Salesforce

Exam Questions MuleSoft-Platform-Architect-I

Salesforce Certified MuleSoft Platform Architect 1 Exam (SP24)

NEW QUESTION 1
Which scenario is suited for MUnit tests instead of integration tests?

A. For read-only interactions to any dependencies (such as other web APIs)
B. When testing does not require knowledge of implementation details
C. When no mocking is permissible
D. For tests that are implemented using SoapUI

Answer: A

NEW QUESTION 2
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 3
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

NEW QUESTION 4
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

NEW QUESTION 5
Which APIs can be used with DataGraph to create a unified schema?

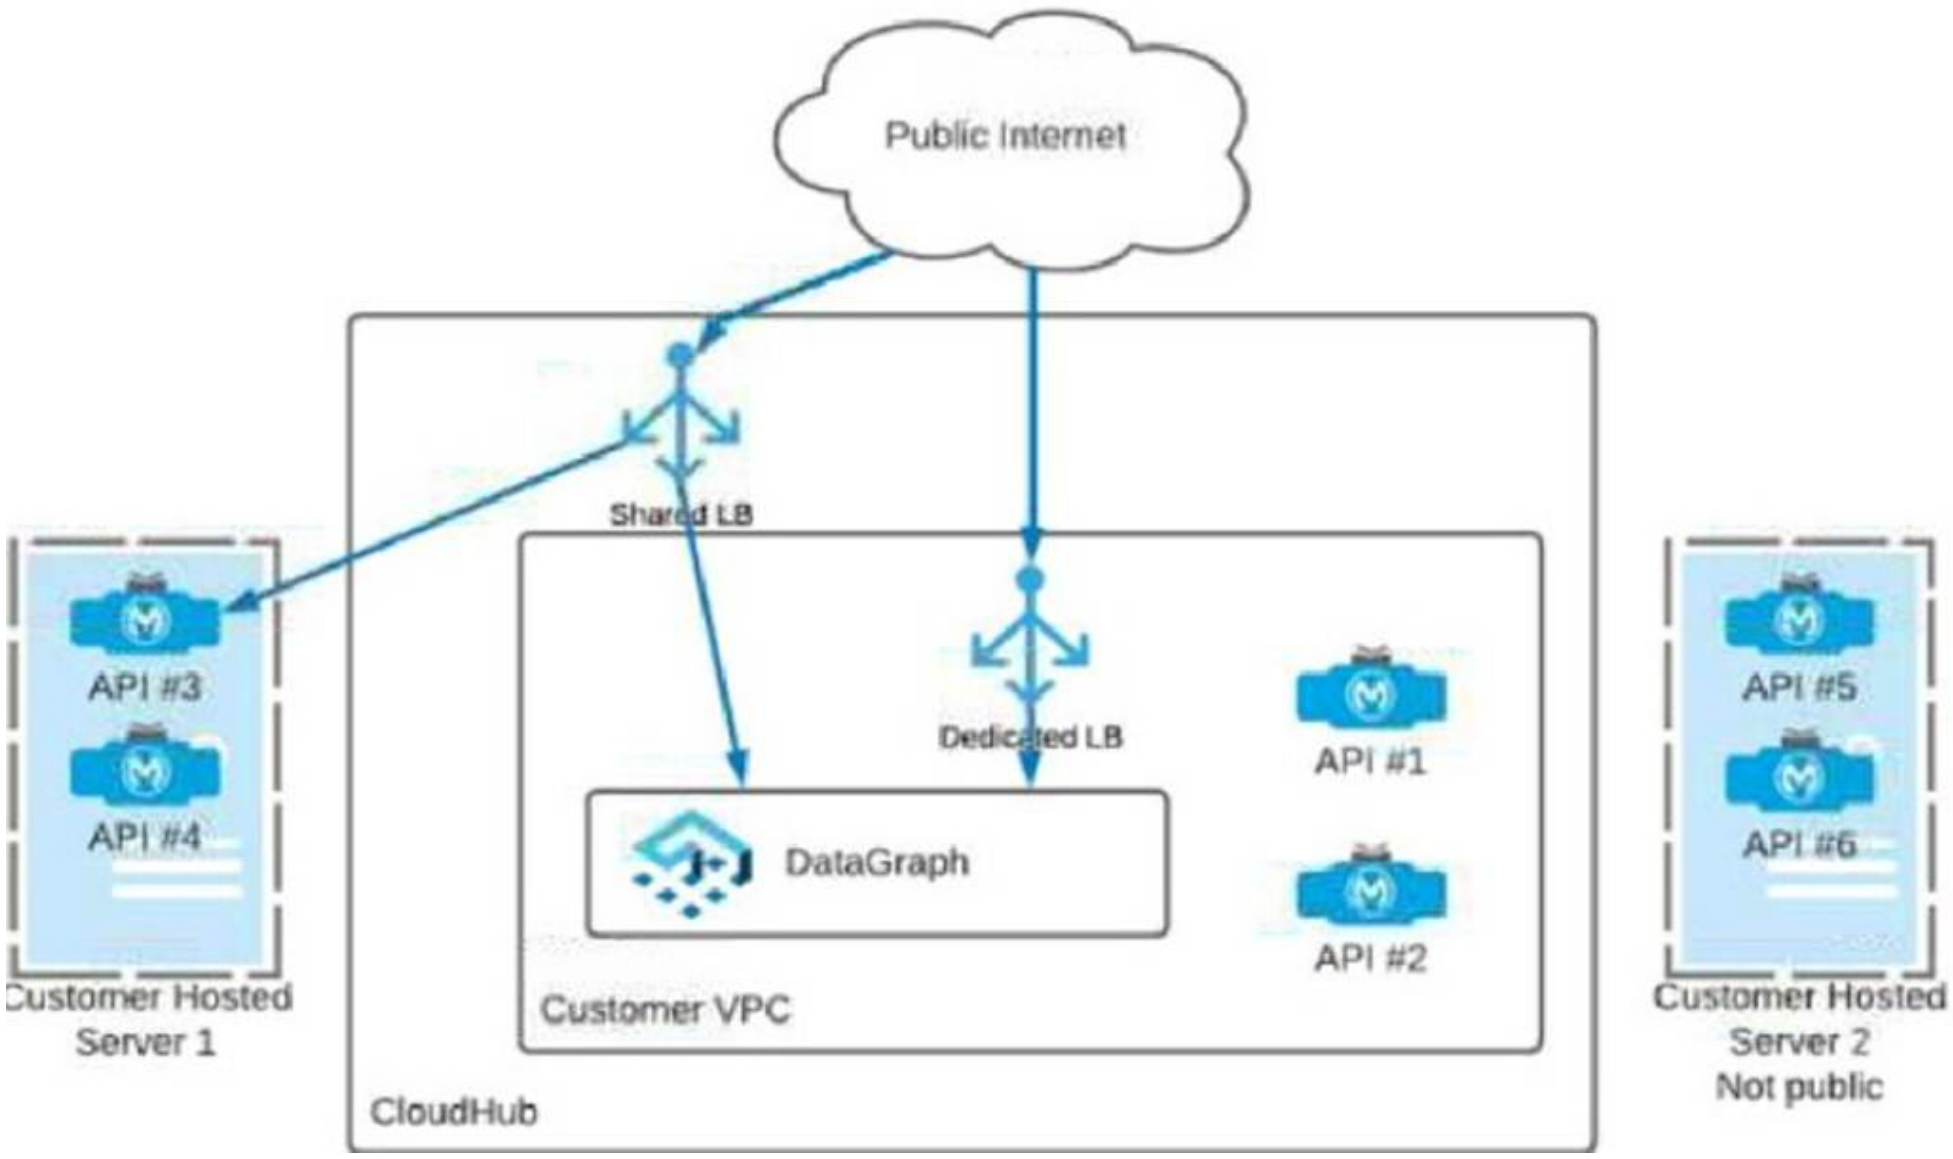

A. APIs 1, 3, 5
B. APIs 2, 4 ,6
C. APIs 1, 2, s5, 6
D. APIs 1, 2, 3, 4

Answer: D

NEW QUESTION 6
A Platinum customer uses the U.S. control plane and deploys applications to CloudHub in Singapore with a default log configuration.
The compliance officer asks where the logs and monitoring data reside?

A. Logs are held in:Singapore and monitoring data is held in the United States
B. Logs and monitoring data are held in the United States
C. Logs are held in the United States and monitoring data is held in Singapore
D. Logs and monitoring data are held in Singapore

Answer: B

NEW QUESTION 7
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 8
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 9

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 10
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 10
An API with multiple API implementations (Mule applications) is deployed to both CloudHub and customer-hosted Mule runtimes. All the deployments are managed by the MuleSoft-hosted control plane. An alert needs to be triggered whenever an API implementation stops responding to API requests, even if no API clients have called the API implementation for some time.
