Salesforce

Exam Questions MuleSoft-Platform-Architect-I

Salesforce Certified MuleSoft Platform Architect 1 Exam (SP24)

NEW QUESTION 1
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

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 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 4
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 5
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 6
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 7
What is true about where an API policy is defined in Anypoint Platform and how it is then applied to API instances?

A. The API policy Is defined In Runtime Manager as part of the API deployment to a Mule runtime, and then ONLY applied to the specific API Instance
B. The API policy Is defined In API Manager for a specific API Instance, and then ONLY applied to the specific API instance
C. The API policy Is defined in API Manager and then automatically applied to ALL API instances
D. The API policy is defined in API Manager, and then applied to ALL API instances in the specified environment

Answer: B

NEW QUESTION 8
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 9
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 10
Which APIs can be used with DataGraph to create a unified schema?

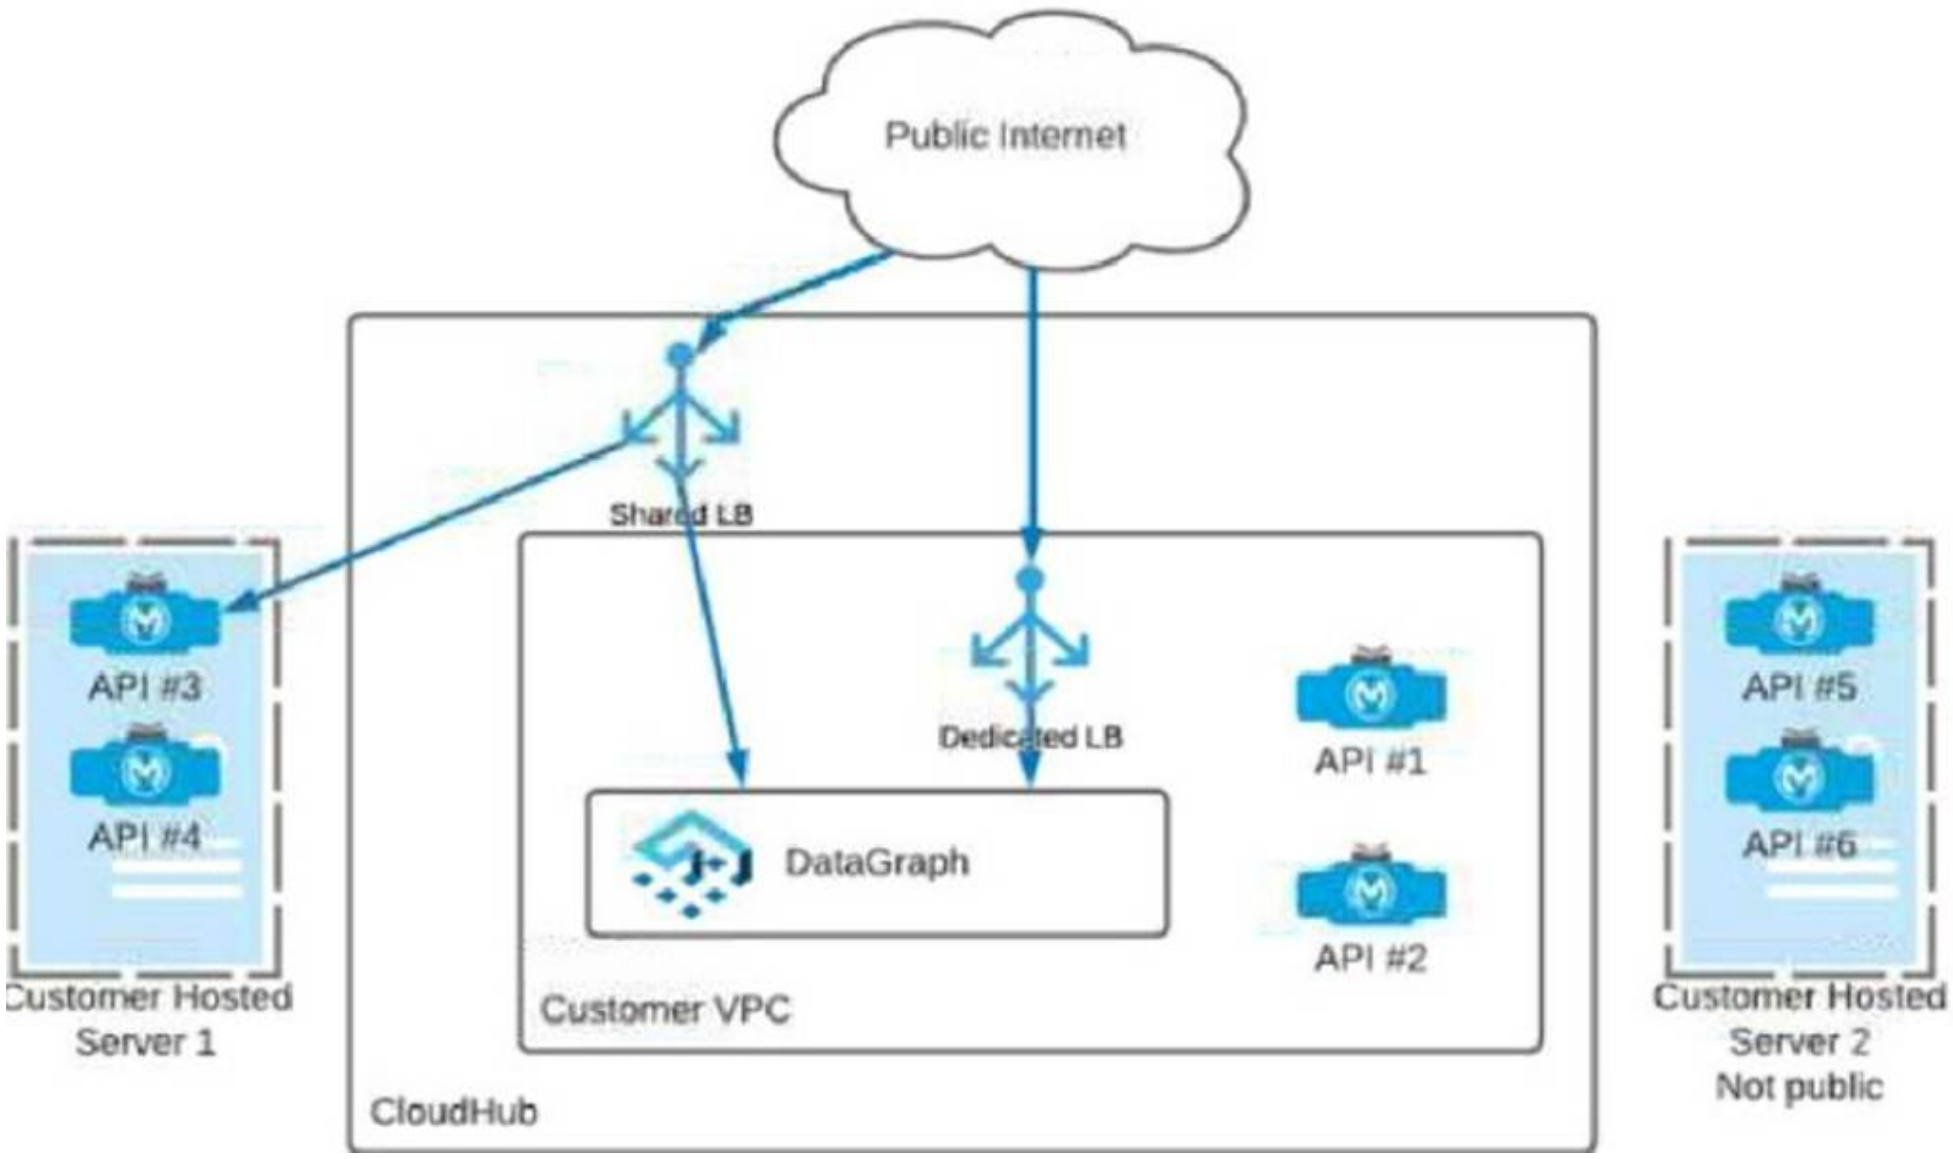

A. APIs 1, 3, 5
B. APIs 2, 4 ,6
C. APIs 1, 2, s5, 6
D. APIs 1, 2, 3, 4

Answer: D

NEW QUESTION 10
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 12
Refer to the exhibit.

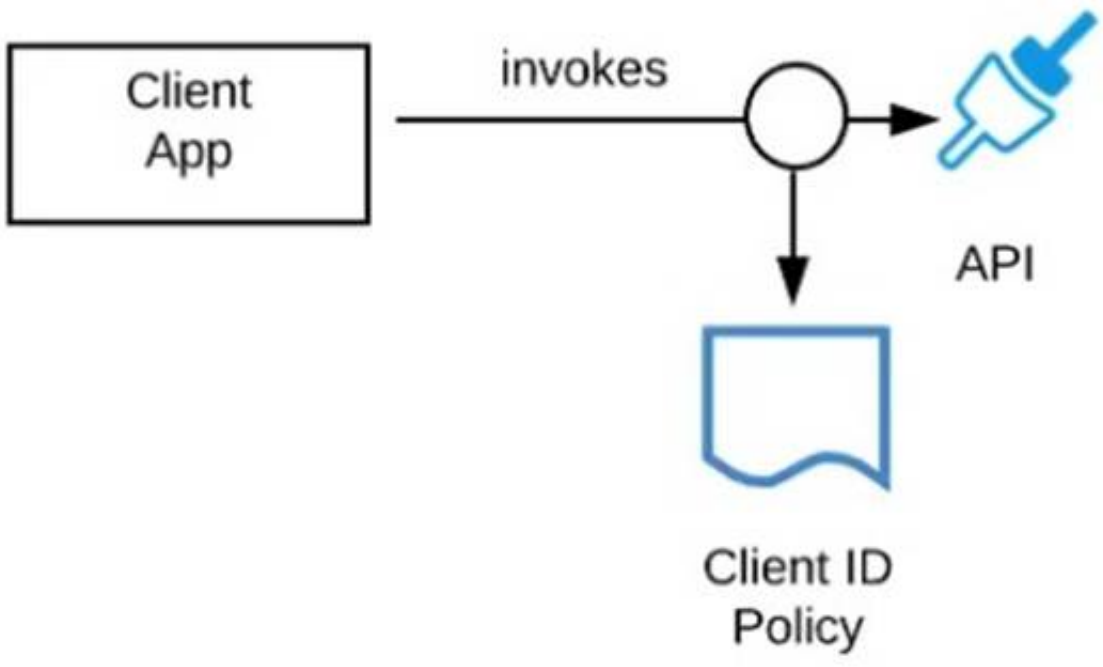

A developer is building a client application to invoke an API deployed to the STAGING environment that is governed by a client ID enforcement policy.
What is required to successfully invoke the API?

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 13
