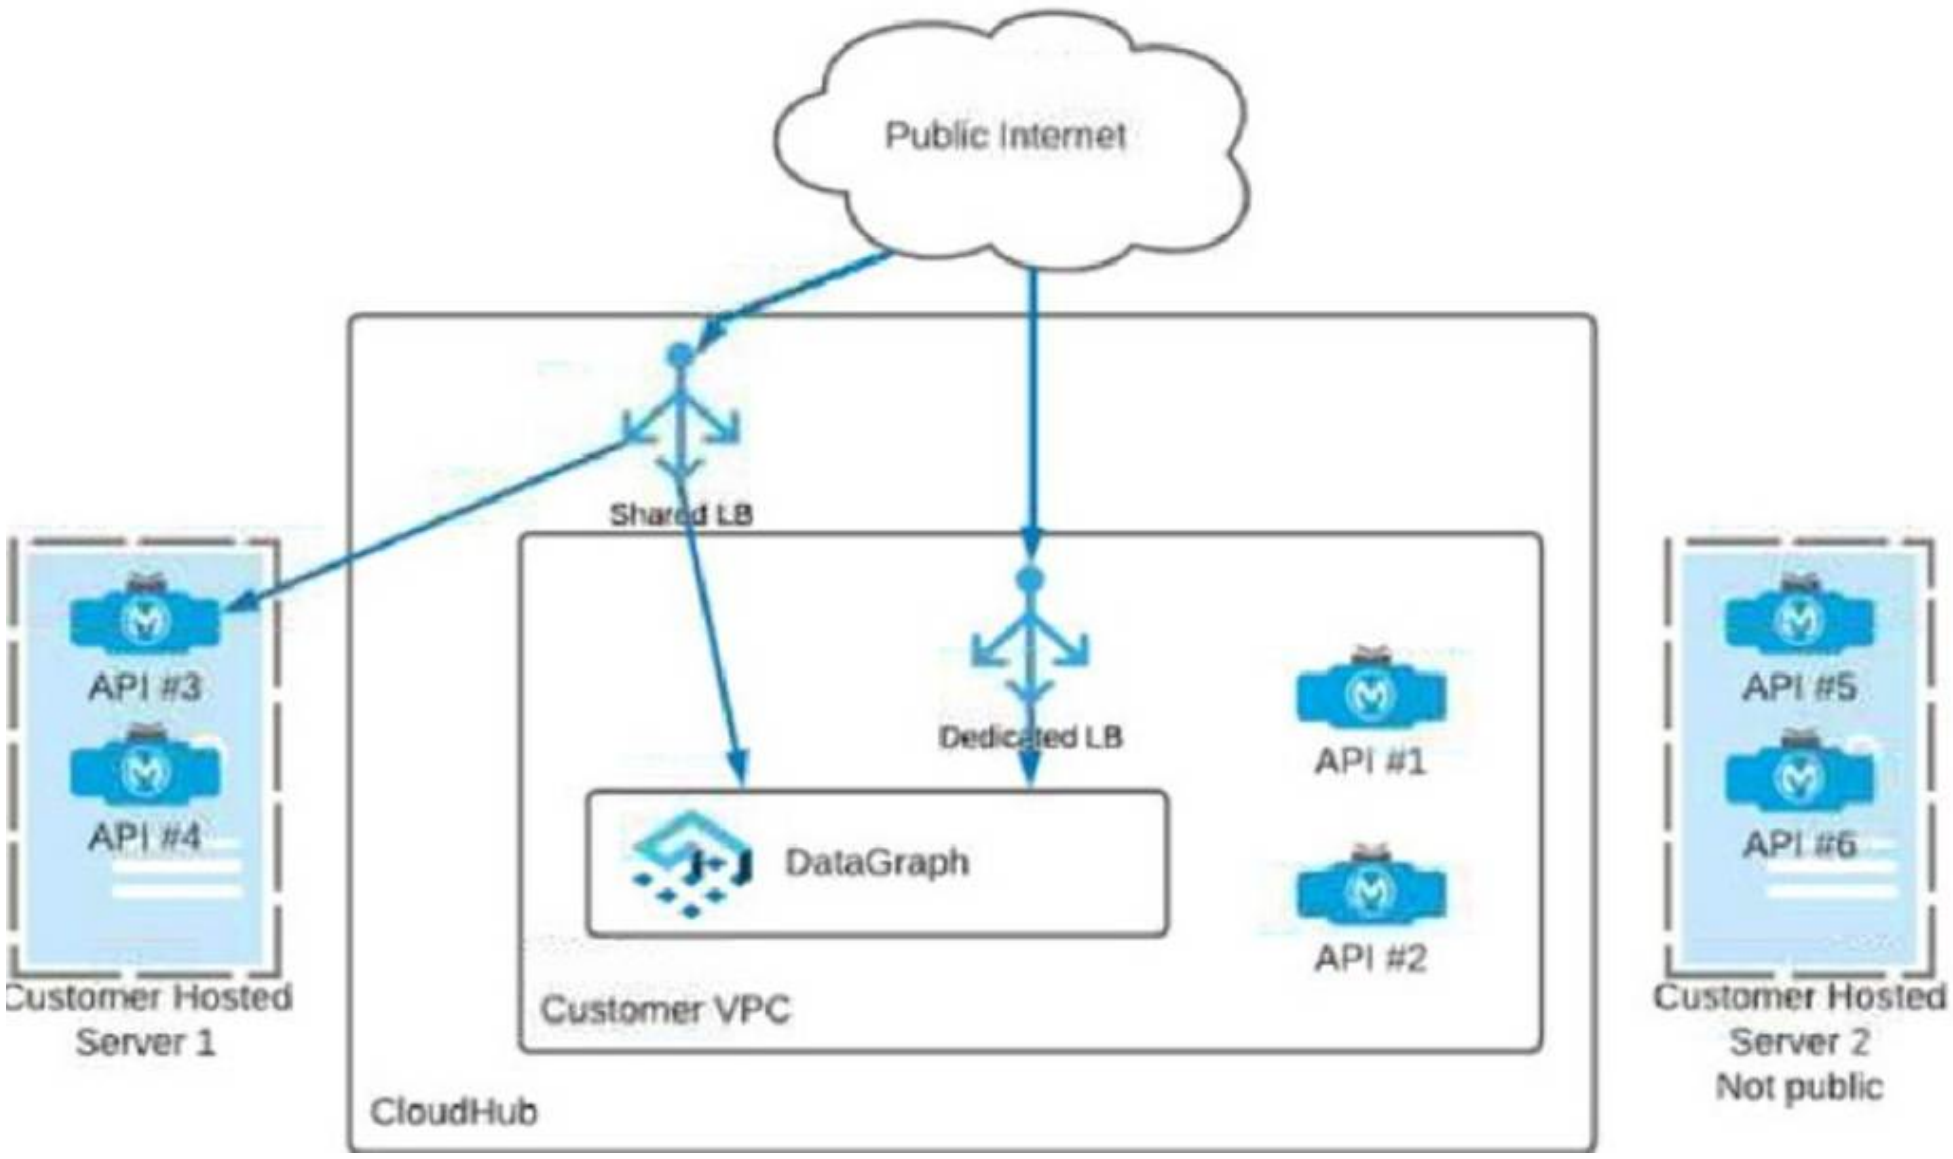

A. APIs 1, 3, 5
B. APIs 2, 4 ,6
C. APIs 1, 2, s5, 6
D. APIs 1, 2, 3, 4

Answer: D

NEW QUESTION 14
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

NEW QUESTION 16
Which of the following sequence is correct?

A. API Client implementes logic to call an API >> API Consumer requests access to API >> API Implementation routes the request to >> API
B. API Consumer requests access to API >> API Client implementes logic to call an API >> API routes the request to >> API Implementation
C. API Consumer implementes logic to call an API >> API Client requests access to API >> API Implementation routes the request to >> API
D. API Client implementes logic to call an API >> API Consumer requests access to API >> API routes the request to >> API Implementation

Answer: B

NEW QUESTION 20
A Rate Limiting policy is applied to an API implementation to protect the back-end system. Recently, there have been surges in demand that cause some API client POST requests to the API implementation to be rejected with policy-related errors, causing delays and complications to the API clients.
How should the API policies that are applied to the API implementation be changed to reduce the frequency of errors returned to API clients, while still protecting the back-end system?

A. Keep the Rate Limiting policy and add 9 Client ID Enforcement policy
B. Remove the Rate Limiting policy and add an HTTP Caching policy
C. Remove the Rate Limiting policy and add a Spike Control policy
D. Keep the Rate Limiting policy and add an SLA-based Spike Control policy

Answer: D

NEW QUESTION 21
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

Answer: C

NEW QUESTION 22
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

NEW QUESTION 25
Refer to the exhibit.

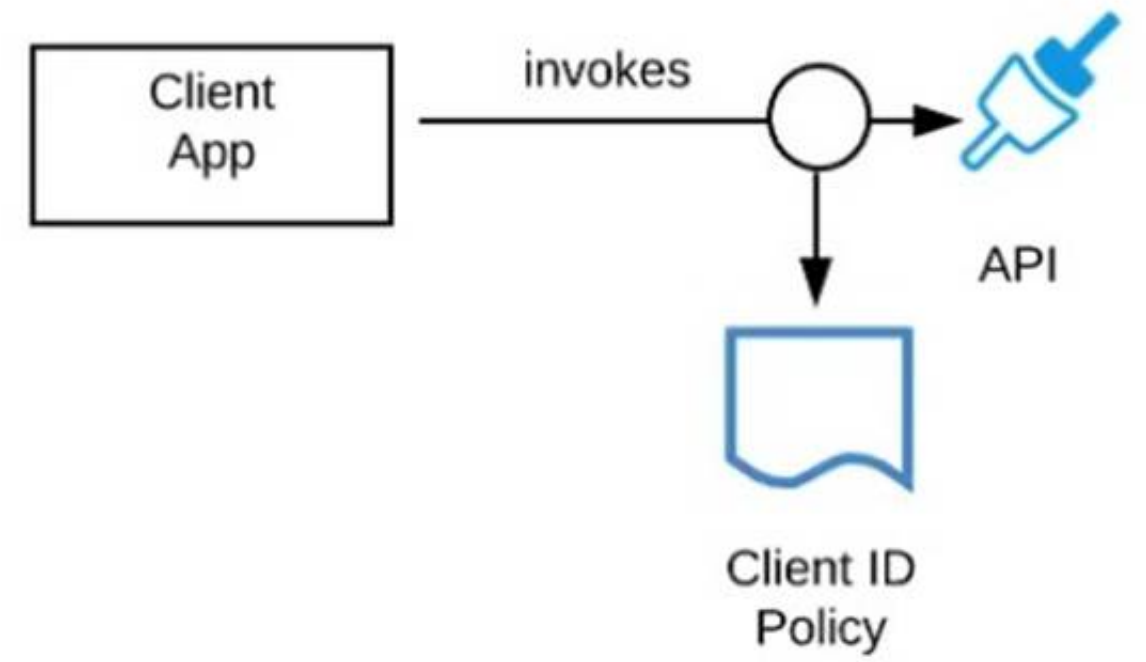

A developer is building a client application to invoke an API deployed to the STAGING environment that is governed by a client ID enforcement policy.
What is required to successfully invoke the API?

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

NEW QUESTION 26
A Mule 4 API has been deployed to CloudHub and a Basic Authentication - Simple policy has been applied to all API methods and resources. However, the API is still accessible by clients without using authentication. How is this possible?

A. The APE Router component is pointing to the incorrect Exchange version of the APT
B. The Autodiscovery element is not present, in the deployed Mule application
C. No?? for client applications have been created of this API
D. One of the application??s CloudHub workers restarted

Answer: B

NEW QUESTION 31
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

Answer: D

NEW QUESTION 32
4 Production environment is running on a dedicated Virtual Private Cloud (VPC) on CloudHub 1,0, and the security team guidelines clearly state no traffic on HTTP.
Which two options support these security guidelines? Choose 2 answers

A. Configure the HTTPS protocol in HTTP listener in the Mule application
B. Create a custom policy to apply to outgoing and incoming HTTP requests to control access to a configured API endpoint
C. Remove the entry from the VPC firewall rule

```
{
  "CIDR Block": "0.0.0.0/0", // (Anywhere)
  "Protocol": "TCP",
  "From port": 8081,
},
{
  "CIDR Block": "10.111.0.0/24", // (Local VPC)
  "Protocol": "TCP",
  "From port": 8091,
}
```

D. Configure the IP Blocklist policy to control access to a configured API endpoint from either a single IP address or a range of IP addresses.
E. Add the entry in the VPC firewall rule.

```
{
  "CIDR Block": "0.0.0.0/0", // (Anywhere)
  "Protocol": "TCP",
  "From port": 8081,
},
{
```

Answer: AC

NEW QUESTION 33
A company is building an application network using MuleSoft's recommendations for various API layers.
What is the main (default) role of a process API in an application network?