What is the most effective out-of-the-box solution to create these alerts to monitor the API implementations?

A. Create monitors in Anypoint Functional Monitoring for the API implementations, where each monitor repeatedly invokes an API implementation endpoint
B. Add code to each API client to send an Anypoint Platform REST API request to generate a custom alert in Anypoint Platform when an API invocation times out
C. Handle API invocation exceptions within the calling API client and raise an alert from that API client when such an exception is thrown
D. Configure one Worker Not Responding alert.in Anypoint Runtime Manager for all API implementations that will then monitor every API implementation

Answer: A

NEW QUESTION 14
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 19
An application updates an inventory running only one process at any given time to keep the inventory consistent. This process takes 200 milliseconds (.2 seconds) to execute; therefore, the scalability threshold of the application is five requests per second.
What is the impact on the application if horizontal scaling is applied, thereby increasing the number of Mule workers?

A. The application scalability threshold is five requests per second regardless of the horizontal scaling
B. The total process execution time is now 100 milliseconds (.1 seconds)
C. The application scalability threshold is now 10 requests per second
D. Horizontal scaling cannot be applied to an already-running application

Answer: A

NEW QUESTION 22
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

NEW QUESTION 26

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

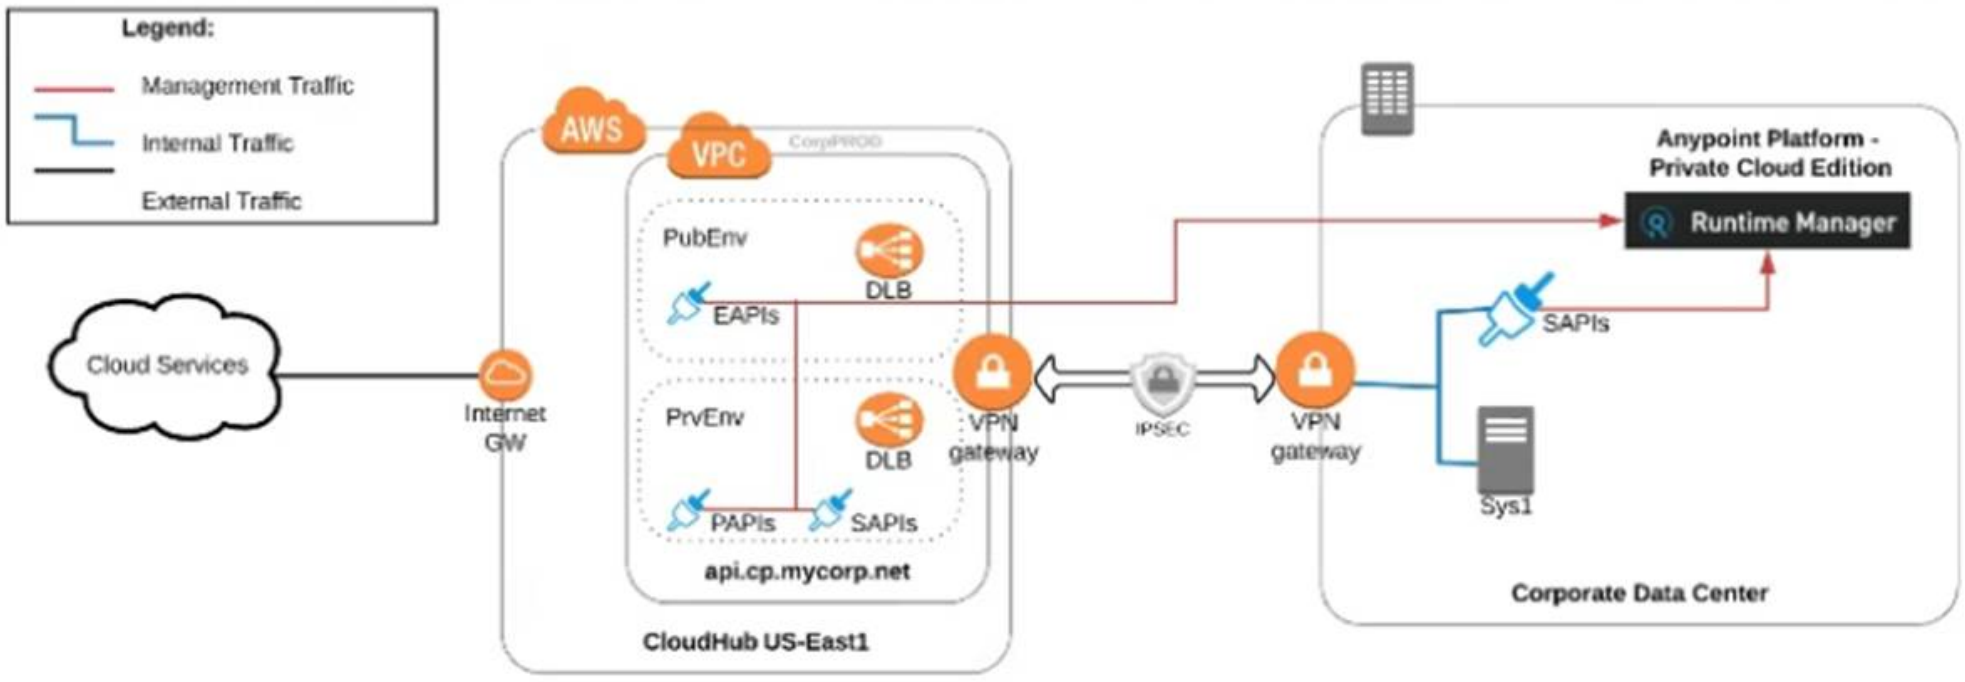

B) Use CloudHub-deployed Mule runtimes in the shared worker cloud managed by the MuleSoft-hosted Anypoint Platform control plane