An API has been updated in Anypoint exchange by its API producer from version 3.1.1 to 3.2.0 following accepted semantic versioning practices and the changes have been communicated via the APIs public portal. The API endpoint does NOT change in the new version. How should the developer of an API client respond to this change?

A. The API producer should be requested to run the old version in parallel with the new one
B. The API producer should be contacted to understand the change to existing functionality
C. The API client code only needs to be changed if it needs to take advantage of the new features
D. The API clients need to update the code on their side and need to do full regression

Answer: C

NEW QUESTION 18
A retail company is using an Order API to accept new orders. The Order API uses a JMS queue to submit orders to a backend order management service. The normal load for orders is being handled using two (2) CloudHub workers, each configured with 0.2 vCore. The CPU load of each CloudHub worker normally runs well below 70%. However, several times during the year the Order API gets four times (4x) the average number of orders. This causes the CloudHub worker CPU load to exceed 90% and the order submission time to exceed 30 seconds. The cause, however, is NOT the backend order management service, which still responds fast enough to meet the response SLA for the Order API. What is the MOST resource-efficient way to configure the Mule application's CloudHub deployment to help the company cope with this performance challenge?

A. Permanently increase the size of each of the two (2) CloudHub workers by at least four times (4x) to one (1) vCore
B. Use a vertical CloudHub autoscaling policy that triggers on CPU utilization greater than 70%
C. Permanently increase the number of CloudHub workers by four times (4x) to eight (8) CloudHub workers
D. Use a horizontal CloudHub autoscaling policy that triggers on CPU utilization greater than 70%

Answer: D

NEW QUESTION 23
Refer to the exhibit.