A. To secure and optimize the data synchronization processing of large data dumps between back-end systems
B. To manage and process the secure direct communication between a back-end system and an end-user client of mobile device in the application network
C. To automate parts of business processes by coordinating and orchestrating the invocation of other APIs in the application network
D. To secure, Manage, and process communication with specific types of end-user client applications or devices in the application network

Answer: C

NEW QUESTION 34
Several times a week, an API implementation shows several thousand requests per minute in an Anypoint Monitoring dashboard, Between these bursts, the dashboard shows between two and five requests per minute. The API implementation is running on Anypoint Runtime Fabric with two non-clustered replicas, reserved vCPU 1.0 and vCPU Limit 2.0.
An API consumer has complained about slow response time, and the dashboard shows the 99 percentile is greater than 120 seconds at the time of the complaint. It also shows greater than 90% CPU usage during these time periods.
In manual tests in the QA environment, the API consumer has consistently reproduced the slow response time and high CPU usage, and there were no other API requests at this time. In a brainstorming session, the engineering team has created several proposals to reduce the response time for requests.
Which proposal should be pursued first?

A. Increase the vCPU resources of the API implementation
B. Modify the API client to split the problematic request into smaller, less-demanding requests
C. Increase the number of replicas of the API implementation
D. Throttle the APT client to reduce the number of requests per minute

Answer: A

NEW QUESTION 36
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

NEW QUESTION 37
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

NEW QUESTION 42
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
Reference: https://docs.mulesoft.com/api-manager/2.x/using-policies

NEW QUESTION 45
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:

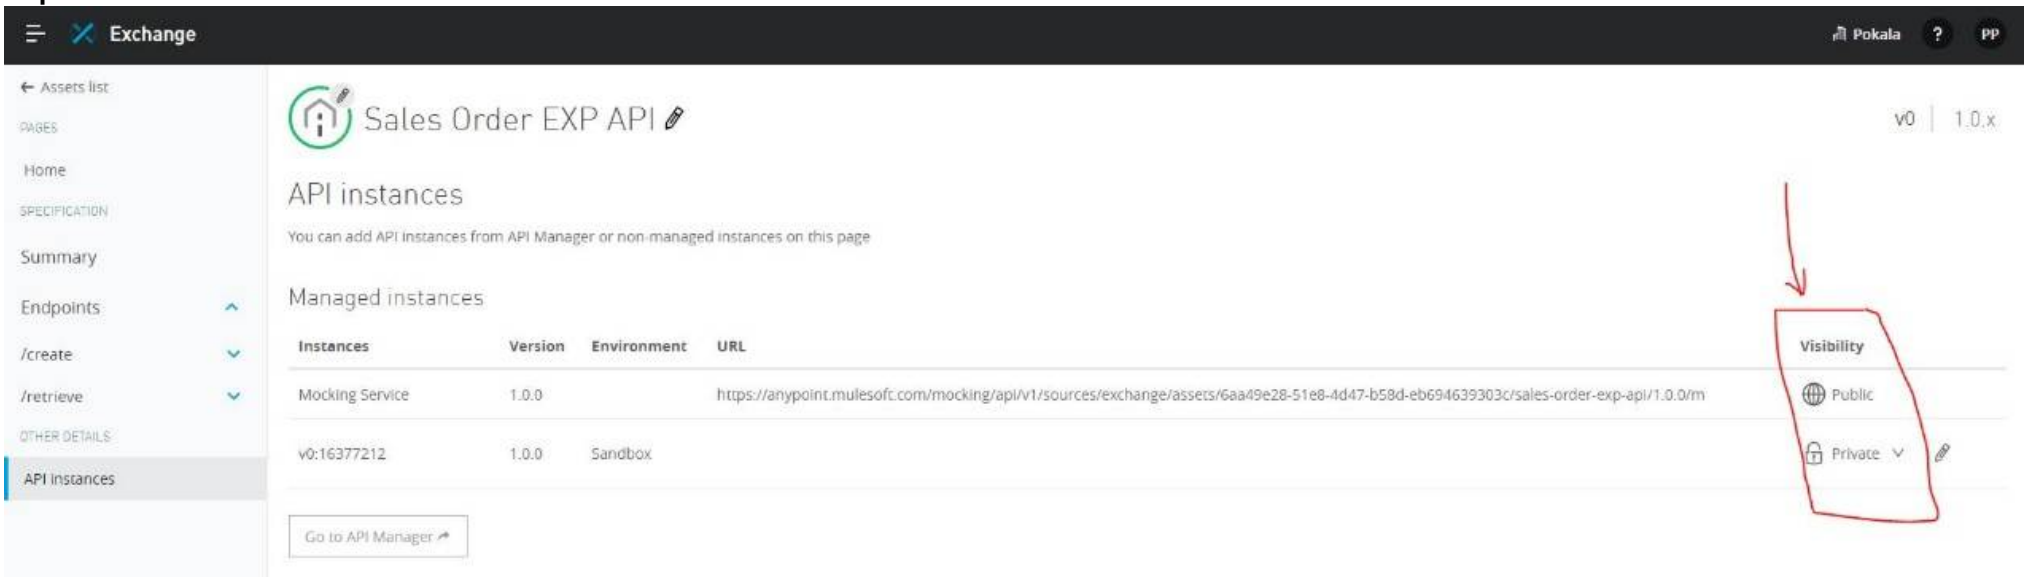

Correct Answer: The visibility level of the API instances of that API that need to be publicly accessible should be set to public visibility.
****
Reference: https://docs.mulesoft.com/exchange/to-share-api-asset-to-portal
https://docs.mulesoft.com/exchange/to-share-api-asset-to-portal

NEW QUESTION 50
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 benegotiated 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

Answer: B

NEW QUESTION 54
An organization wants to create a Center for Enablement (C4E). The IT director schedules a series of meetings with IT senior managers.
What should be on the agenda of the first meeting?

A. Define C4E objectives, mission statement, guiding principles, a
B. Explore API monetization options based on identified use cases through MuleSoft
C. A walk through of common-services best practices for logging, auditing, exception handling, caching, security via policy, and rate limiting/throttling via policy
D. Specify operating model for the MuleSoft Integrations division

Answer: A

NEW QUESTION 58
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

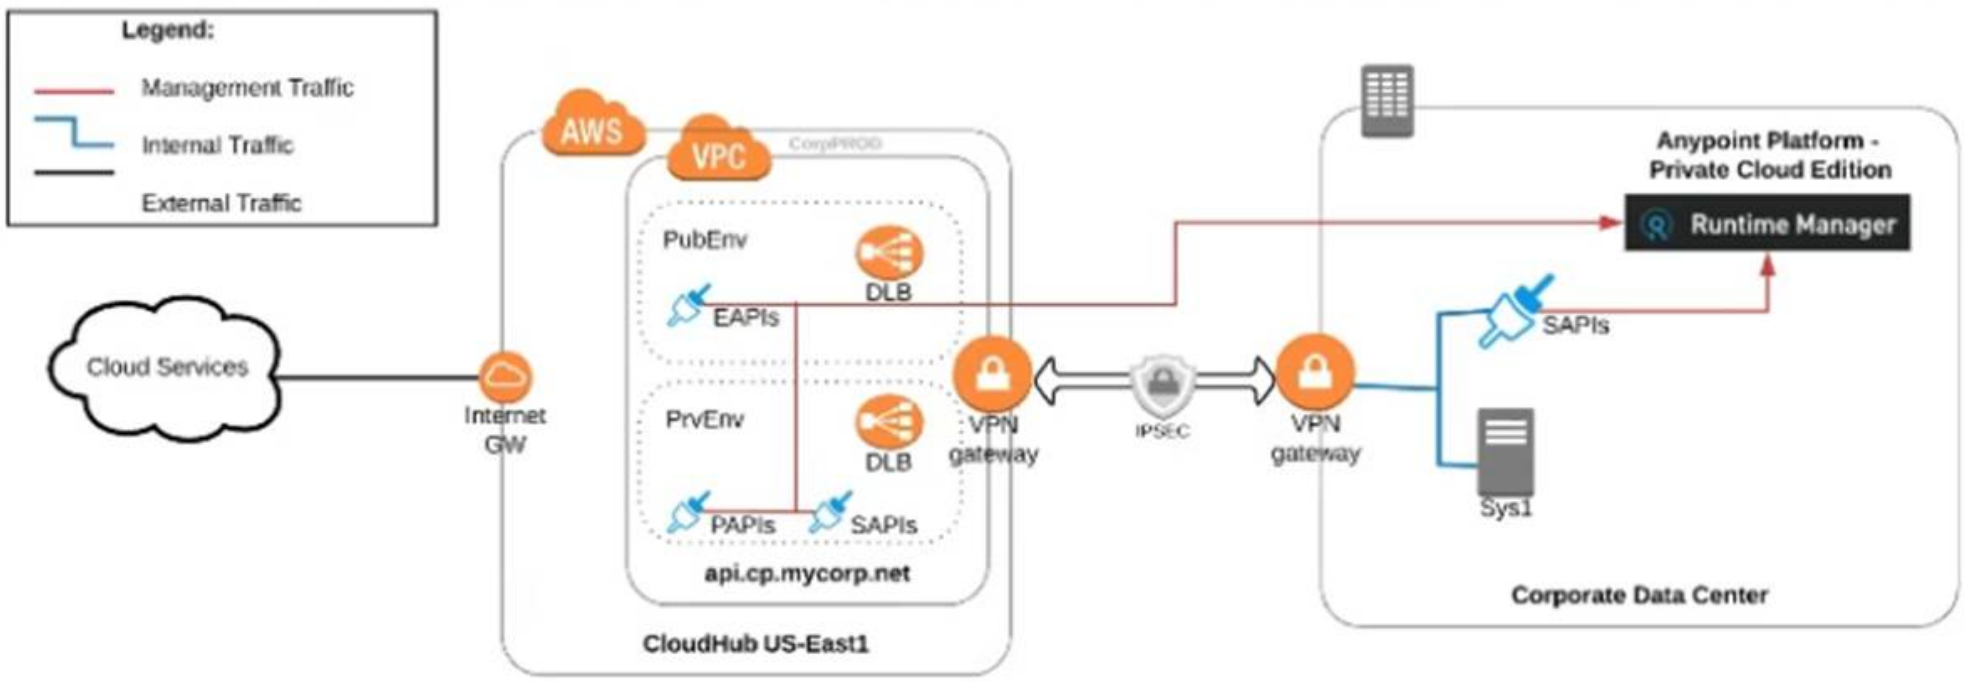

B) Use CloudHub-deployed Mule runtimes in the shared worker cloud managed by the MuleSoft-hosted Anypoint Platform control plane