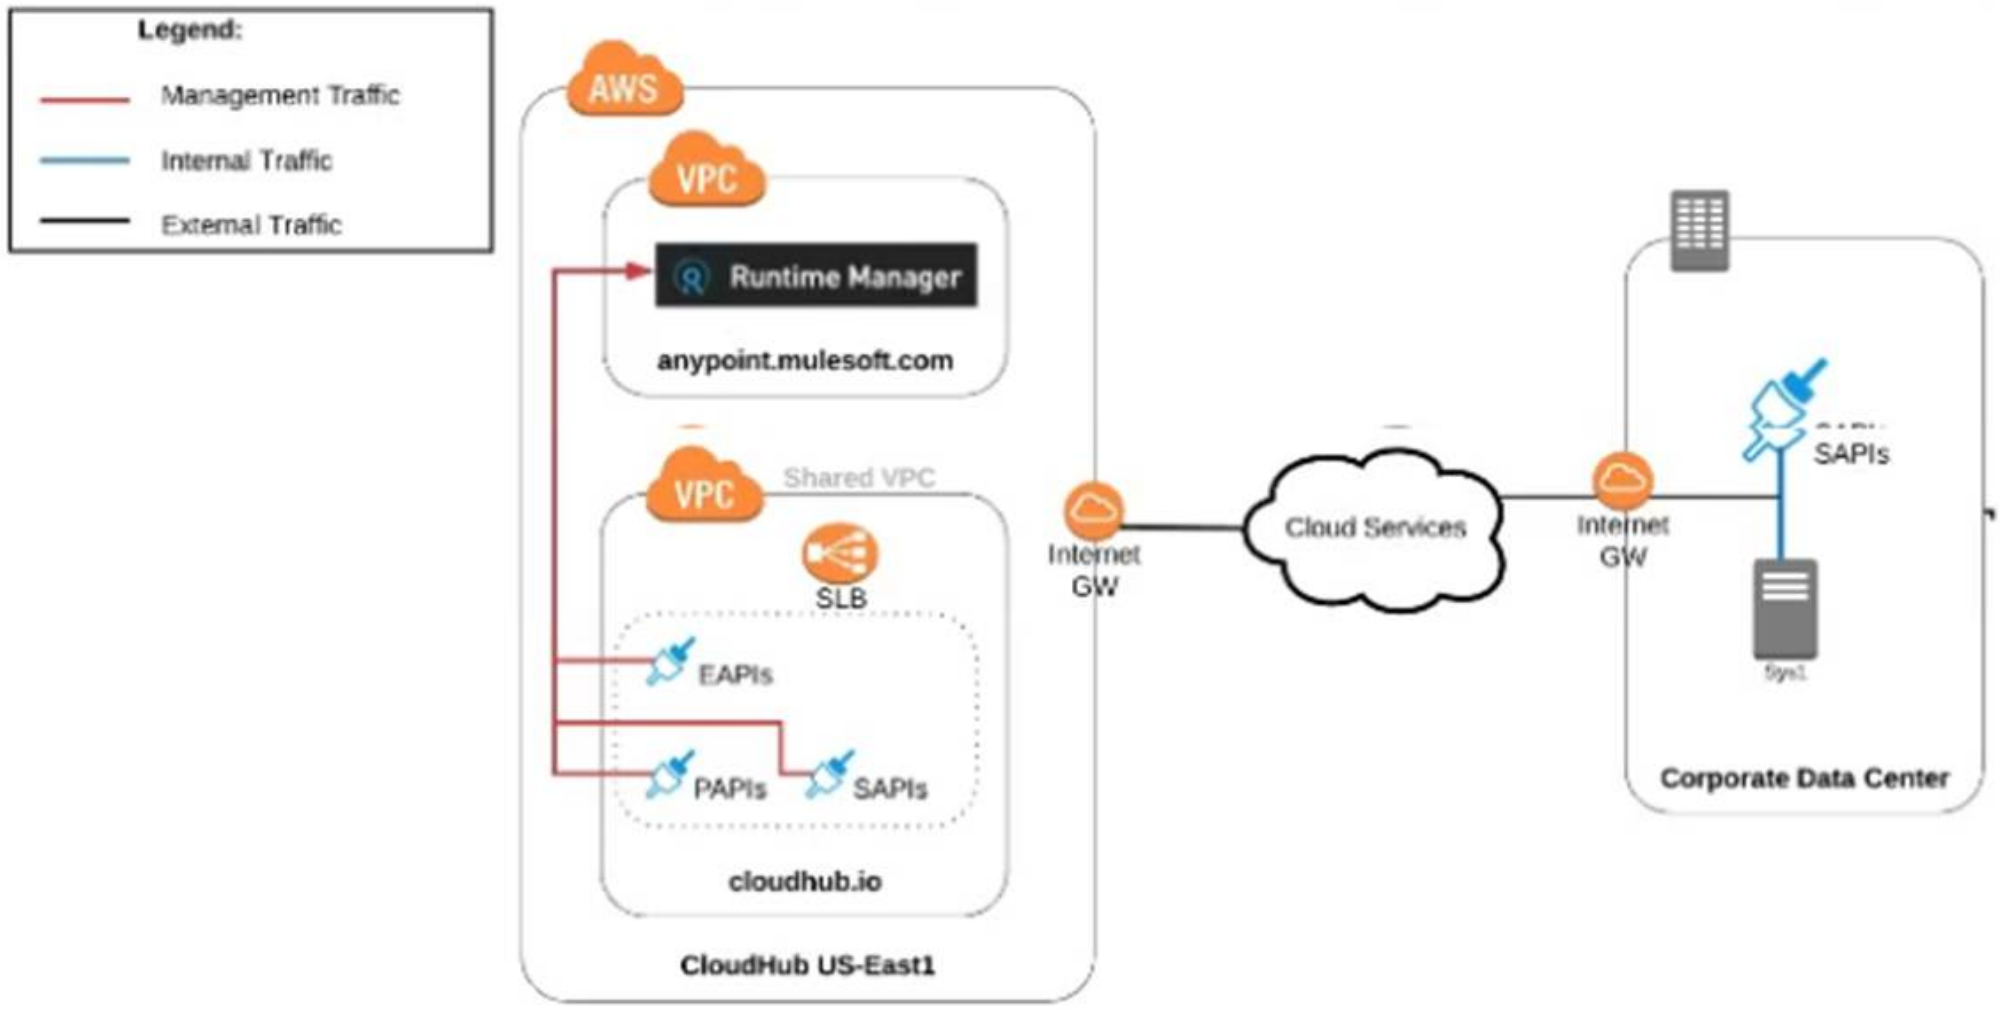

C) Use an on-premises installation of Mule runtimes that are completely isolated with NO external network access, managed by the Anypoint Platform Private Cloud Edition control plane