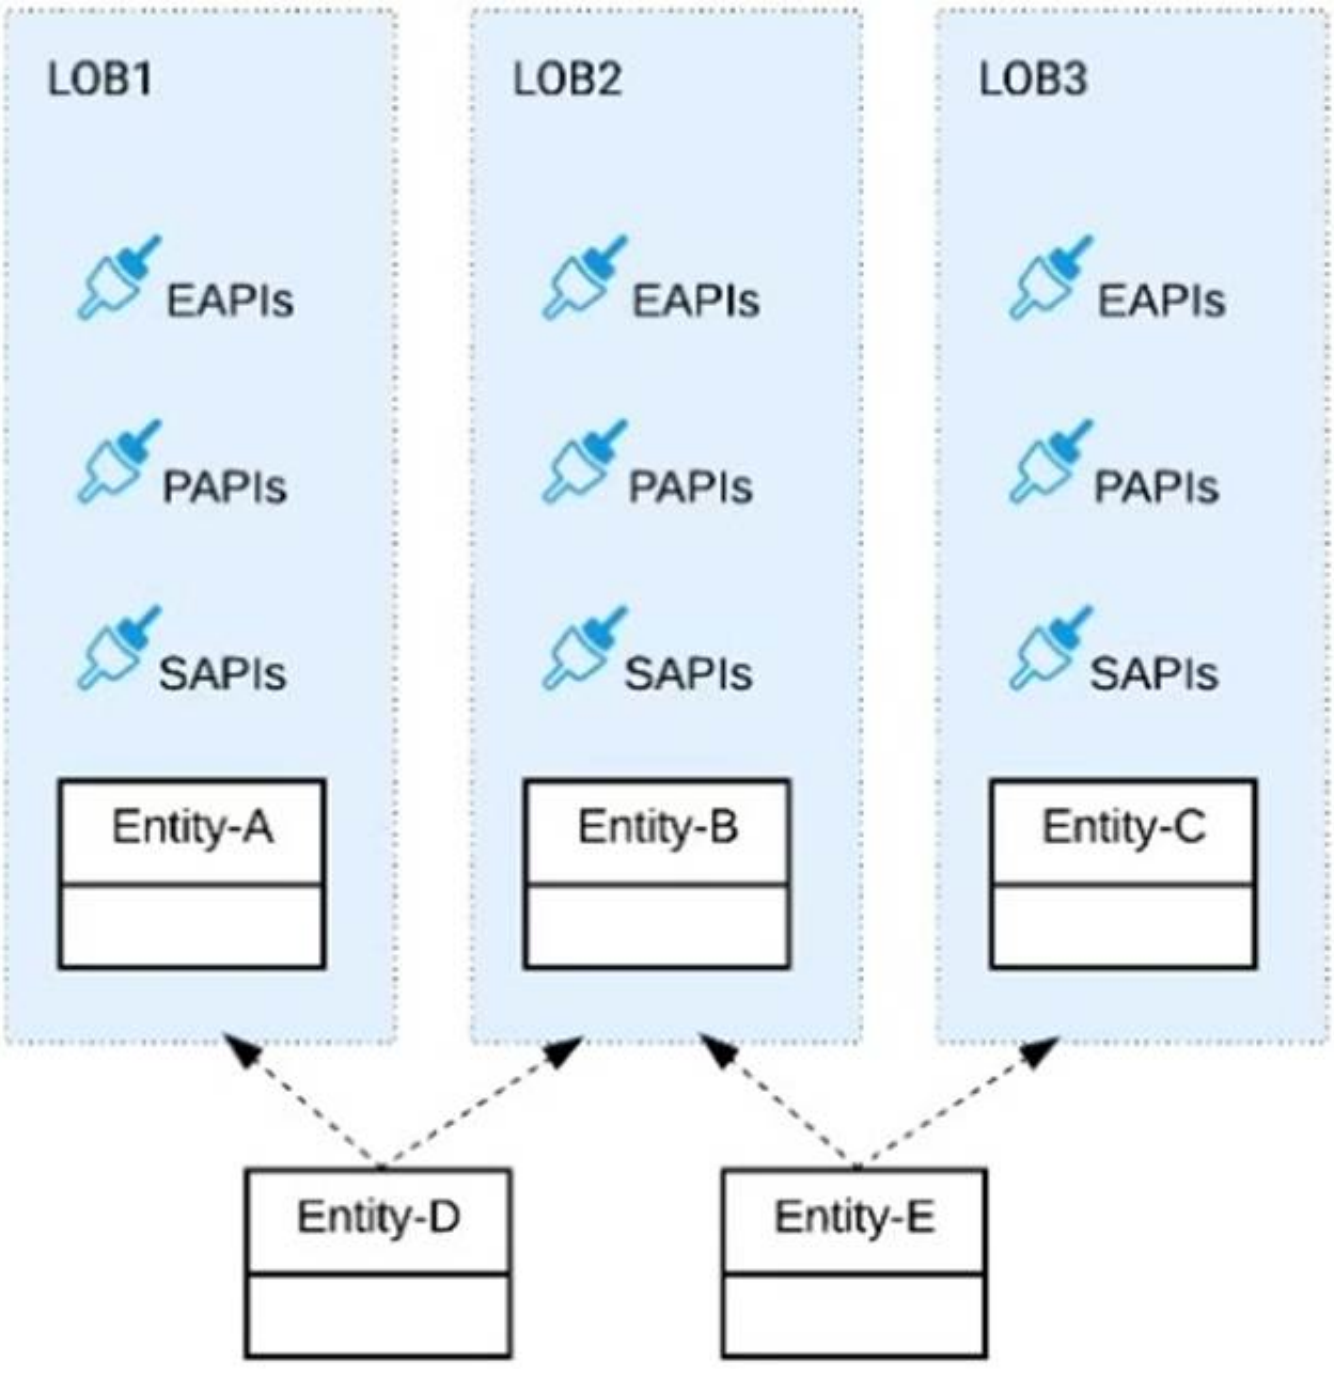

Three business processes need to be implemented, and the implementations need to communicate with several different SaaS applications.
These processes are owned by separate (siloed) LOBs and are mainly independent of each other, but do share a few business entities. Each LOB has one development team 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