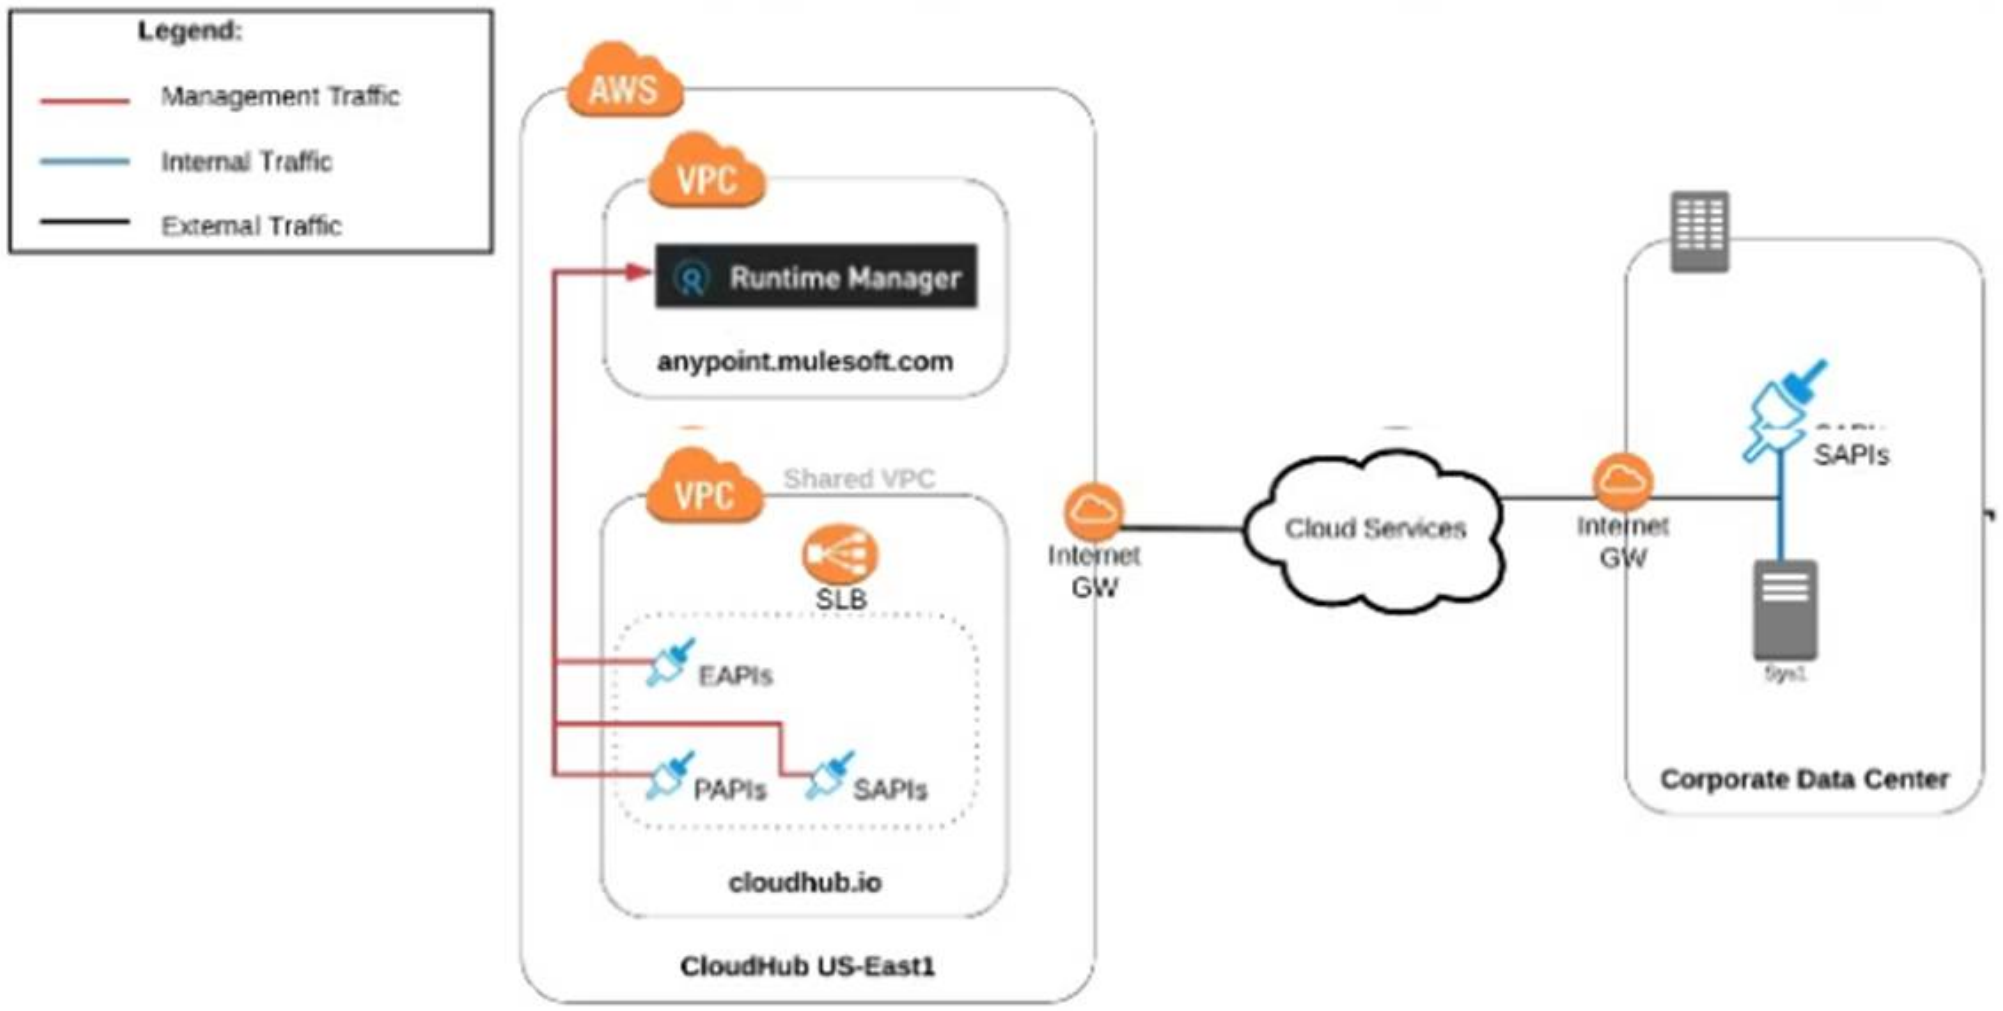

C) Use an on-premises installation of Mule runtimes that are completely isolated with NO external network access, managed by the Anypoint Platform Private Cloud Edition control plane

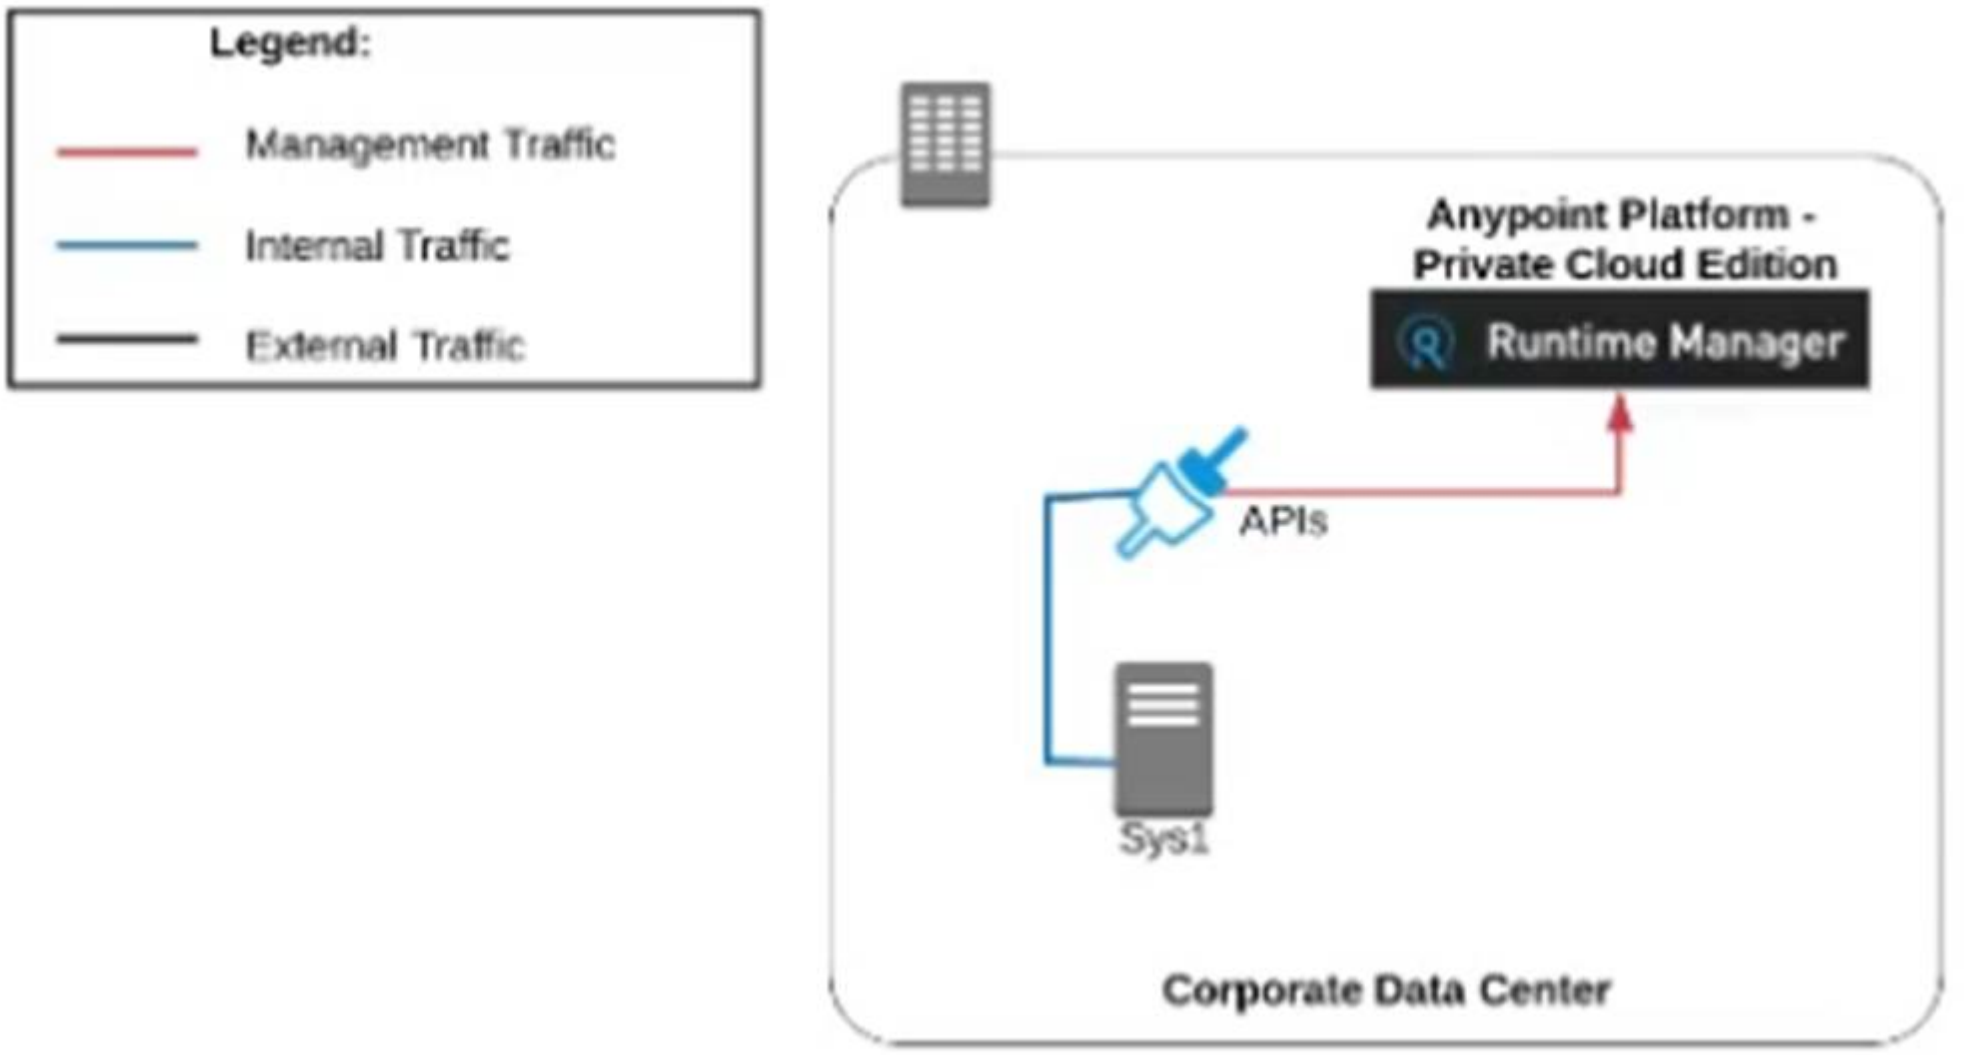

D) Use a combination of Cloud Hub-deployed and manually provisioned on-premises Mule runtimes managed by the MuleSoft-hosted Anypoint Platform control plane