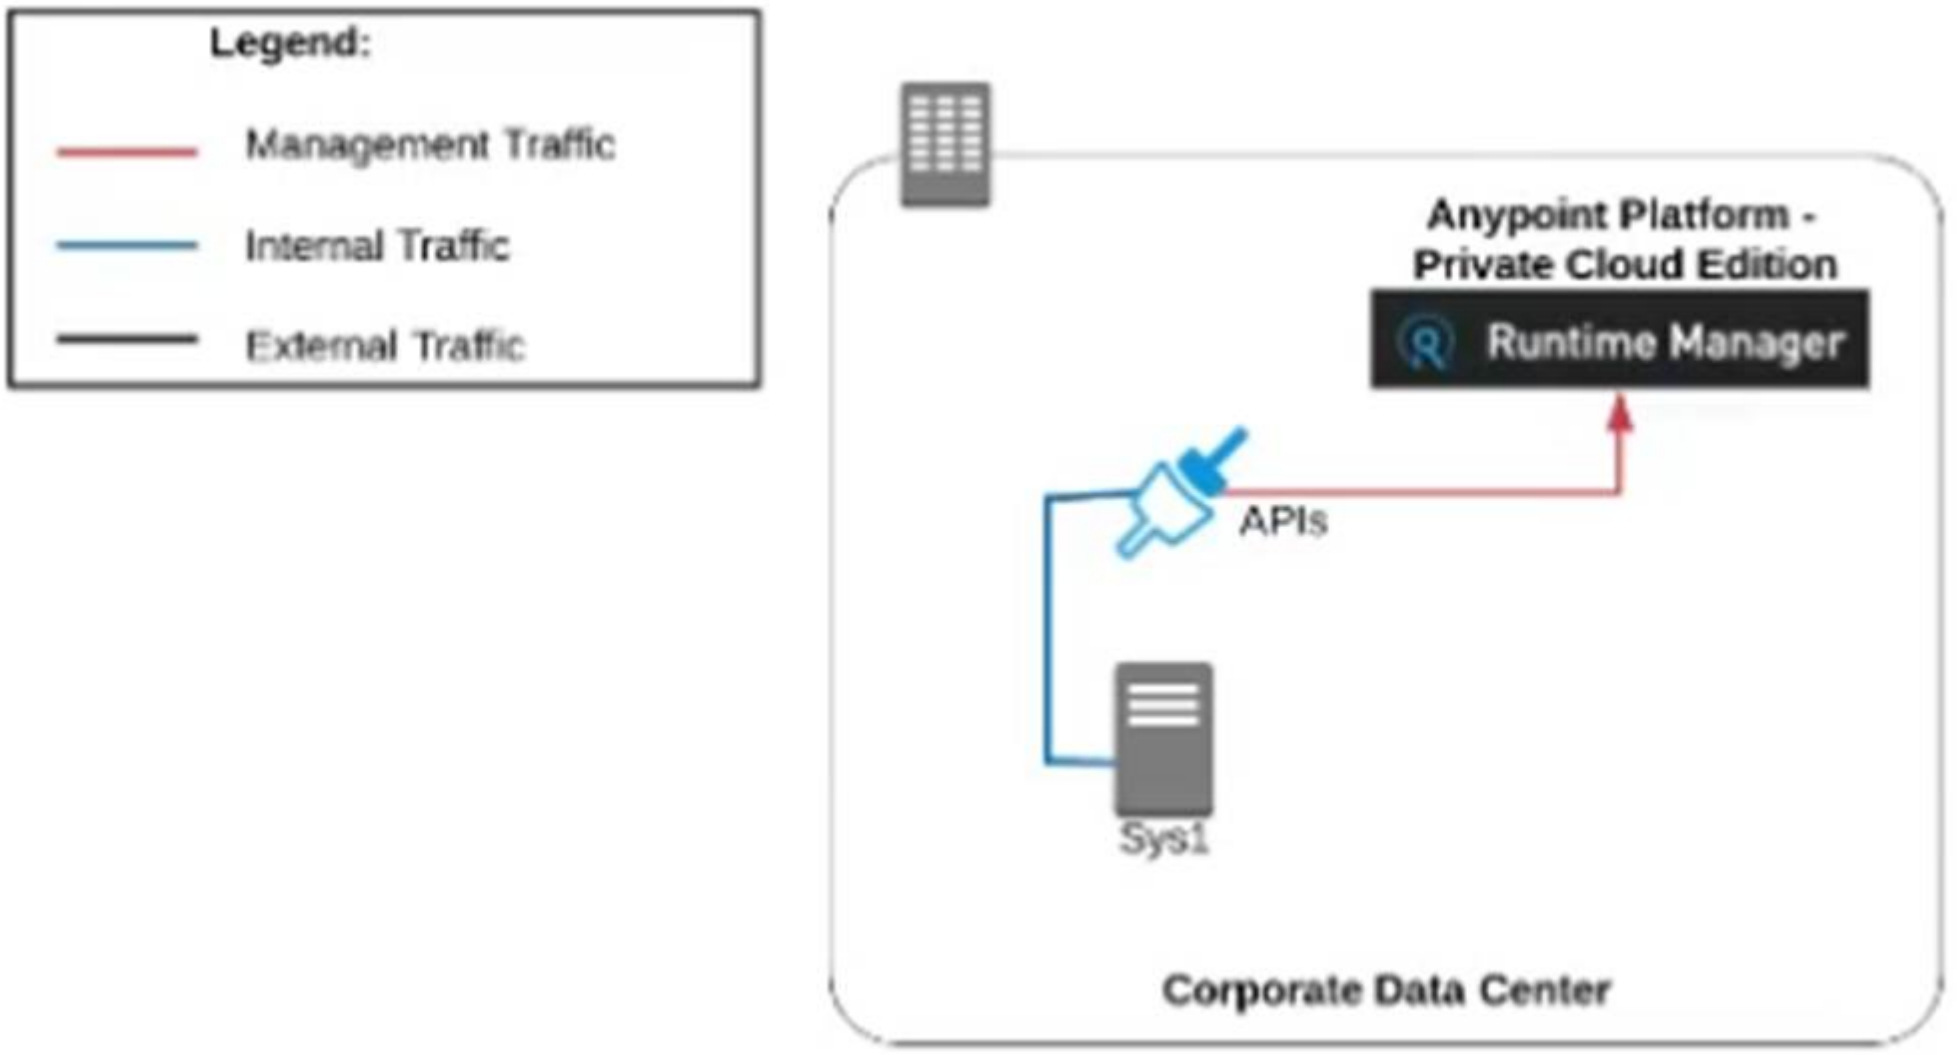

D) Use a combination of Cloud Hub-deployed and manually provisioned on-premises Mule runtimes managed by the MuleSoft-hosted Anypoint Platform control plane