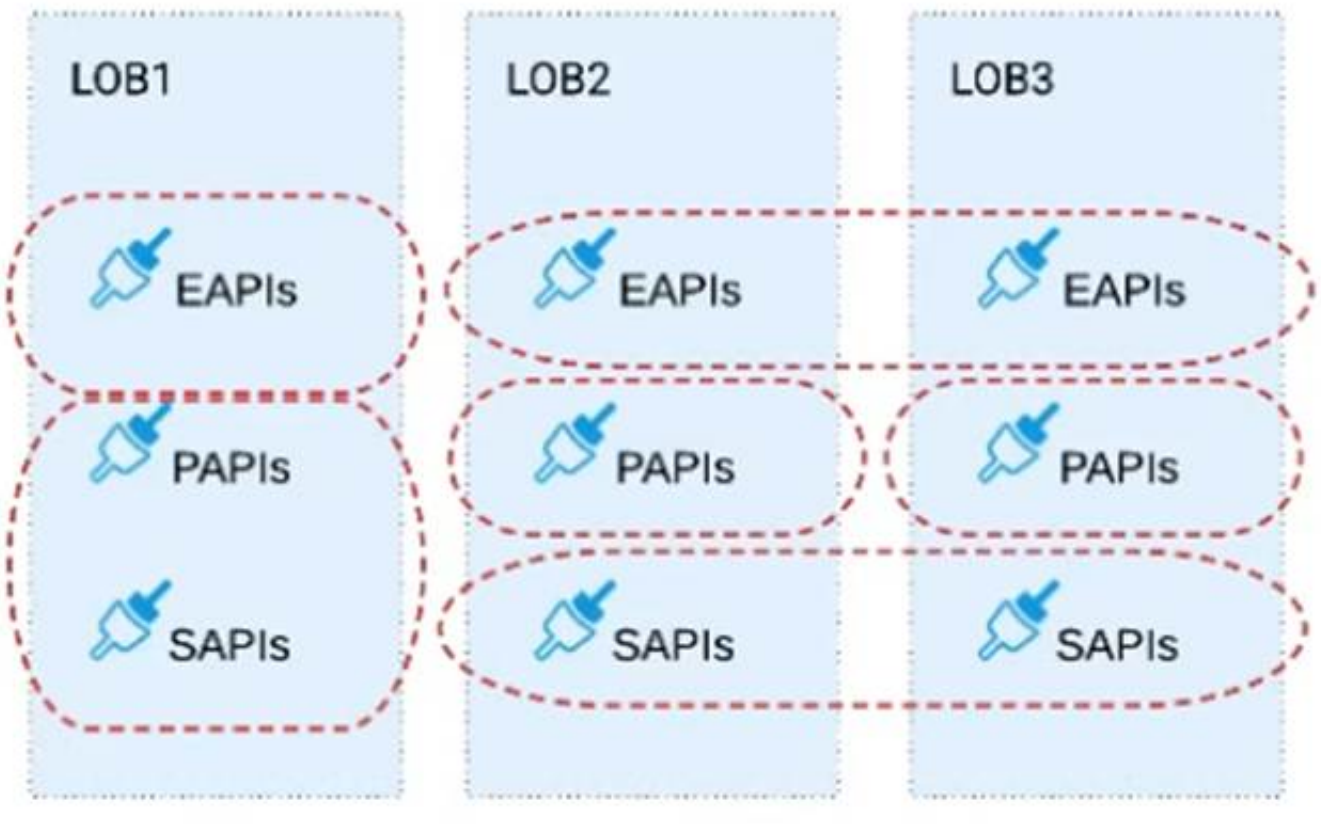

B) Build distinct data models for each API to follow established micro-services and Agile API-centric practices

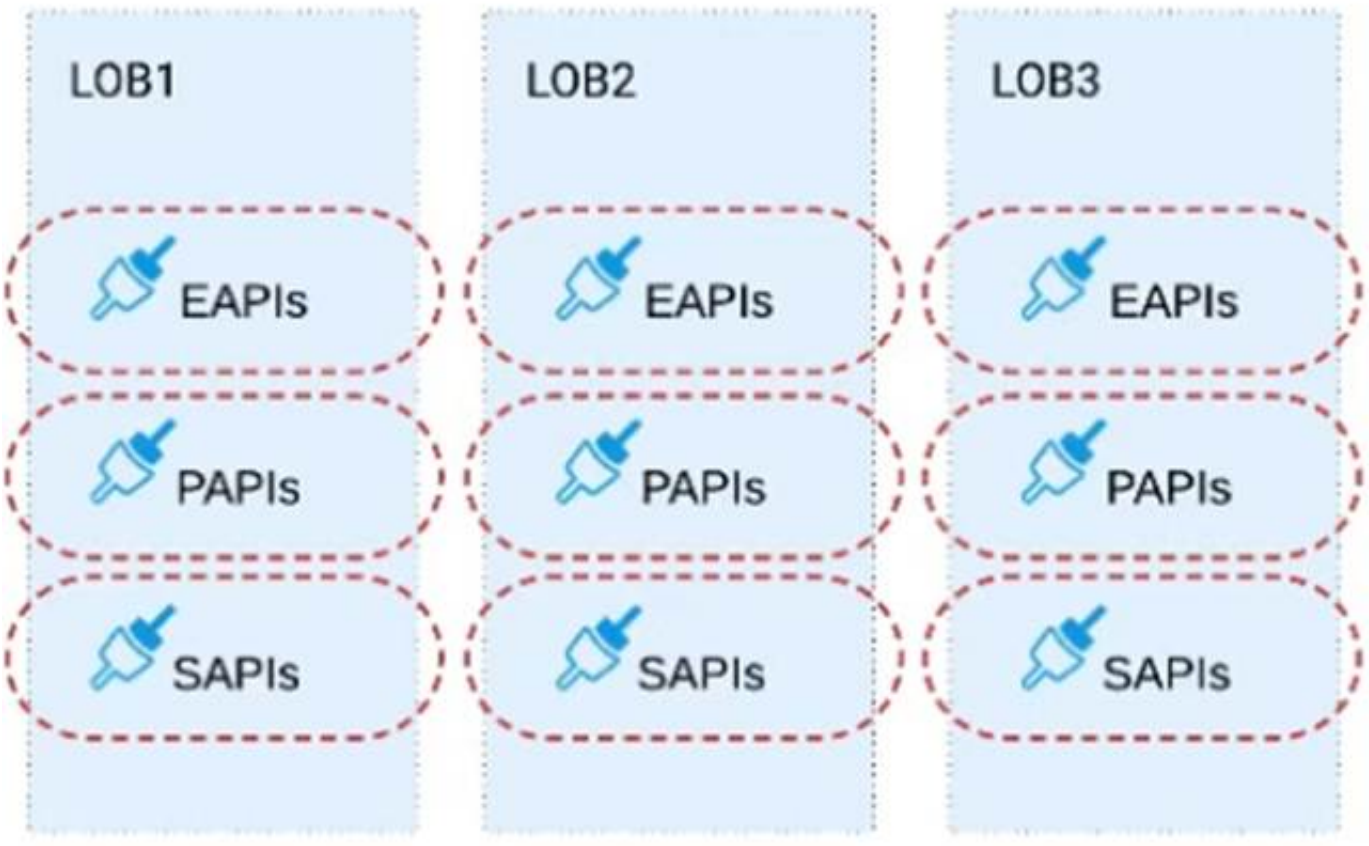

C) Build all API data models using XML schema to drive consistency and reuse across the organization