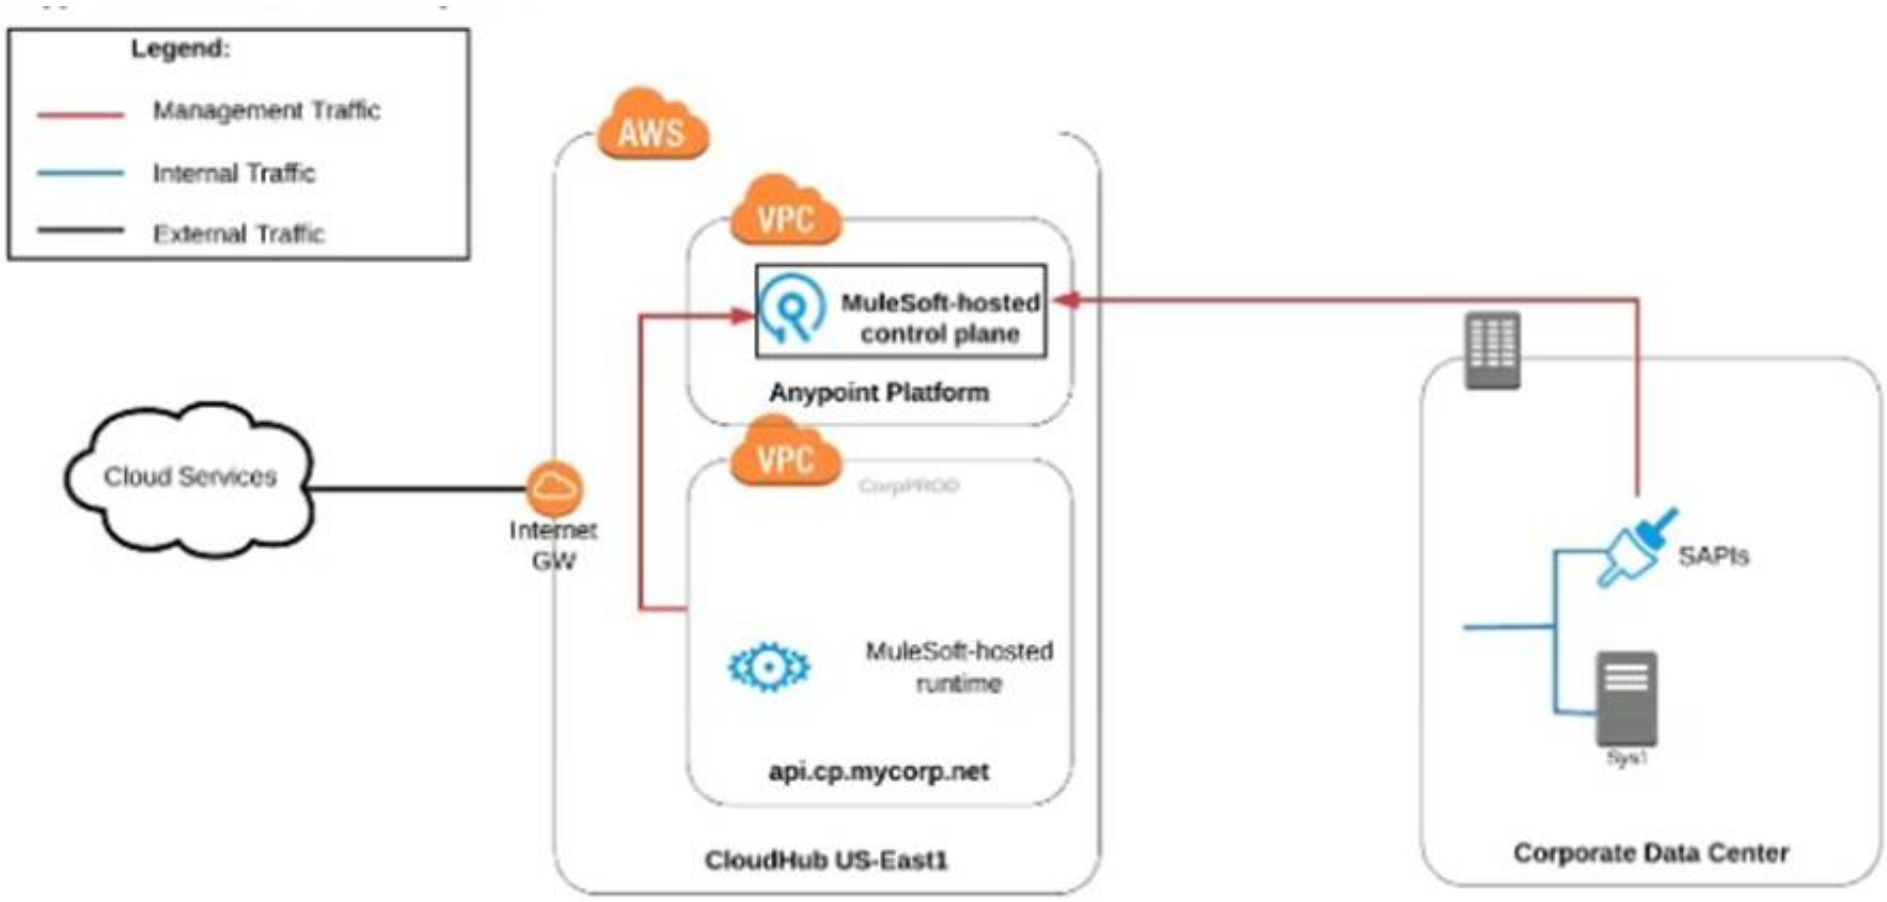

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

NEW QUESTION 61
What best explains the use of auto-discovery in API implementations?

A. It makes API Manager aware of API implementations and hence enables it to enforce policies
B. It enables Anypoint Studio to discover API definitions configured in Anypoint Platform
C. It enables Anypoint Exchange to discover assets and makes them available for reuse
D. It enables Anypoint Analytics to gain insight into the usage of APIs

Answer: A

NEW QUESTION 63
A customer has an ELA contract with MuleSoft. An API deployed to CloudHub is consistently experiencing performance issues. Based on the root cause analysis, it is determined that autoscaling needs to be applied. How can this be achieved?

A. Configure a policy so that when the number of HTTP requests reaches a certain threshold the number of workers/replicas increases (horizontal scaling)
B. Configure two separate policies: When CPU and memory reach certain threshold, increase the worker/replica type (vertical sealing) and the number of workers/replicas (horizontal sealing)
C. Configure a policy based on CPU usage so that CloudHub auto-adjusts the number of workers/replicas (horizontal scaling)
D. Configure a policy so that when the response time reaches a certain threshold the worker/replica type increases (vertical scaling)

Answer: C

NEW QUESTION 67
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

NEW QUESTION 68
Refer to the exhibits.

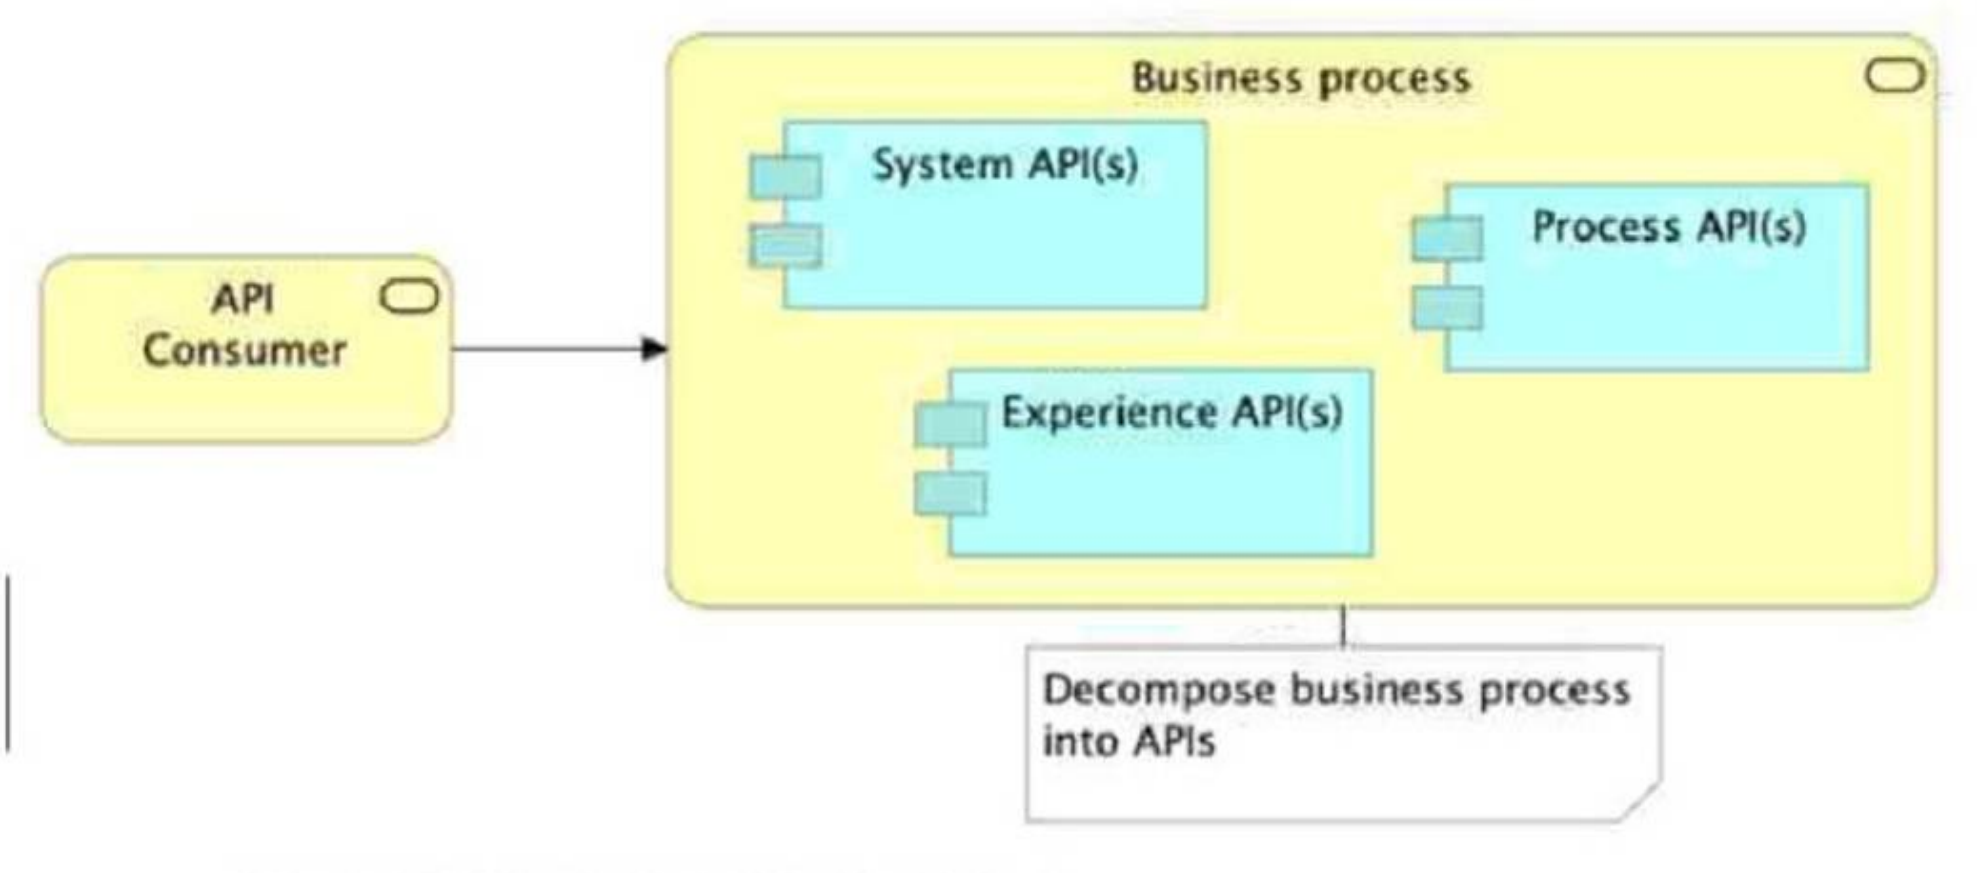

Which architectural constraint is compatible with the API-led connectivity architectural style?

A. Always use a tiered approach by creating exactly one API for each of the three layers (Experience, Process, and System)

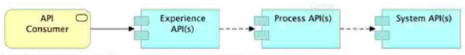

B. Use a Process API to-orchestrate calls to multiple System APIs but not to other Process APIs:
C. Allow System APIs to return data that is not currently required by the identified Process or Experience APIs

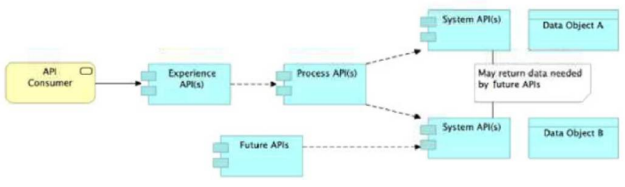

D. Handle customizations for the end-user application at the Process layer rather than at the Experience layer

Answer: B

NEW QUESTION 69
A Mule application exposes an HTTPS endpoint and is deployed to three CloudHub workers that do not use static IP addresses. The Mule application expects a high volume of client requests in short time periods. What is the most cost-effective infrastructure component that should be used to serve the high volume of client requests?