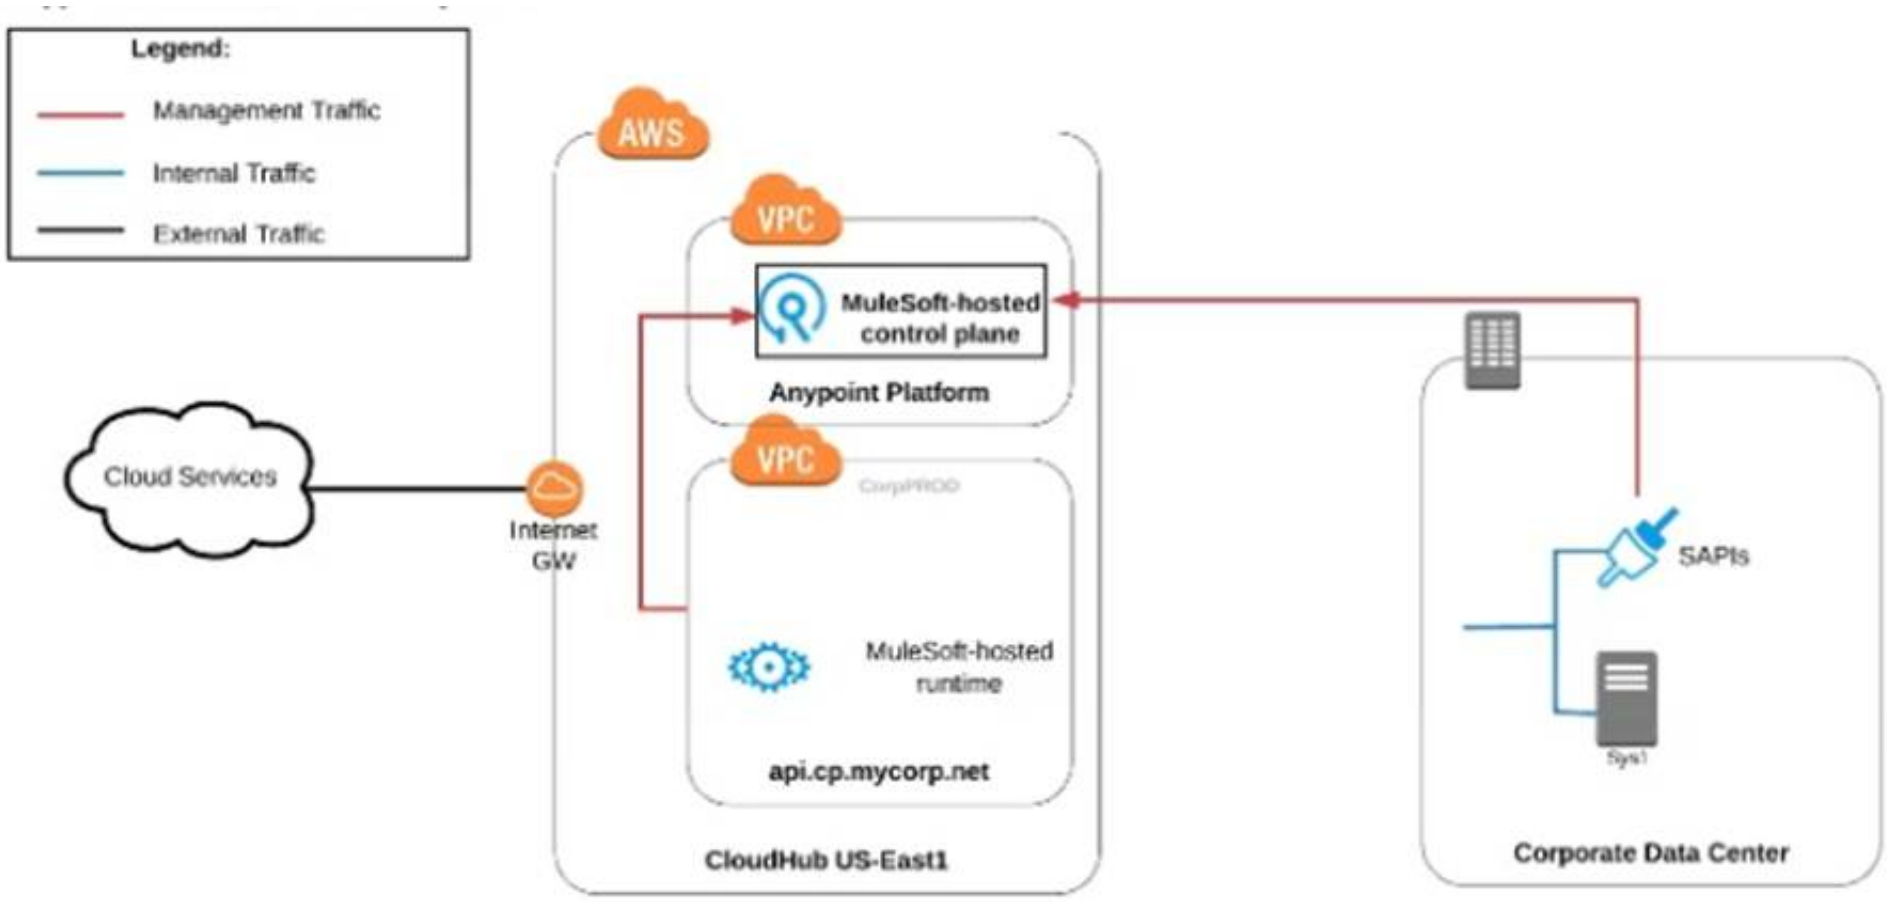

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

NEW QUESTION 28
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 31
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 36
In which layer of API-led connectivity, does the business logic orchestration reside?

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

Answer: C

NEW QUESTION 39
Refer to the exhibits.