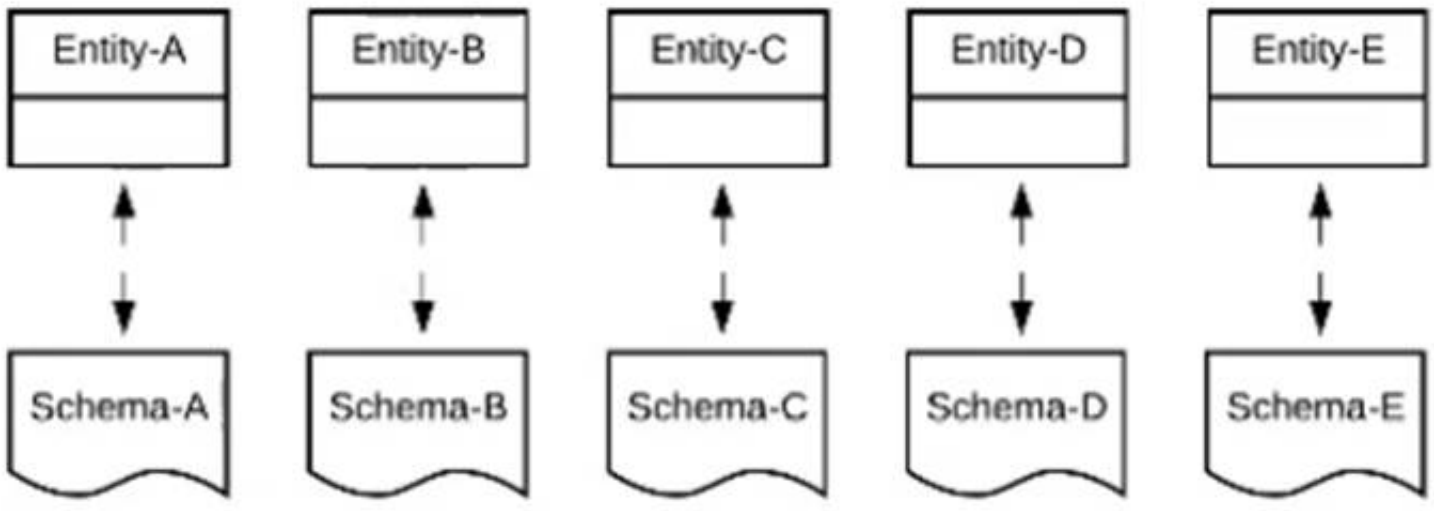

D) Build one centralized Canonical Data Model (Enterprise Data Model) that unifies all the data types from all three business processes, ensuring the data model is consistent and non-redundant

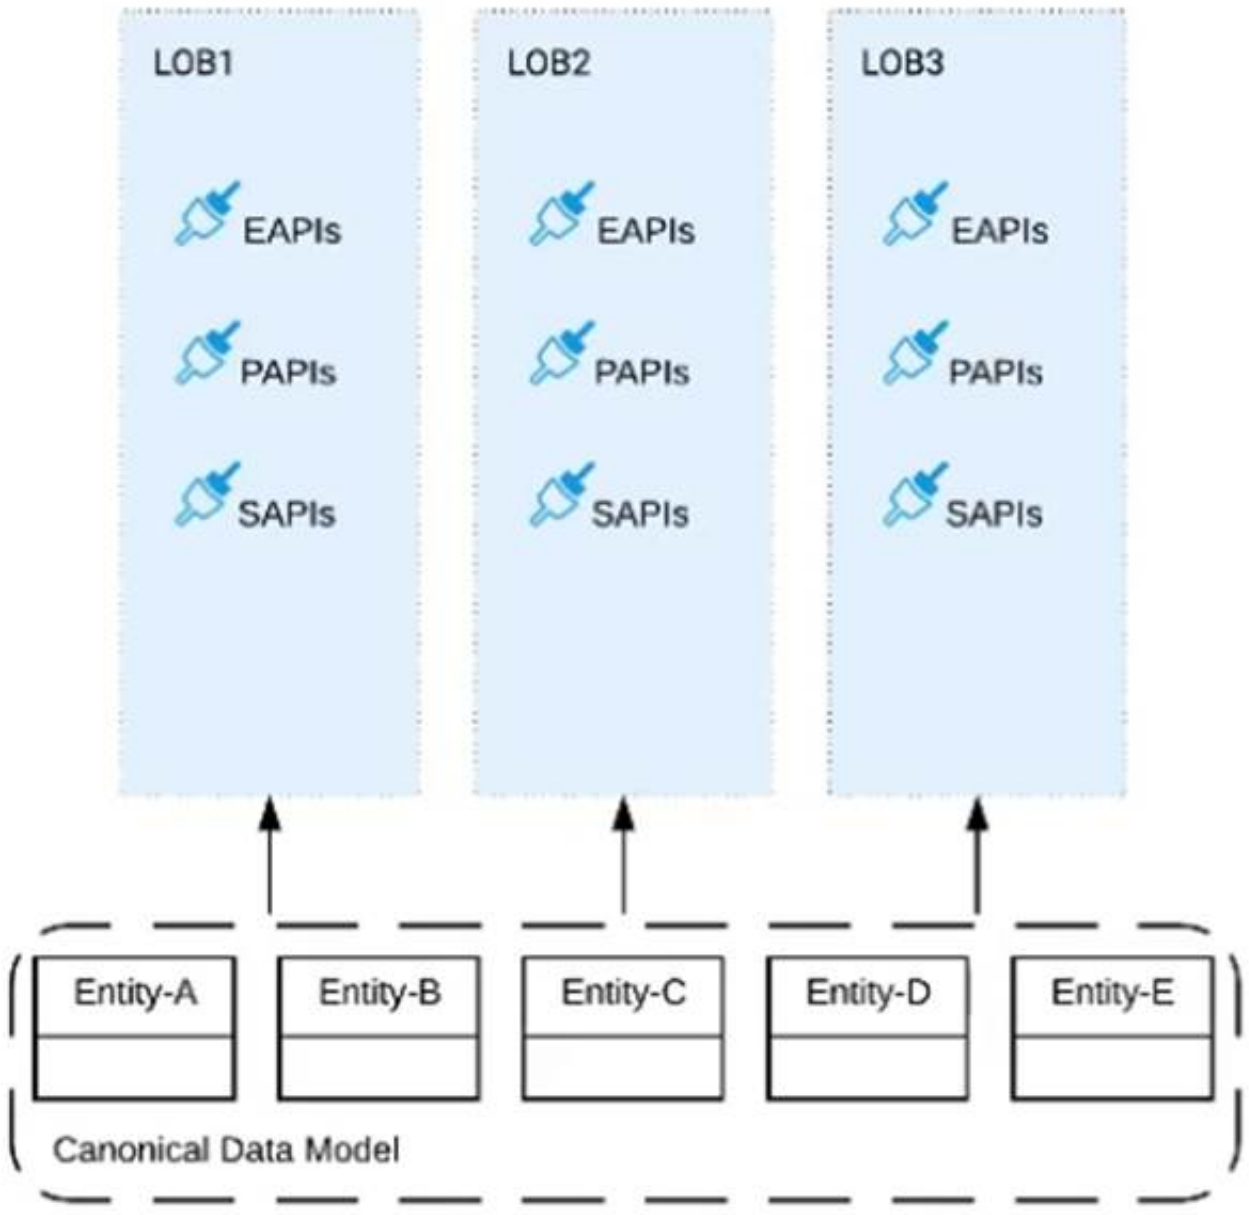

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

NEW QUESTION 27
A developer from the Central IT team has created an initial version of the RAML definition in Design Center for an OAuth 2.0-protected System API and published it to Exchange. Another developer from LoB IT discovered the System API in Exchange and would like to leverage it in the Process API.
What is the MuleSoft-recommended approach for Process API to invoke the System API?