A. A customer-hosted load balancer
B. The CloudHub shared load balancer
C. An API proxy
D. Runtime Manager autoscaling

Answer: B

NEW QUESTION 74
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

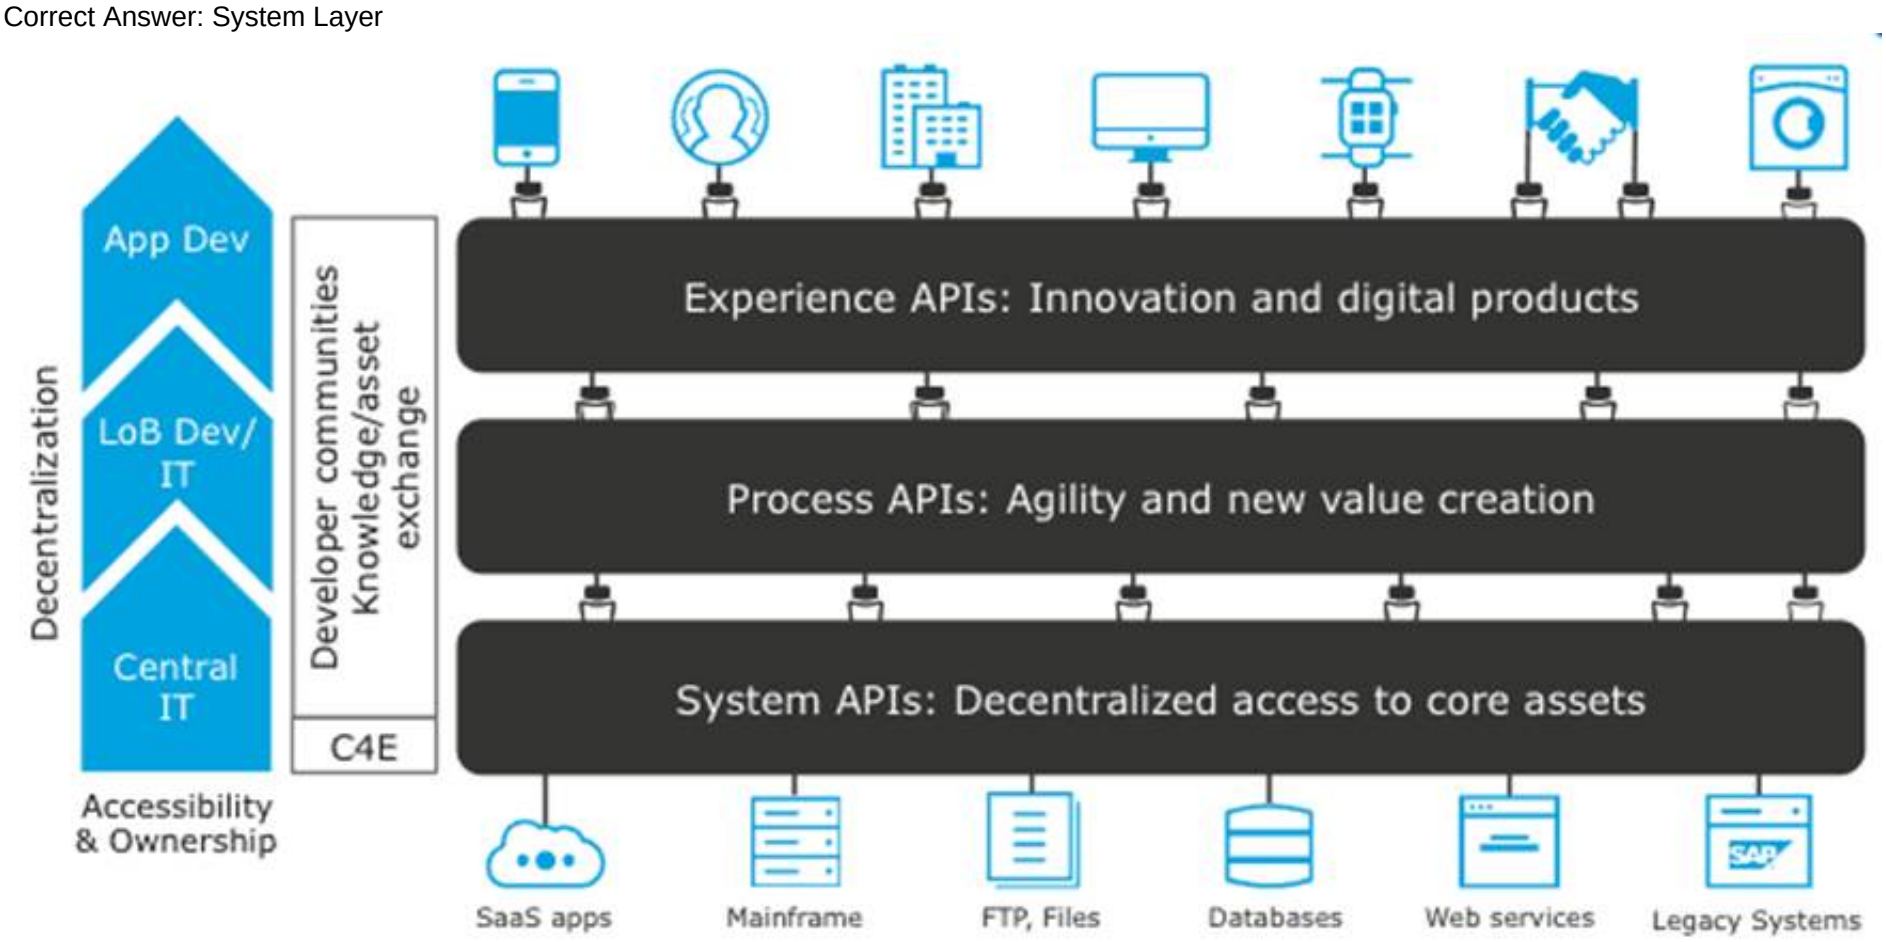

The APIs used in an API-led approach to connectivity fall into three categories:
System APIs – these usually access the core systems of record and provide a means of insulating the user from the complexity or any changes to the underlying systems. Once built, many users, can access data without any need 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 76
A company wants to move its Mule API implementations into production as quickly as possible. To protect access to all Mule application data and metadata, the company requires that all Mule applications be deployed to the company's customer-hosted infrastructure within the corporate firewall. What combination of runtime plane and control plane options meets these project lifecycle goals?

A. Manually provisioned customer-hosted runtime plane and customer-hosted control plane
B. MuleSoft-hosted runtime plane and customer-hosted control plane
C. Manually provisioned customer-hosted runtime plane and MuleSoft-hosted control plane
D. iPaaS provisioned customer-hosted runtime plane and MuleSoft-hosted control plane

Answer: A

NEW QUESTION 80
Refer to the exhibit.

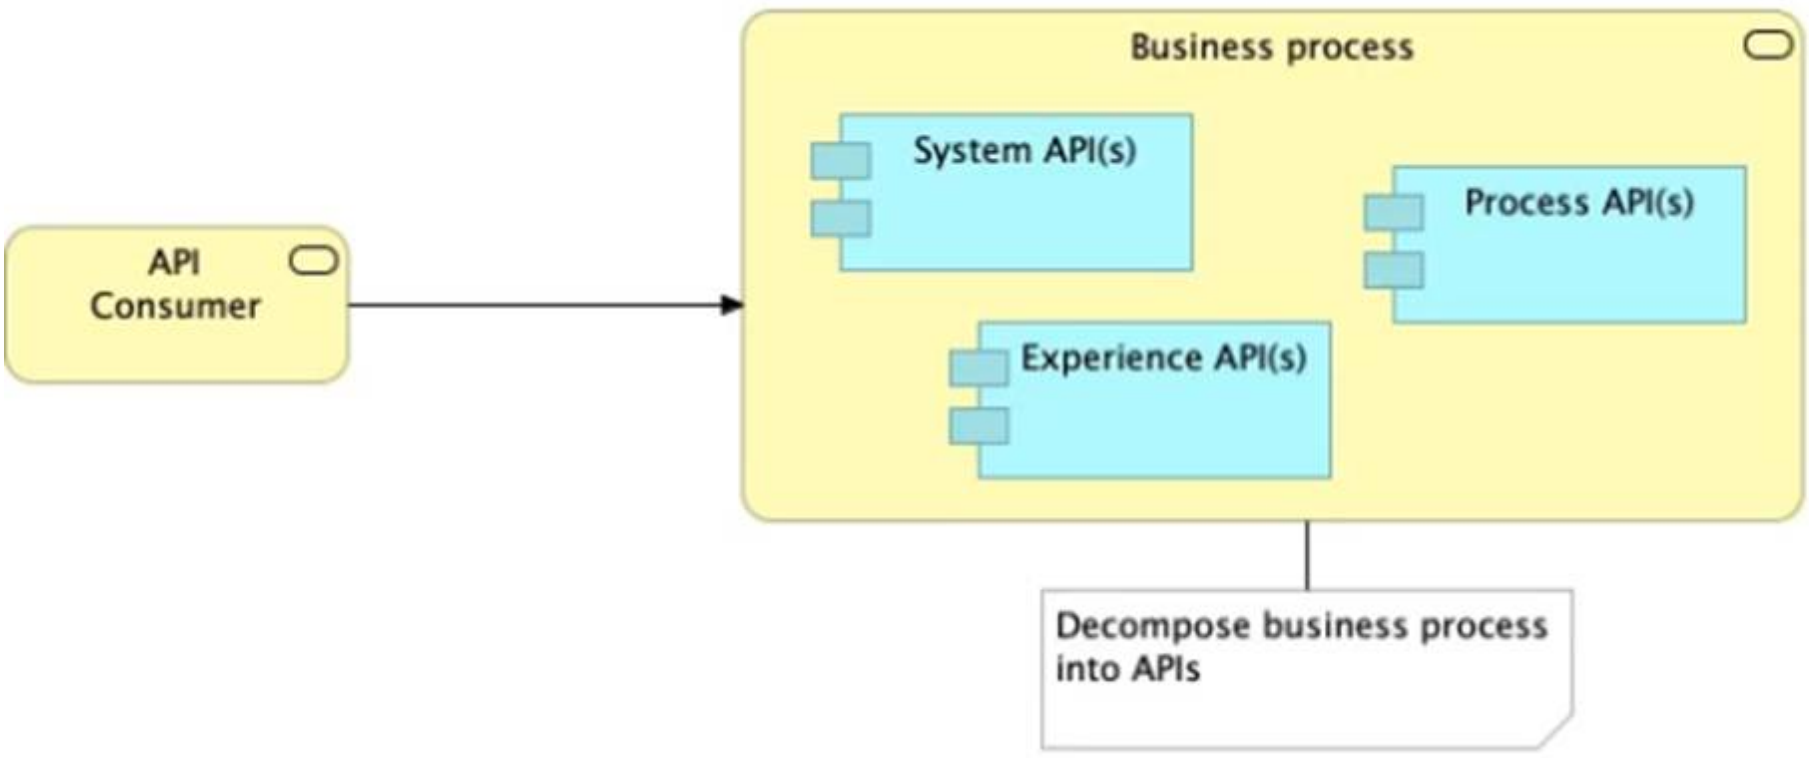

What is the best way to decompose one end-to-end business process into a collaboration of Experience, Process, and System APIs?
A) Handle customizations for the end-user application at the Process API level rather than the Experience API level

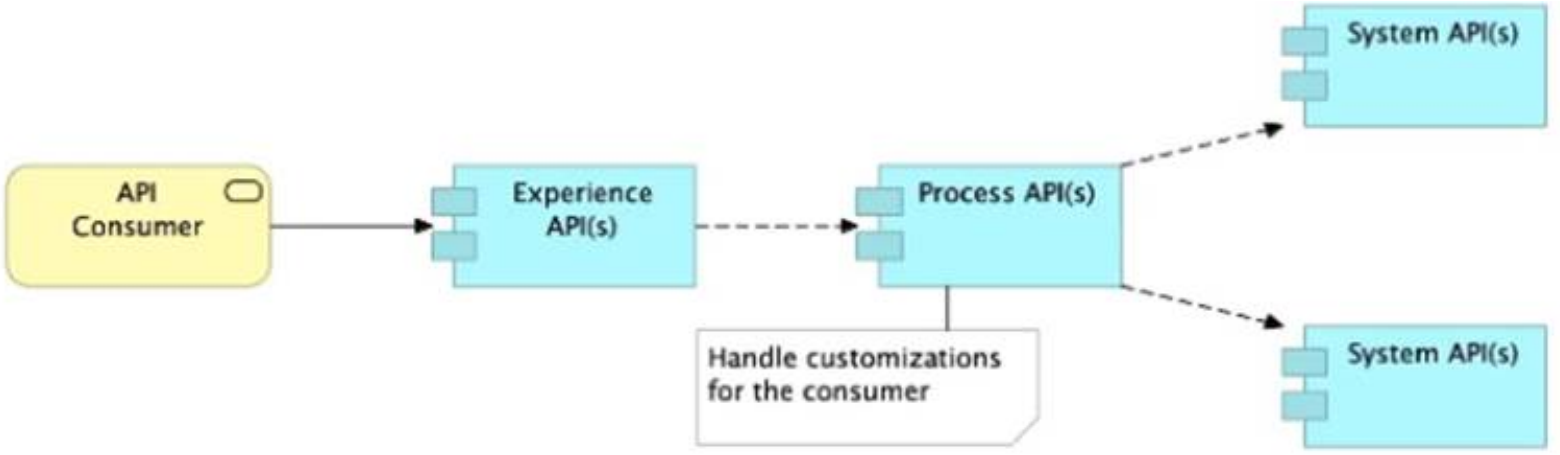

B) Allow System APIs to return data that is NOT currently required by the identified Process or Experience APIs

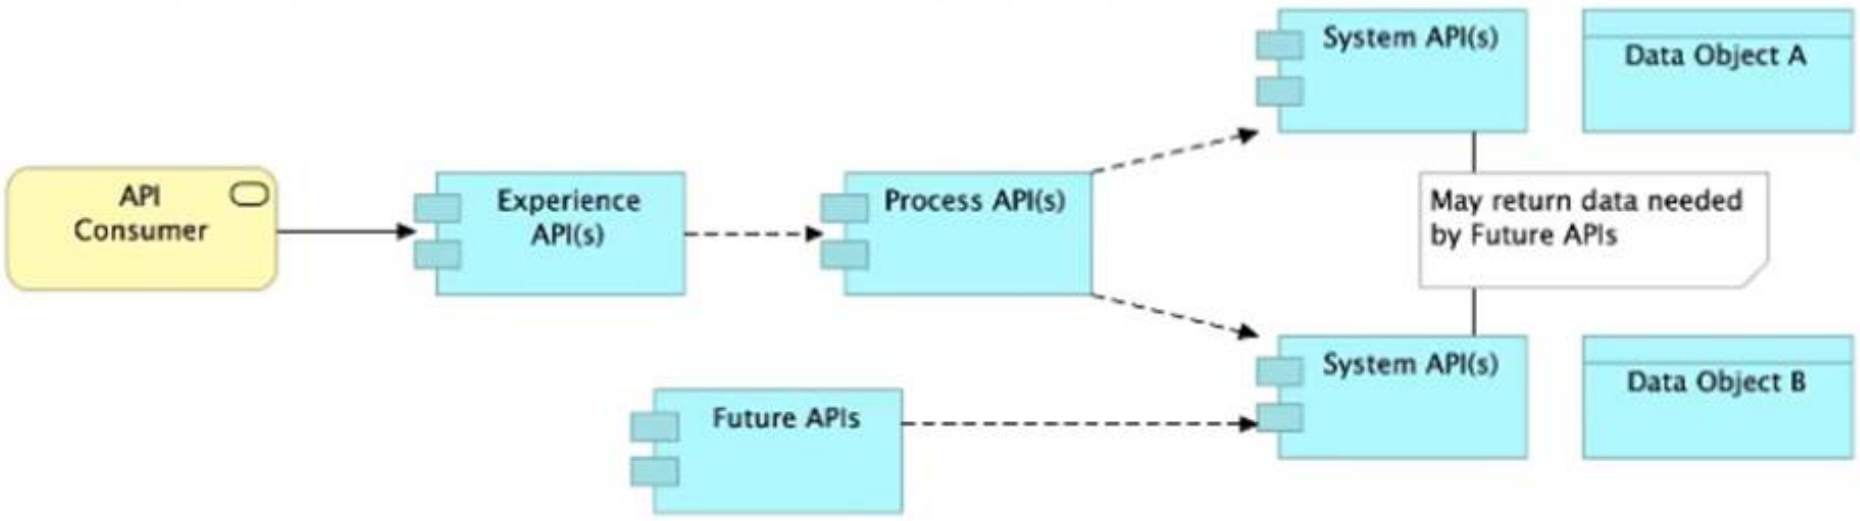

C) Always use a tiered approach by creating exactly one API for each of the 3 layers (Experience, Process and System APIs)