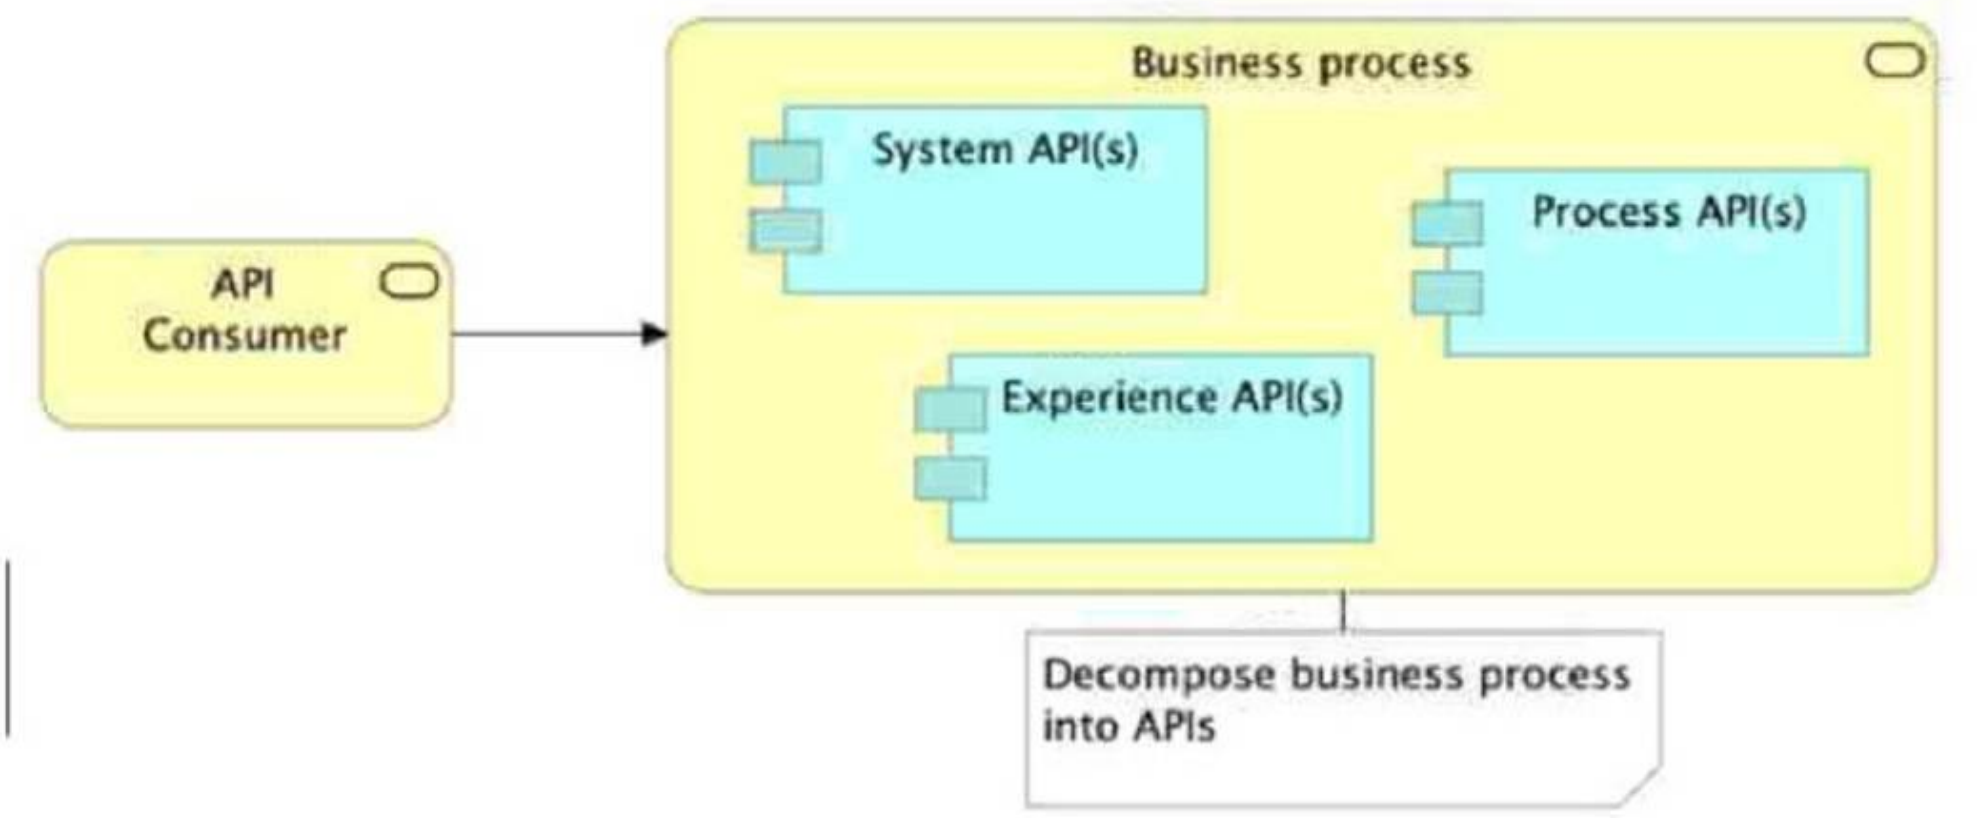

Which architectural constraint is compatible with the API-led connectivity architectural style?

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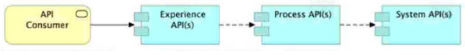

B. Use a Process API to-orchestrate calls to multiple System APIs but not to other Process APIs:

C. Allow System APIs to return data that is not currently required by the identified Process or Experience APIs 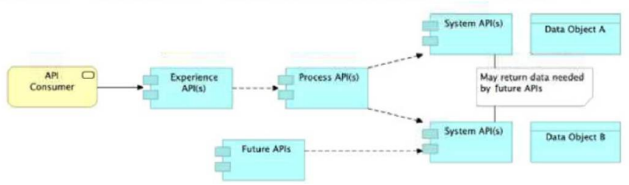

D. Handle customizations for the end-user application at the Process layer rather than at the Experience layer

Answer: B

NEW QUESTION 42
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 44
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 49
Due to a limitation in the backend system, a system API can only handle up to 500 requests per second. What is the best type of API policy to apply to the system API to avoid overloading the backend system?

A. Rate limiting
B. HTTP caching
C. Rate limiting - SLA based
D. Spike control

Answer: D

NEW QUESTION 54
An online store's marketing team has noticed an increase in customers leaving online baskets without checking out. They suspect a technology issue is at the root cause of the baskets being left behind. They approach the Center for Enablement to ask for help identifying the issue. Multiple APIs from across all the layers of their application network are involved in the shopping application.
Which feature of the Anypoint Platform can be used to view metrics from all involved APIs at the same time?

A. Custom dashboards
B. Built-in dashboards
C. Functional monitoring
D. API Manager

Answer: B

NEW QUESTION 58
A Platform Architect inherits a legacy monolithic SOAP-based web service that performs a number of tasks, including showing all policies belonging to a client. The service connects to two back-end systems — a life-insurance administration system and a general-insurance administration system — and then queries for insurance policy information within each system, aggregates the results, and presents a SOAP-based response to a user interface (UI).
The architect wants to break up the monolithic web service to follow API-led conventions. Which part of the service should be put into the process layer?