A. The Process API needs to import an CAuth 2.0 module from Exchange first and update it with OAuth 2.0 credentials before the System API can be invoked
B. The Process API uses property YAML files to store the System API URLs and uses the HTTP Request Connector to invoke the Systerm API
C. The Process APL uses the REST Connect Connector autogenerated in Exchange for the System API
D. The Process API manually updates the Process API POM file to include the System API as a dependency

Answer: C

NEW QUESTION 28
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 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 33
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 38
Which component monitors APIs and endpoints at scheduled intervals, receives reports about whether tests pass or fail, and displays statistics about API and endpoint performance?

A. API Analytics
B. Anypoint Monitoring dashboards
C. APT Functional Monitoring
D. Anypoint Runtime Manager alerts

Answer: C

NEW QUESTION 40
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 API's 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 update should be identified as a project risk and full regression testing of the functionality that uses this API should be run
B. The API producer should be contacted to understand the change to existing functionality
C. The API producer should be requested to run the old version in parallel with the new one
D. The API client code ONLY needs to be changed if it needs to take advantage of new features

Answer: D

NEW QUESTION 42
When could the API data model of a System API reasonably mimic the data model exposed by the corresponding backend system, with minimal improvements over the backend system's data model?

A. When there is an existing Enterprise Data Model widely used across the organization
B. When the System API can be assigned to a bounded context with a corresponding data model
C. When a pragmatic approach with only limited isolation from the backend system is deemed appropriate
D. When the corresponding backend system is expected to be replaced in the near future

Answer: C

NEW QUESTION 45
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

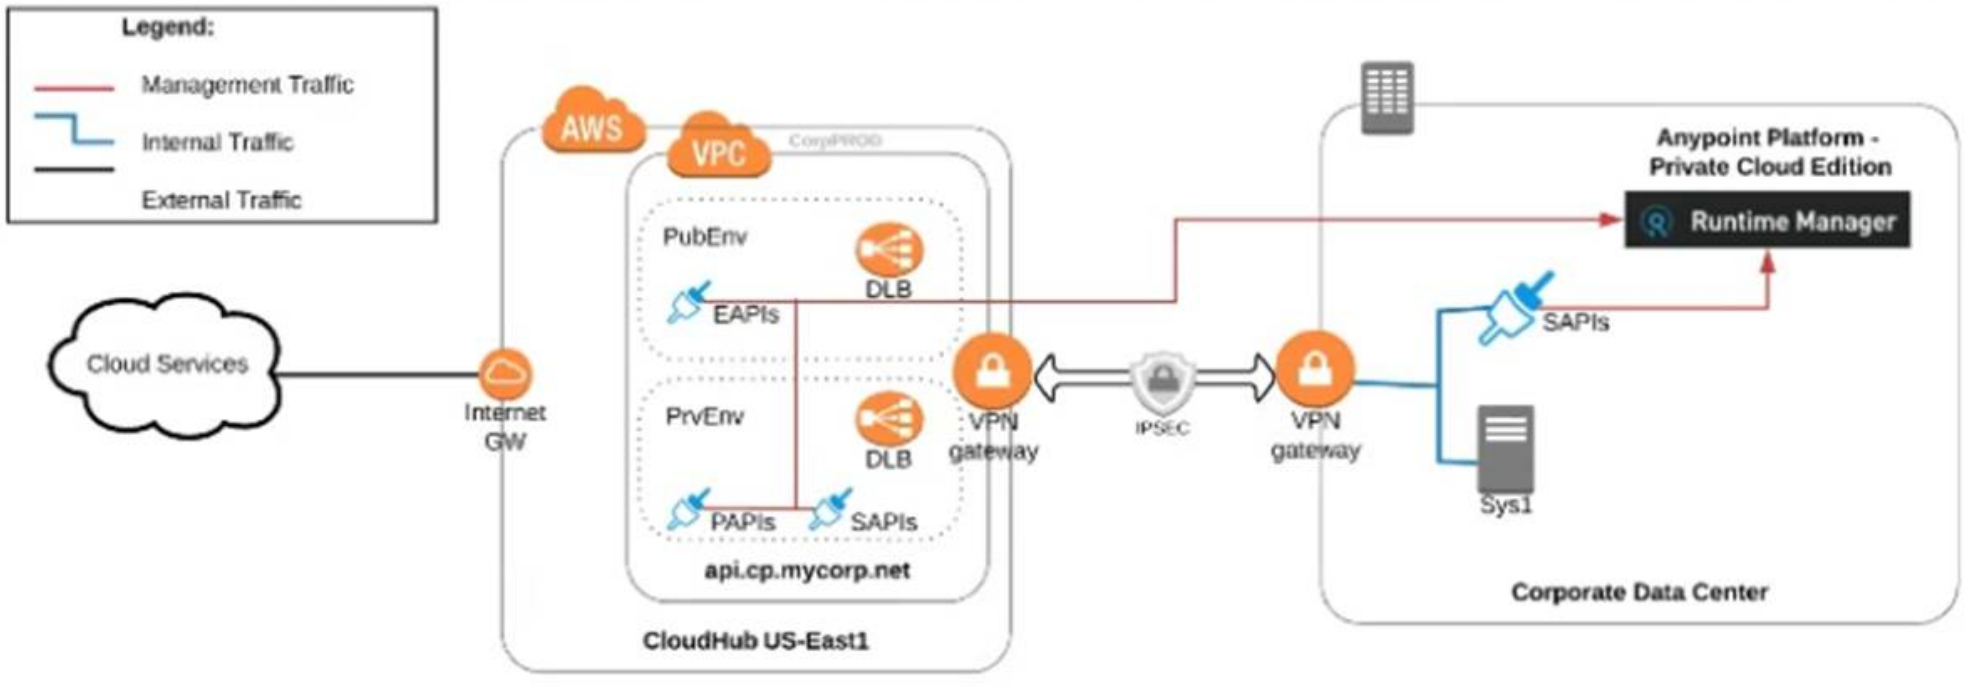

B) Use CloudHub-deployed Mule runtimes in the shared worker cloud managed by the MuleSoft-hosted Anypoint Platform control plane