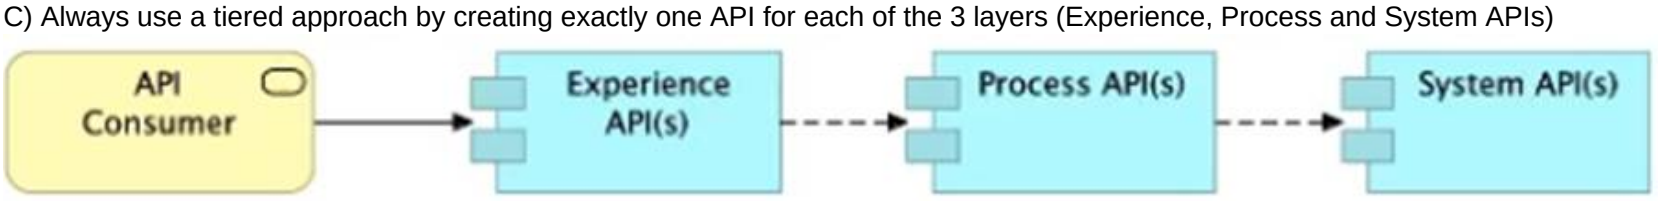

D) Use a Process API to orchestrate calls to multiple System APIs, but NOT to other Process APIs

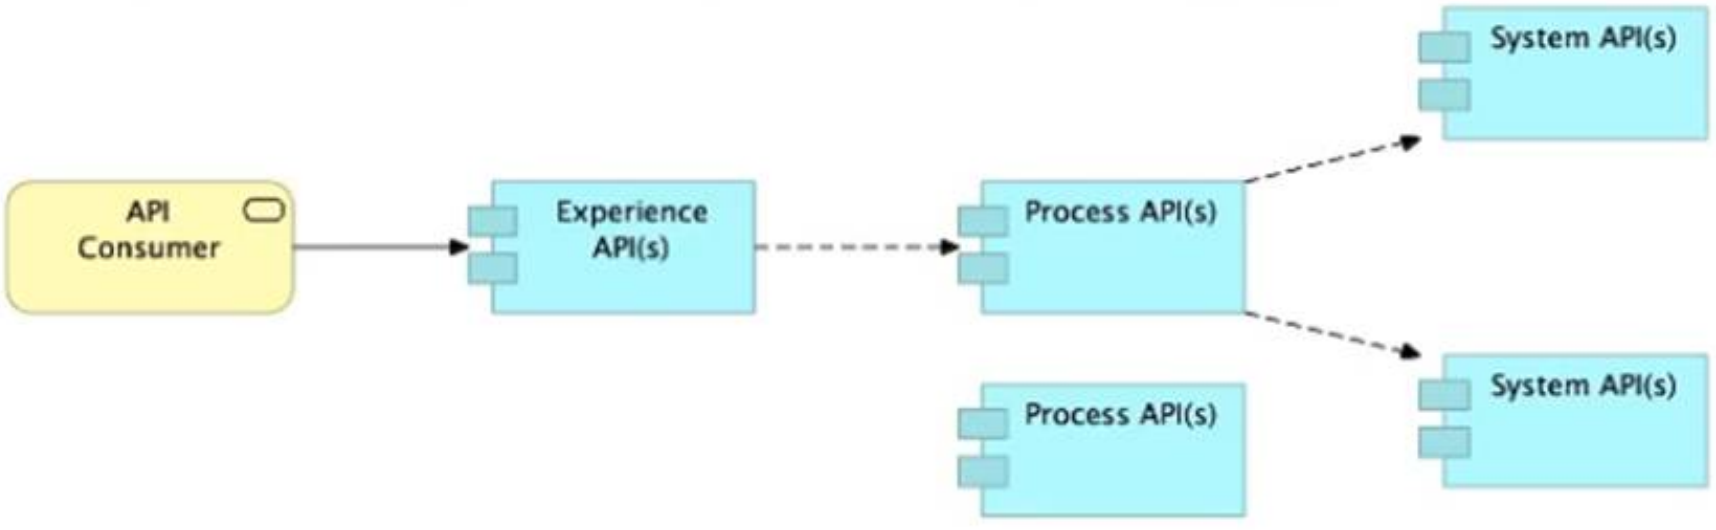

A. Option A
B. Option B

C. Option C
D. Option D

Answer: B

NEW QUESTION 84
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

NEW QUESTION 85
Which three tools automate the deployment of Mule applications? Choose 3 answers

A. Runtime Manager
B. Anypoint Platform CLI
C. Platform APIs
D. Anypoint Studio
E. Mule Mayen plugin
F. API Community Manager

Answer: ABC

NEW QUESTION 87
An operations team is analyzing the effort needed to set up monitoring of their application network. They are looking at which API invocation metrics can be used to identify and predict trouble without having to write custom scripts or install additional analytics software or tools.
Which type of metrics can satisfy this goal of directly identifying and predicting failures?

A. The number and types of API policy violations per day
B. The effectiveness of the application network based on the level of reuse
C. The number and types of past API invocations across the application network
D. The ROI from each APT invocation

Answer: A

NEW QUESTION 88
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 API invocations to an API implementation?

A. When the API invocations are sent directly to the internal DNS record of the API implementation
B. When the API invocations are not over-a- secure TLS/SSL communication channel
C. When the APL invecations originate from a geography different than the API
D. When the number of API invocations are below a threshold

Answer: D

NEW QUESTION 90
A system API is deployed to a primary environment as well as to a disaster recovery (DR) environment, with different DNS names in each environment. A process API is a client to the system API and is being rate limited by the system API, with different limits in each of the environments. The system API's DR environment provides only 20% of the rate limiting offered by the primary environment. What is the best API fault-tolerant invocation strategy to reduce overall errors in the process API, given these conditions and constraints?

A. Invoke the system API deployed to the primary environment; add timeout and retry logic to the process API to avoid intermittent failures; if it still fails, invoke the system API deployed to the DR environment
B. Invoke the system API deployed to the primary environment; add retry logic to the process API to handle intermittent failures by invoking the system API deployed to the DR environment
C. In parallel, invoke the system API deployed to the primary environment and the system API deployed to the DR environment; add timeout and retry logic to the process API to avoid intermittent failures; add logic to the process API to combine the results
D. Invoke the system API deployed to the primary environment; add timeout and retry logic to the process API to avoid intermittent failures; if it still fails, invoke a copy of the process API deployed to the DR environment

Answer: A

NEW QUESTION 94
Which statement is true about Spike Control policy and Rate Limiting policy?

A. All requests are rejected after the limit is reached in Rate Limiting policy, whereas the requests are queued in Spike Control policy after the limit is reached
B. In a clustered environment, the Rate Limiting.and Spike Control policies are applied to each node in the cluster
C. To protect Experience APIs by limiting resource consumption, Rate Limiting policy must be applied
D. In order to apply Rate Limiting and Spike Control policies, a contract to bind client application and API is needed for both

Answer: B

NEW QUESTION 97
A company deployed an API to a single worker/replica in the shared cloud in the U.S. West Region.
What happens when the Availability Zone experiences an outage?

A. CloudHub will auto-redeploy the APL in the U.
B. East Region
C. The APT will be unavailable until the availability comes back online, at which time the worker/replica will be auto-restarted
D. CloudHub will auto-redeploy the API in another Availability Zone in the U.
E. West Region
F. The Anypoint Platform admin is alerted when the AP] is experiencing an outage and needs the trigger the CI/CD pipeline to redeploy to the U
G. East Region

Answer: B

NEW QUESTION 102
What API policy would be LEAST LIKELY used when designing an Experience API that is intended to work with a consumer mobile phone or tablet application?

A. OAuth 2.0 access token enforcement
B. Client ID enforcement
C. JSON threat protection
D. IPwhitellst

Answer: D

NEW QUESTION 105
Once an API Implementation is ready and the API is registered on API Manager, who should request the access to the API on Anypoint Exchange?

A. None
B. Both
C. API Client
D. API Consumer

Answer: D

NEW QUESTION 110
When can CloudHub Object Store v2 be used?

A. To store an unlimited number of key-value pairs
B. To store payloads with an average size greater than 15MB
C. To store information in Mule 4 Object Store v1
D. To store key-value pairs with keys up to 300 characters

Answer: D

NEW QUESTION 114
An organization is deploying their new implementation of the OrderStatus System API to multiple workers in CloudHub. This API fronts the organization's on-premises Order Management System, which is accessed by the API implementation over an IPsec tunnel.
What type of error typically does NOT result in a service outage of the OrderStatus System API?

A. A CloudHub worker fails with an out-of-memory exception
B. API Manager has an extended outage during the initial deployment of the API implementation
C. The AWS region goes offline with a major network failure to the relevant AWS data centers
D. The Order Management System is Inaccessible due to a network outage in the organization's on-premises data center

Answer: A

NEW QUESTION 115
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:

* 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 119
To minimize operation costs, a customer wants to use a CloudHub 1.0 solution. The customer's requirements are:
* Separate resources with two Business groups
* High-availability (HA) for all APIs
* Route traffic via Dedicated load balancer (DLBs)
* Separate environments into production and non-production Which solution meets the customer's needs?

A. One production and one non-production Virtual Private Cloud (VPC). Use availability zones to differentiate between Business groups.Allocate maximum CIDR per VPCs to ensure HA across availability zones
B. One production and one non-production Virtual Private Cloud (VPC) per Business group.Minimize CIDR aligning with projected application total.Choose a MuleSoft CloudHub 1.0 region with multiple availability zone
C. Deploy multiple workers for HA,
D. One production and one non-production Virtual Private Cloud (VPC) per Business group.Minimize CIDR aligning with projected application tota
E. Divide availability zones during deployment of APIs for HA.
F. One production and one non-production Virtual Private Claud (VPC). Configure subnet to differentiate between business groups.Allocate maximum CIDR per VPCs to make it easier to add Child group
G. Span VPC to cover three availability zones.

Answer: B

NEW QUESTION 124
Refer to the exhibit.

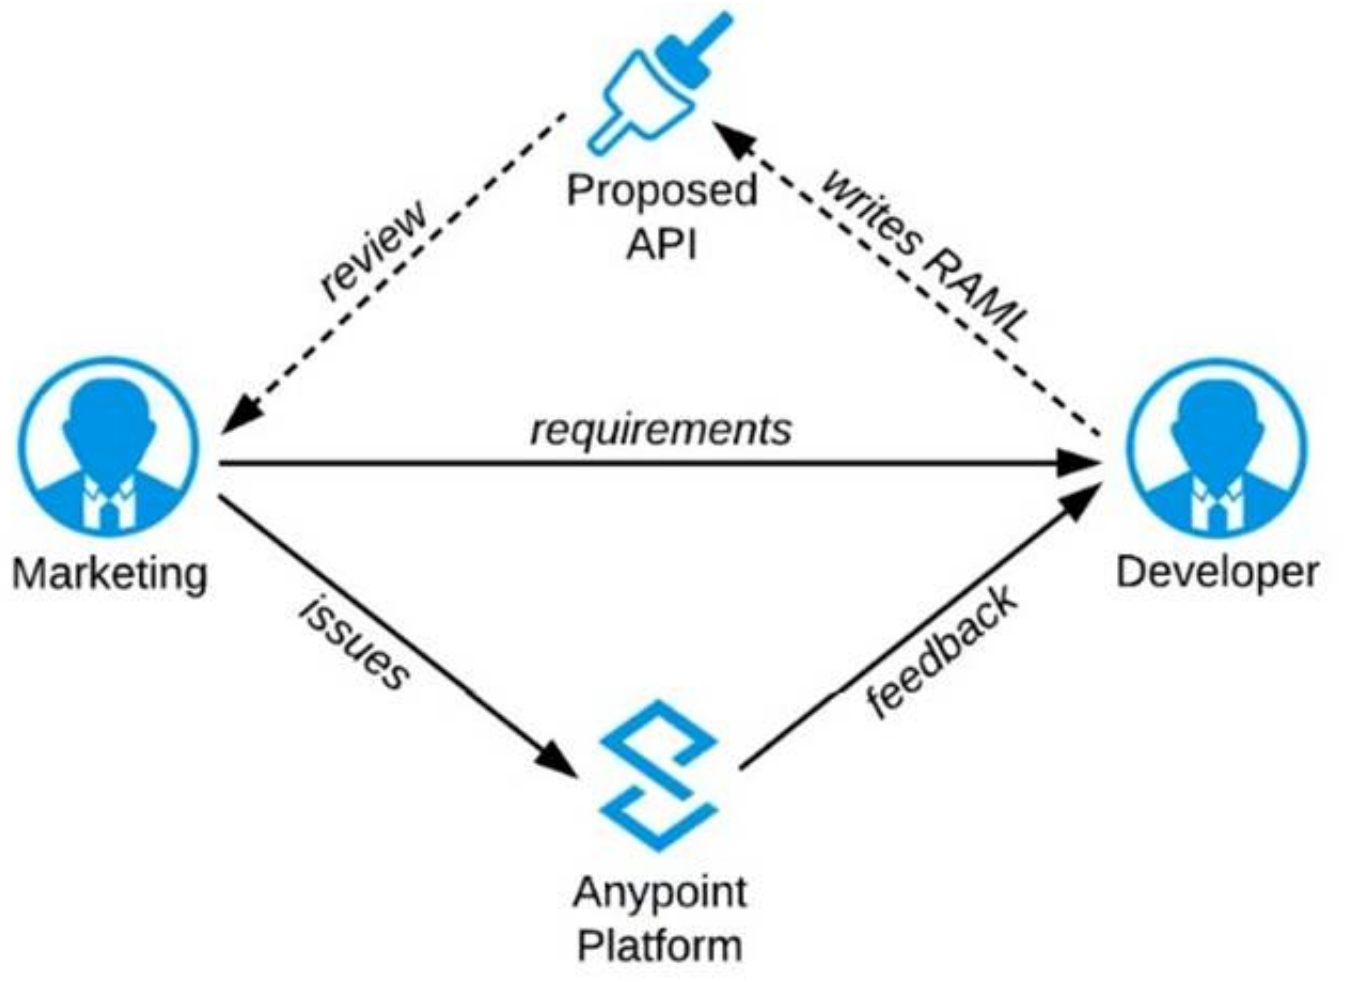

A RAML definition has been proposed for a new Promotions Process API, and has been published to Anypoint Exchange.
The Marketing Department, who will be an important consumer of the Promotions API, has important requirements and expectations that must be met.
What is the most effective way to use Anypoint Platform features to involve the Marketing Department in this early API design phase?
A) Ask the Marketing Department to interact with a mocking implementation of the API using the automatically generated API Console