A. Combining the insurance policy information from the administration systems
B. Presenting the SOAP-based response to the UI
C. Authenticating and maintaining connections to each of the back-end administration systems
D. Querying the data from the administration systems

Answer: A

NEW QUESTION 62
A code-centric API documentation environment should allow API consumers to investigate and execute API client source code that demonstrates invoking one or more APIs as part of representative scenarios.
What is the most effective way to provide this type of code-centric API documentation environment using Anypoint Platform?

A. Enable mocking services for each of the relevant APIs and expose them via their Anypoint Exchange entry
B. Ensure the APIs are well documented through their Anypoint Exchange entries and API Consoles and share these pages with all API consumers
C. Create API Notebooks and include them in the relevant Anypoint Exchange entries
D. Make relevant APIs discoverable via an Anypoint Exchange entry

Answer: C

NEW QUESTION 65
4A developer for a transportation organization is implementing exactly one processing functionality in a Reservation Mule application to process and store passenger records. This Reservation application will be deployed to multiple CloudHub workers/replicas. It is possible that several external systems could send duplicate passenger records to the Reservation application.
An appropriate storage mechanism must be selected to help the Reservation application process each passenger record exactly once as much as possible. The selected storage mechanism must be shared by all the CloudHub workers/replicas in order to synchronize the state information to assist attempting exactly once processing of each passenger record by the deployed Reservation Mule application.
Which type of simple storage mechanism in Anypoint Platform allows the Reservation Mule application to update and share data between the CloudHub workers/replicas exactly
once, with minimal development effort?

A. Persistent Object Store
B. Runtime Fabric Object Store
C. Non-persistent Object Store
D. In-memory Mule Object Store

Answer: A

NEW QUESTION 68

Refer to the exhibit. An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields.
The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?

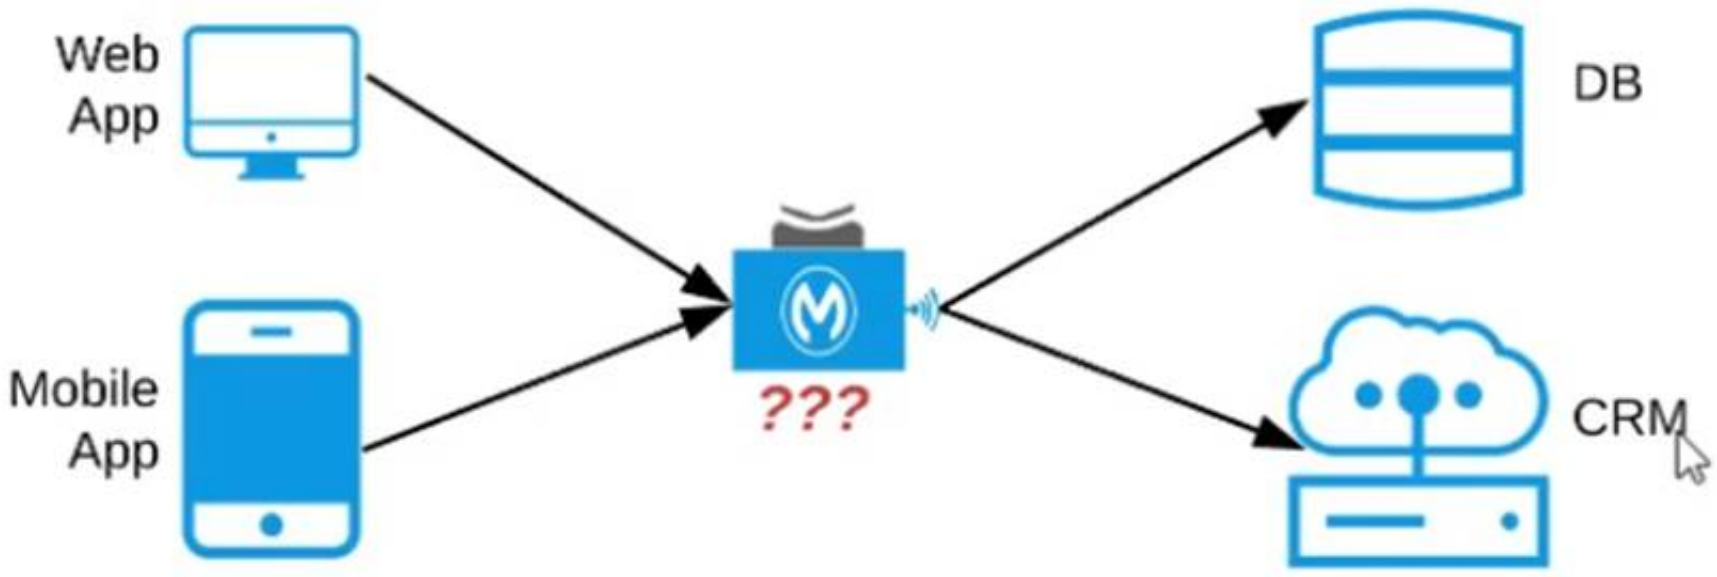

A) A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes

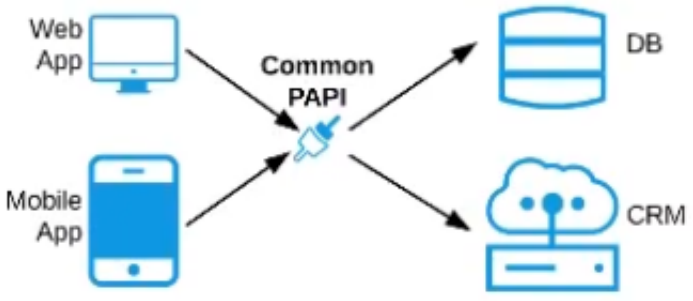

B) One set of APIs (Experience API, Process API, and System API) for the web app, and another set for the mobile app