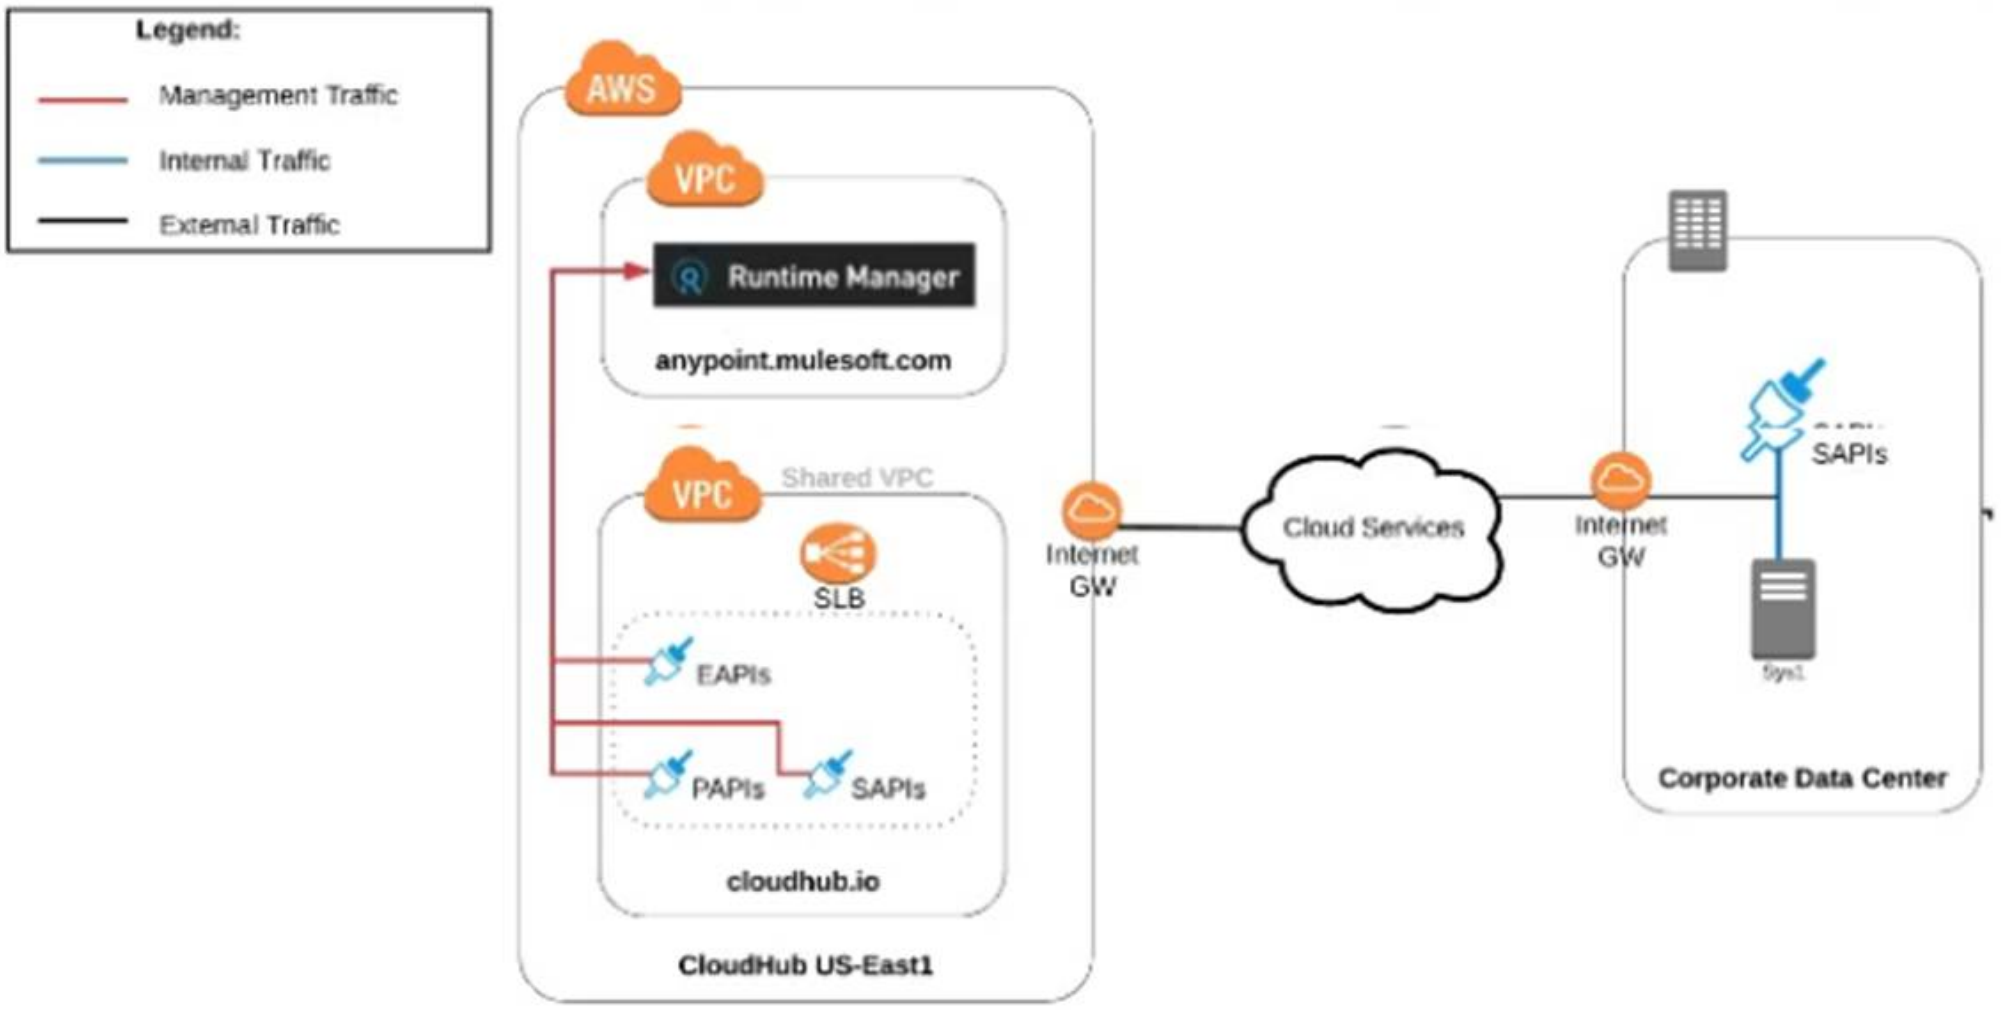

C) Use an on-premises installation of Mule runtimes that are completely isolated with NO external network access, managed by the Anypoint Platform Private Cloud Edition control plane