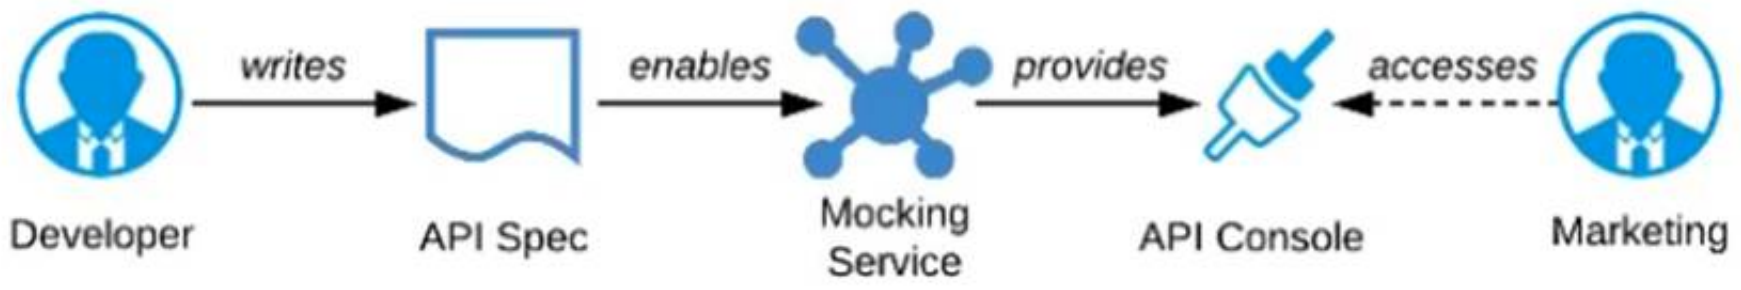

B) Organize a design workshop with the DBAs of the Marketing Department in which the database schema of the Marketing IT systems is translated into RAML

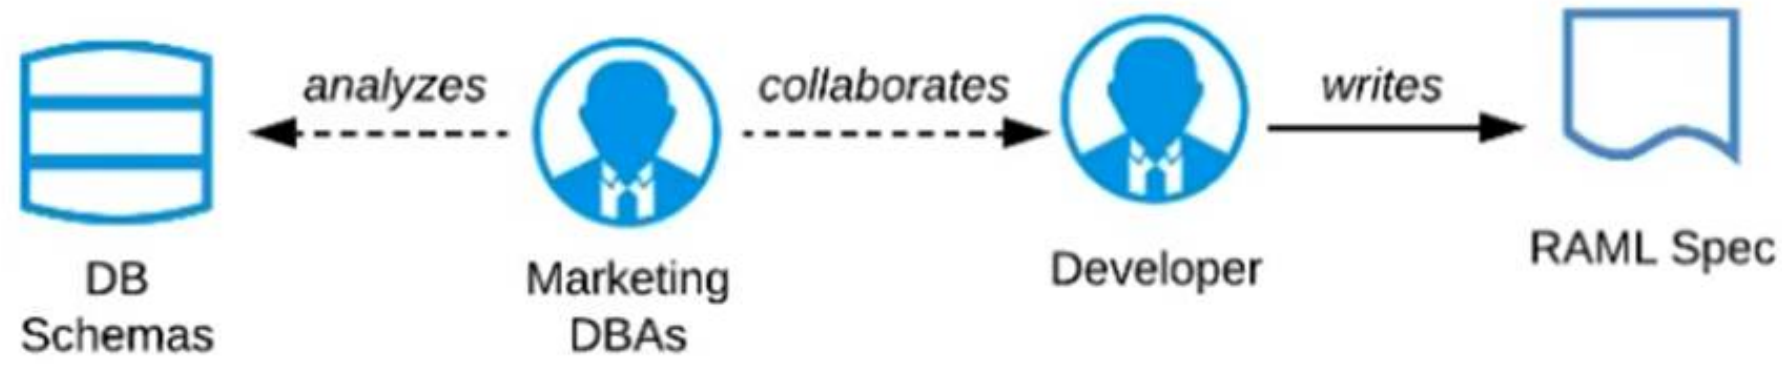

C) Use Anypoint Studio to Implement the API as a Mule application, then deploy that API implementation to CloudHub and ask the Marketing Department to interact with it

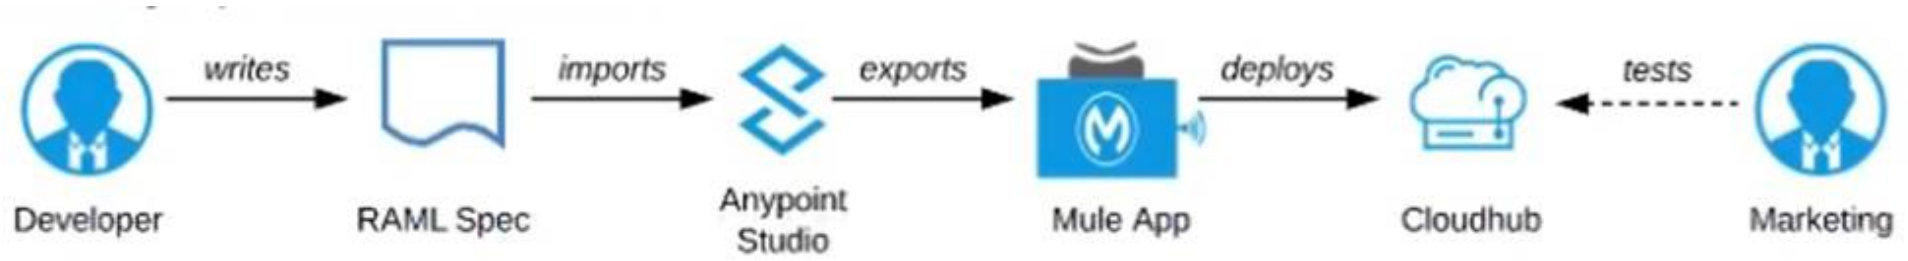

D) Export an integration test suite from API designer and have the Marketing Department execute the tests In that suite to ensure they pass

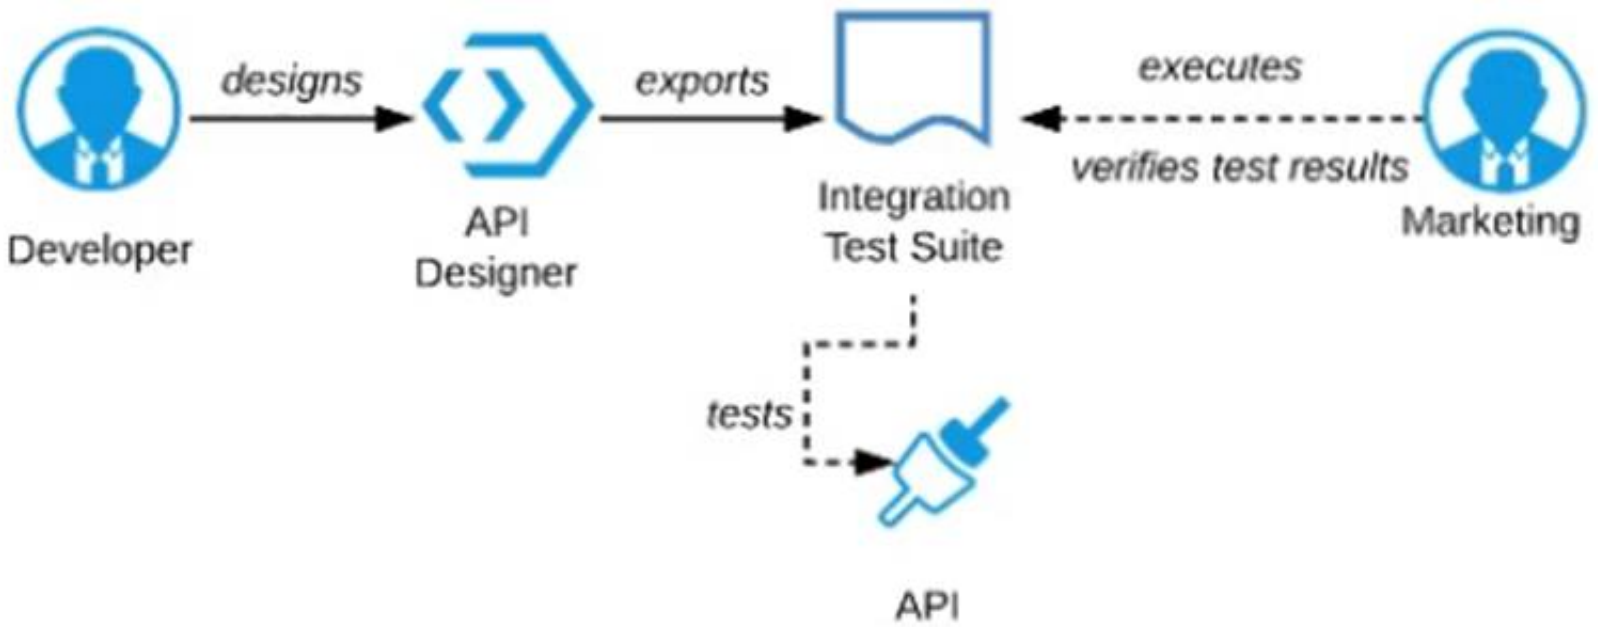

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

NEW QUESTION 126
Say, there is a legacy CRM system called CRM-Z which is offering below functions:
* 1. Customer creation
* 2. Amend details of an existing customer
* 3. Retrieve details of a customer
* 4. Suspend a customer

A. Implement a system API named customerManagement which has all the functionalities wrapped in it as various operations/resources
B. Implement different system APIs named createCustomer, amendCustomer, retrieveCustomer and suspendCustomer as they are modular and has seperation of concerns
C. Implement different system APIs named createCustomerInCRMZ, amendCustomerInCRMZ, retrieveCustomerFromCRMZ and suspendCustomerInCRMZ as they are modular and has seperation of concerns

Answer: B

NEW QUESTION 129
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

NEW QUESTION 130
A large company wants to implement IT infrastructure in its own data center, based on the corporate IT policy requirements that data and metadata reside locally.
Which combination of Mule control plane and Mule runtime plane(s) meets the requirements?

A. Anypoint Platform Private Cloud Edition for the control plane and the MuleSoft-hosted runtime plane
B. The MuleSoft-hosted control plane and Anypoint Runtime Fabric for the runtime plane
C. The MuleSoft-hosted control plane and customer-hosted Mule runtimes for the runtime plane
D. Anypoint Platform Private Cloud Edition for the control plane and customer-hosted Mule runtimes for the runtime plane

Answer: D

NEW QUESTION 131
A large organization with an experienced central IT department is getting started using MuleSoft. There is a project to connect a siloed back-end system to a new Customer Relationship Management (CRM) system. The Center for Enablement is coaching them to use API-led connectivity.
What action would support the creation of an application network using API-led connectivity?

A. Invite the business analyst to create a business process model to specify the canonicaldata model between the two systems
B. Determine if the new CRM system supports the creation of custom: REST APIs, establishes 4 private network with CloudHub, and supports GAuth 2.0 authentication
C. To expedite this project, central IT should extend the CRM system and back-end systems to connect to one another using built in integration interfaces
D. Create a System API to unlock the data on the back-end system using a REST API

Answer: D

NEW QUESTION 135
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

Answer: B

NEW QUESTION 140
Version 3.0.1 of a REST API implementation represents time values in PST time using ISO 8601 hh:mm:ss format. The API implementation needs to be changed to instead represent time values in CEST time using ISO 8601 hh:mm:ss format. When following the semver.org semantic versioning specification, what version should be assigned to the updated API implementation?