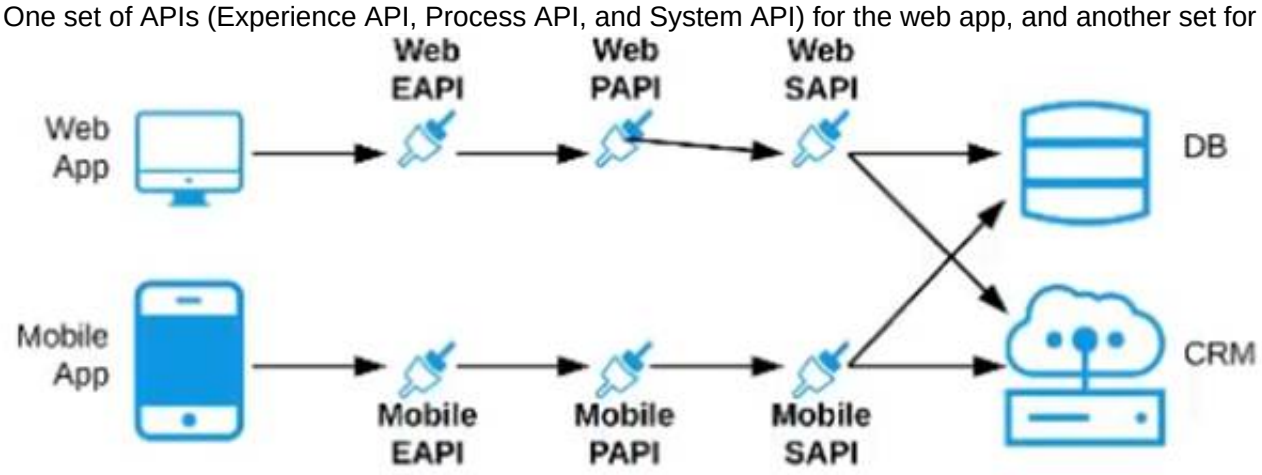

C) Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system

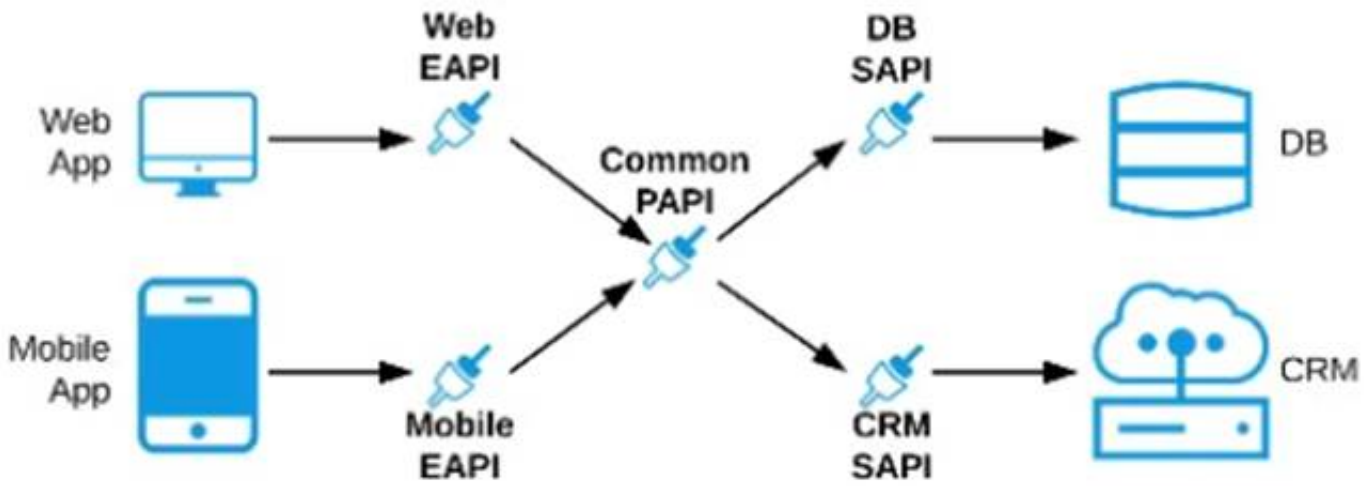

D) A common Experience API used by both the web and mobile apps, but separate Process APIs for the web and mobile apps that interact with the database and the CRM System

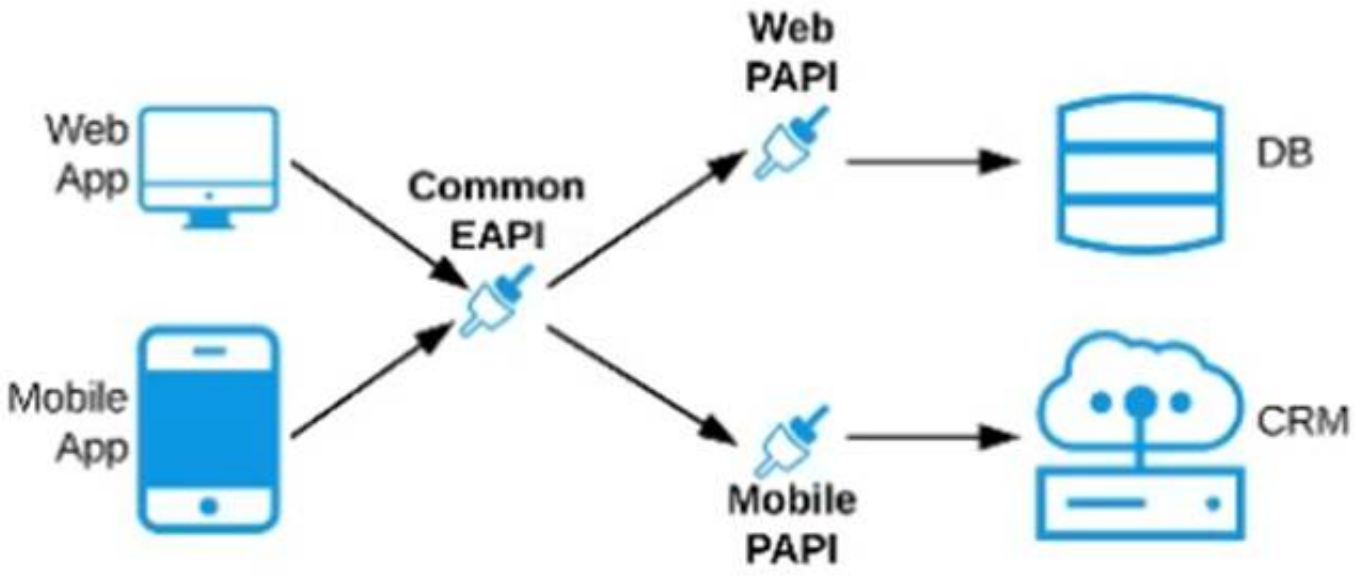

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

NEW QUESTION 72
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 75
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 78
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 83
An API client calls one method from an existing API implementation. The API implementation is later updated. What change to the API implementation would require the API client's invocation logic to also be updated?

A. When the data type of the response is changed for the method called by the API client
B. When a new method is added to the resource used by the API client
C. When a new required field is added to the method called by the API client
D. When a child method is added to the method called by the API client

Answer: C

NEW QUESTION 85
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 87
What is true about API implementations when dealing with legal regulations that require all data processing to be performed within a certain jurisdiction (such as in the USA or the EU)?

A. They must avoid using the Object Store as it depends on services deployed ONLY to the US East region
B. They must use a Jurisdiction-local external messaging system such as Active MQ rather than Anypoint MQ
C. They must te deployed to Anypoint Platform runtime planes that are managed by Anypoint Platform control planes, with both planes in the same Jurisdiction
D. They must ensure ALL data is encrypted both in transit and at rest

Answer: C

NEW QUESTION 88
When should idempotency be taken into account?

A. When making requests to update currently locked entities
B. When storing the results of s previous request for use in response to subsequent requests
C. When sending concurrent update requests for the same entity
D. When preventing duplicate processing from multiple sent requests

Answer: D

NEW QUESTION 90
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 94
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 97
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 99
.....

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