A. 3.0.2
B. 4.0.0
C. 3.1.0
D. 3.0.1

Answer: B

NEW QUESTION 145
A Mule application implements an API. The Mule application has an HTTP Listener whose connector configuration sets the HTTPS protocol and hard-codes the port value. The Mule application is deployed to an Anypoint VPC and uses the CloudHub 1.0 Shared Load Balancer (SLB) for all incoming traffic.
Which port number must be assigned to the HTTP Listener's connector configuration so that the Mule application properly receives HTTPS API invocations routed through the SLB?

A. 8082
B. 8092
C. 80
D. 443

Answer: B

NEW QUESTION 147
An API is protected with a Client ID Enforcement policy and uses the default configuration.
Access is requested for the client application to the API, and an approved contract now exists between the client application and the API
How can a consumer of this API avoid a 401 error "Unauthorized or invalid client application credentials"?

A. Send the obtained token as a header in every call
B. Send the obtained: client_id and client_secret in the request body
C. Send the obtained clent_id and clent_secret as URI parameters in every call
D. Send the obtained clent_id and client_secret in the header of every API Request call

Answer: C

NEW QUESTION 150
Refer to the exhibit.

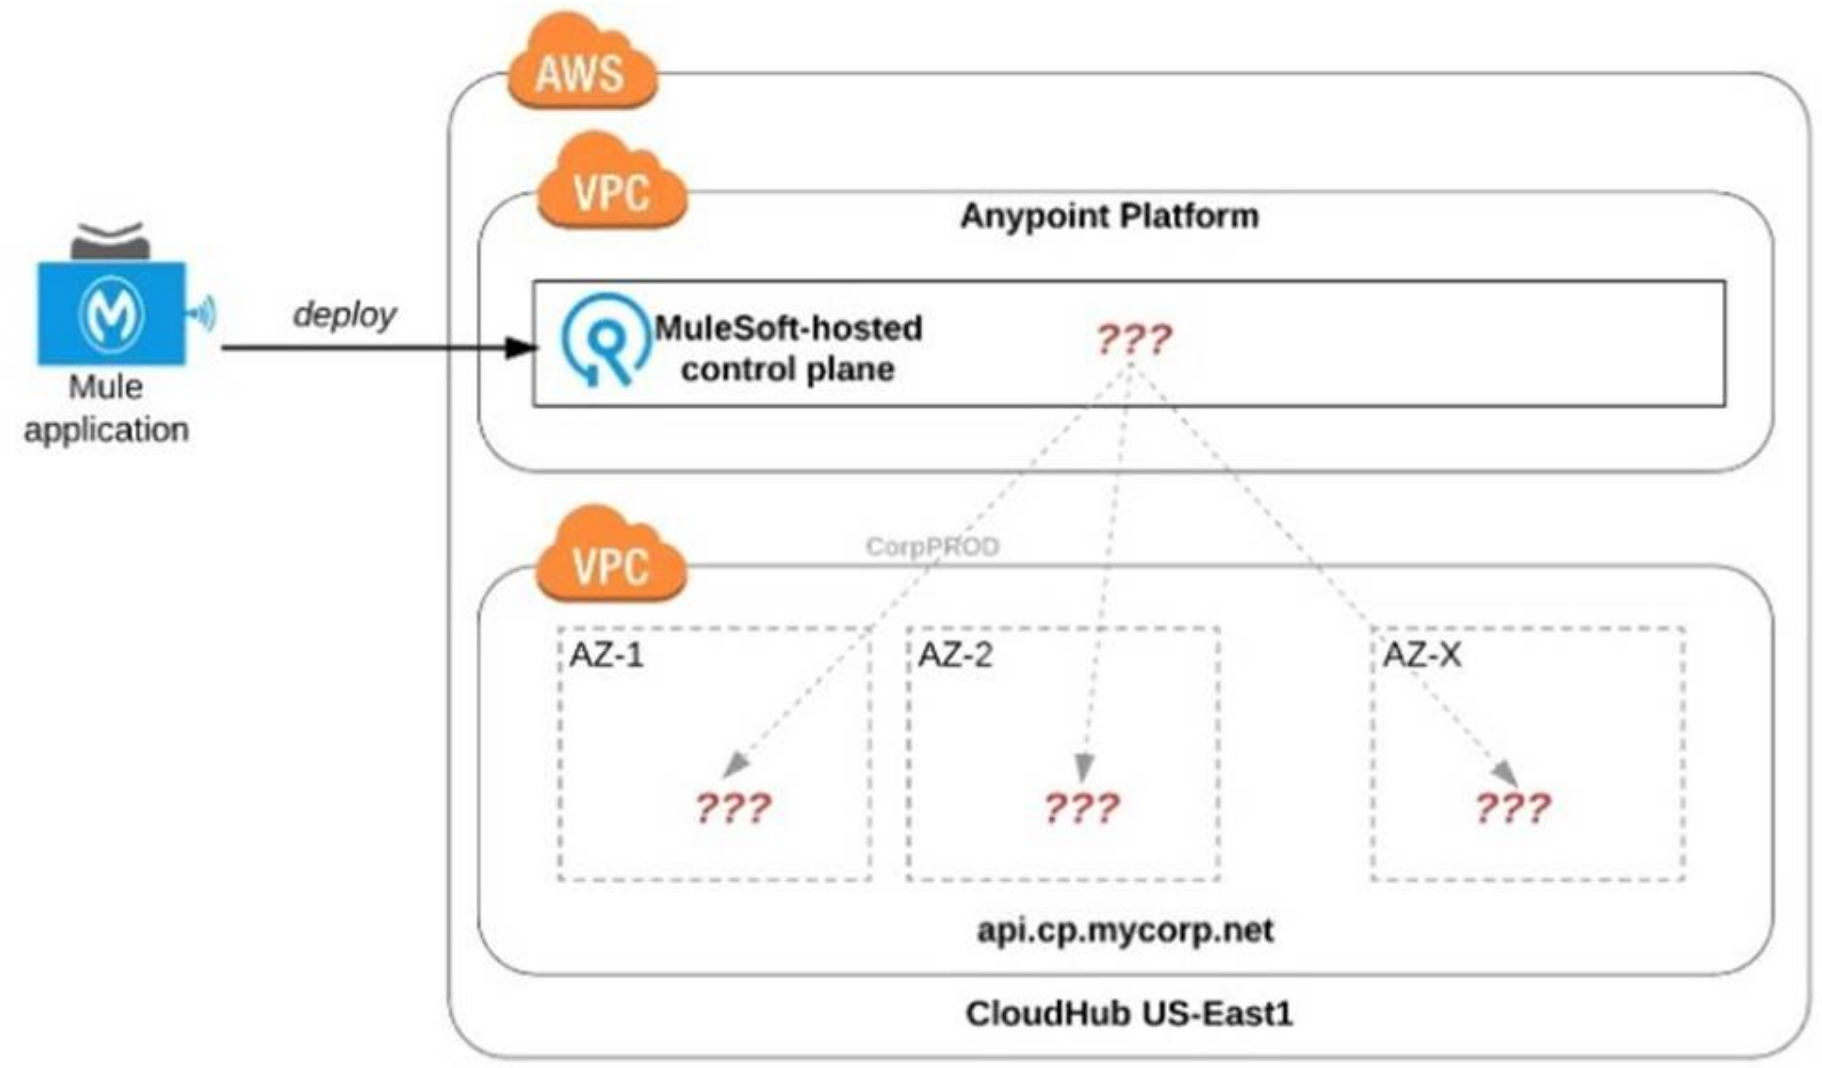

An organization uses one specific CloudHub (AWS) region for all CloudHub deployments. How are CloudHub workers assigned to availability zones (AZs) when the organization's
Mule applications are deployed to CloudHub in that region?

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

NEW QUESTION 152
An organization requires several APIs to be secured with OAuth 2.0, and PingFederate has been identified as the identity provider for API client authorization, The PingFederate Client Provider is configured in access management, and the PingFederate OAuth 2.0 Token Enforcement policy is configured for the API instances required by the organization. The API instances reside in two business groups (Group A and Group B) within the Master Organization (Master Org).
What should be done to allow API consumers to access the API instances?

A. The API administrator should configure the correct client discovery URL in both child business groups, and the API consumer should request access to the API in Ping Identity
B. The API administrator should grant access to the API consumers by creating contracts in the relevant API instances in API Manager
C. The APL consumer should create a client application and request access to the APT in Anypoint Exchange, and the API administrator should approve the request
D. The APT consumer should create a client application and request access to the API in Ping Identity, and the organization's Ping Identity workflow will grant access

Answer: C

NEW QUESTION 157
What are the major benefits of MuleSoft proposed IT Operating Model?

A. * 1. Decrease the IT delivery gap* 2. Meet various business demands without increasing the IT capacity* 3. Focus on creation of reusable assets firs
B. Upon finishing creation of all the possible assets then inform the LOBs in the organization to start using them
C. * 1. Decrease the IT delivery gap* 2. Meet various business demands by increasing the IT capacity and forming various IT departments* 3. Make consumption of assets at the rate of production
D. * 1. Decrease the IT delivery gap* 2. Meet various business demands without increasing the IT capacity* 3. Make consumption of assets at the rate of production

Answer: C

NEW QUESTION 160
An organization has built an application network following the API-led connectivity approach recommended by MuleSoft. To protect the application network against attacks from malicious external API clients, the organization plans to apply JSON Threat Protection policies.
To which API-led connectivity layer should the JSON Threat Protection policies most commonly be applied?

A. All layers

B. System layer
C. Process layer
D. Experience layer

Answer: D

NEW QUESTION 165
A REST API is being designed to implement a Mule application.
What standard interface definition language can be used to define REST APIs?

A. Web Service Definition Language(WSDL)
B. OpenAPI Specification (OAS)
C. YAML
D. AsyncAPI Specification

Answer: B

NEW QUESTION 167
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

NEW QUESTION 171
Which two statements are true about the technology architecture of an Anypoint Virtual Private Cloud (VPC)? Choose 2 answers

A. Ports 8081 and 8082 are used
B. CIDR blacks are used
C. Anypoint VPC is responsible for load balancing the applications
D. Round-robin load balancing is used to distribute client requests across different applications
E. By default, HTTP requests can be made from the public internet to workers at port 6091

Answer: BE

NEW QUESTION 173
An established communications company is beginning its API-led connectivity journey, The company has been using a successful Enterprise Data Model for many years. The company has identified a self-service account management app as the first effort for API- led, and it has identified the following APIs.
Experience layer: Mobile Account Management EAPI, Browser Account Management EAPI Process layer: Customer Lookup PAPI, Service Lookup PAPI, Account Lookup PAPI System layer: Customer SAPI, Account SAPI, Product SAPI, Service SAPI According to MuleSoft's API-led connectivity approach, which API would not be served by the Enterprise Data Model?

A. Customer SAPI
B. Customer Lookup PAPI
C. Mobile Account Management EAPI
D. Service SAPI

Answer: C

NEW QUESTION 176
A customer wants to host their MuleSoft applications in CloudHub 1.0, and these applications should be available at the domain https://api.acmecorp.com.
After creating a dedicated load balancer (DLB) called acme-dib-prod, which further action must the customer take to complete the configuration?

A. Configure the DLB with a TLS certificate for api.acmecorp.com and create an A record for api.acmecorp.com to the public IP addresses associated with their DLB
B. Configure the DLB with a TLS certificate for api.acmecorp.com and create a CNAME record from api.acmecorp.com to acme-dib-prod.|lb.anypointdns.net
C. Configure the DLB with a TLS certificate for acme-dib-prod.Jb.anypointdns.net and create a CNAME record from api.acmecorp:com to acme-dlb-prod.lb.anypointdns.net
D. Configure the DLB with a TLS certificate for aplacmecorp.com and create a CNAME record from api.aomecorp.com to acme-dib-prod.ei.cloubhub.io

Answer: B

NEW QUESTION 177
A team is planning to enhance an Experience API specification, and they are following API- led connectivity design principles.
What is their motivation for enhancing the API?

A. The primary API consumer wants certain kinds of endpoints changed from the Center for Enablement standard to the consumer system standard
B. The underlying System API is updated to provide more detailed data for several heavily used resources

C. An IP Allowlist policy is being added to the API instances in the Development and Staging environments
D. A Canonical Data Model is being adopted that impacts several types of data included in the API

Answer: D

NEW QUESTION 179
An API implementation is being designed that must invoke an Order API, which is known to repeatedly experience downtime.
For this reason, a fallback API is to be called when the Order API is unavailable.
What approach to designing the invocation of the fallback API provides the best resilience?

A. Search Anypoint Exchange for a suitable existing fallback API, and then implement invocations to this fallback API in addition to the Order API
B. Create a separate entry for the Order API in API Manager, and then invoke this API as a fallback API if the primary Order API is unavailable
C. Redirect client requests through an HTTP 307 Temporary Redirect status code to the fallback API whenever the Order API is unavailable
D. Set an option in the HTTP Requester component that invokes the Order API to instead invoke a fallback API whenever an HTTP 4xx or 5xx response status code is returned from the Order API

Answer: A

NEW QUESTION 180
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

Answer: A

NEW QUESTION 183
.....

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

MuleSoft-Platform-Architect-I Practice Exam Features:

* MuleSoft-Platform-Architect-I Questions and Answers Updated Frequently

* MuleSoft-Platform-Architect-I Practice Questions Verified by Expert Senior Certified Staff

* MuleSoft-Platform-Architect-I Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* MuleSoft-Platform-Architect-I Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year