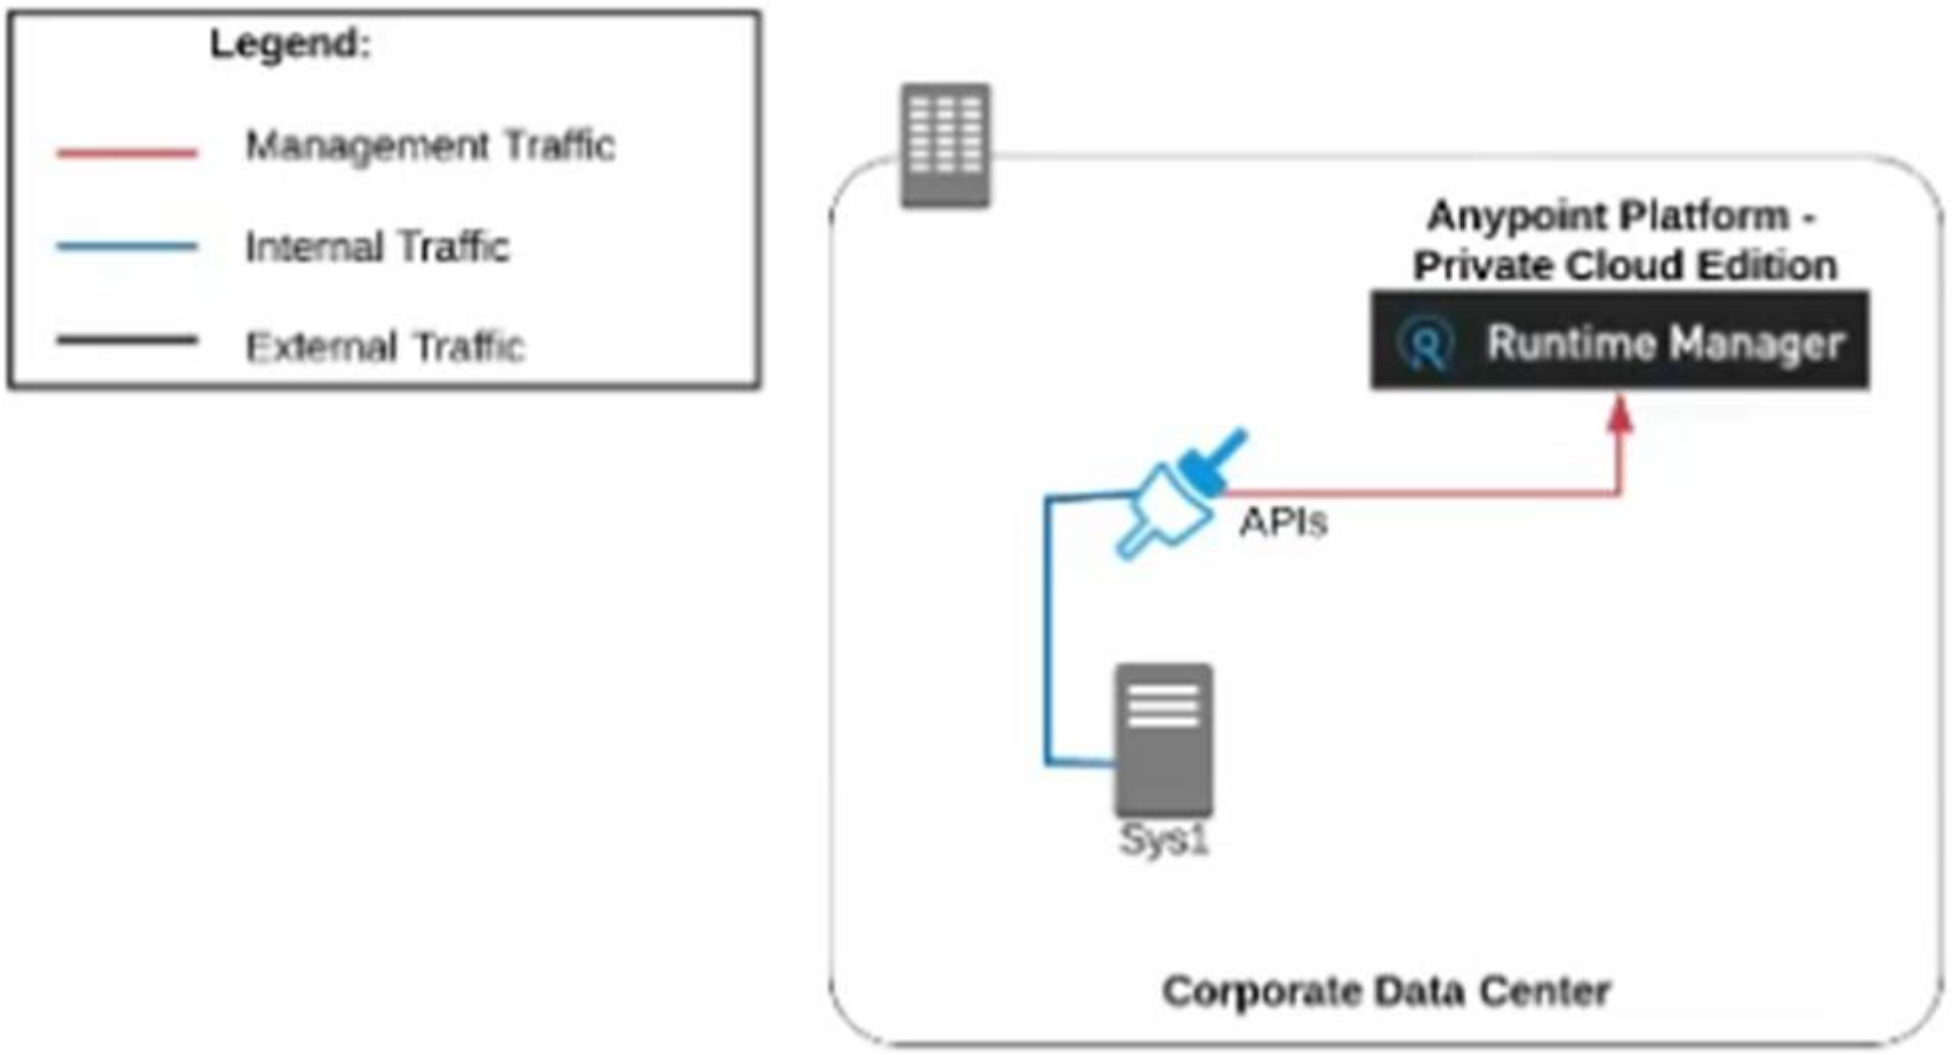

D) Use a combination of Cloud Hub-deployed and manually provisioned on-premises Mule runtimes managed by the MuleSoft-hosted Anypoint Platform control plane

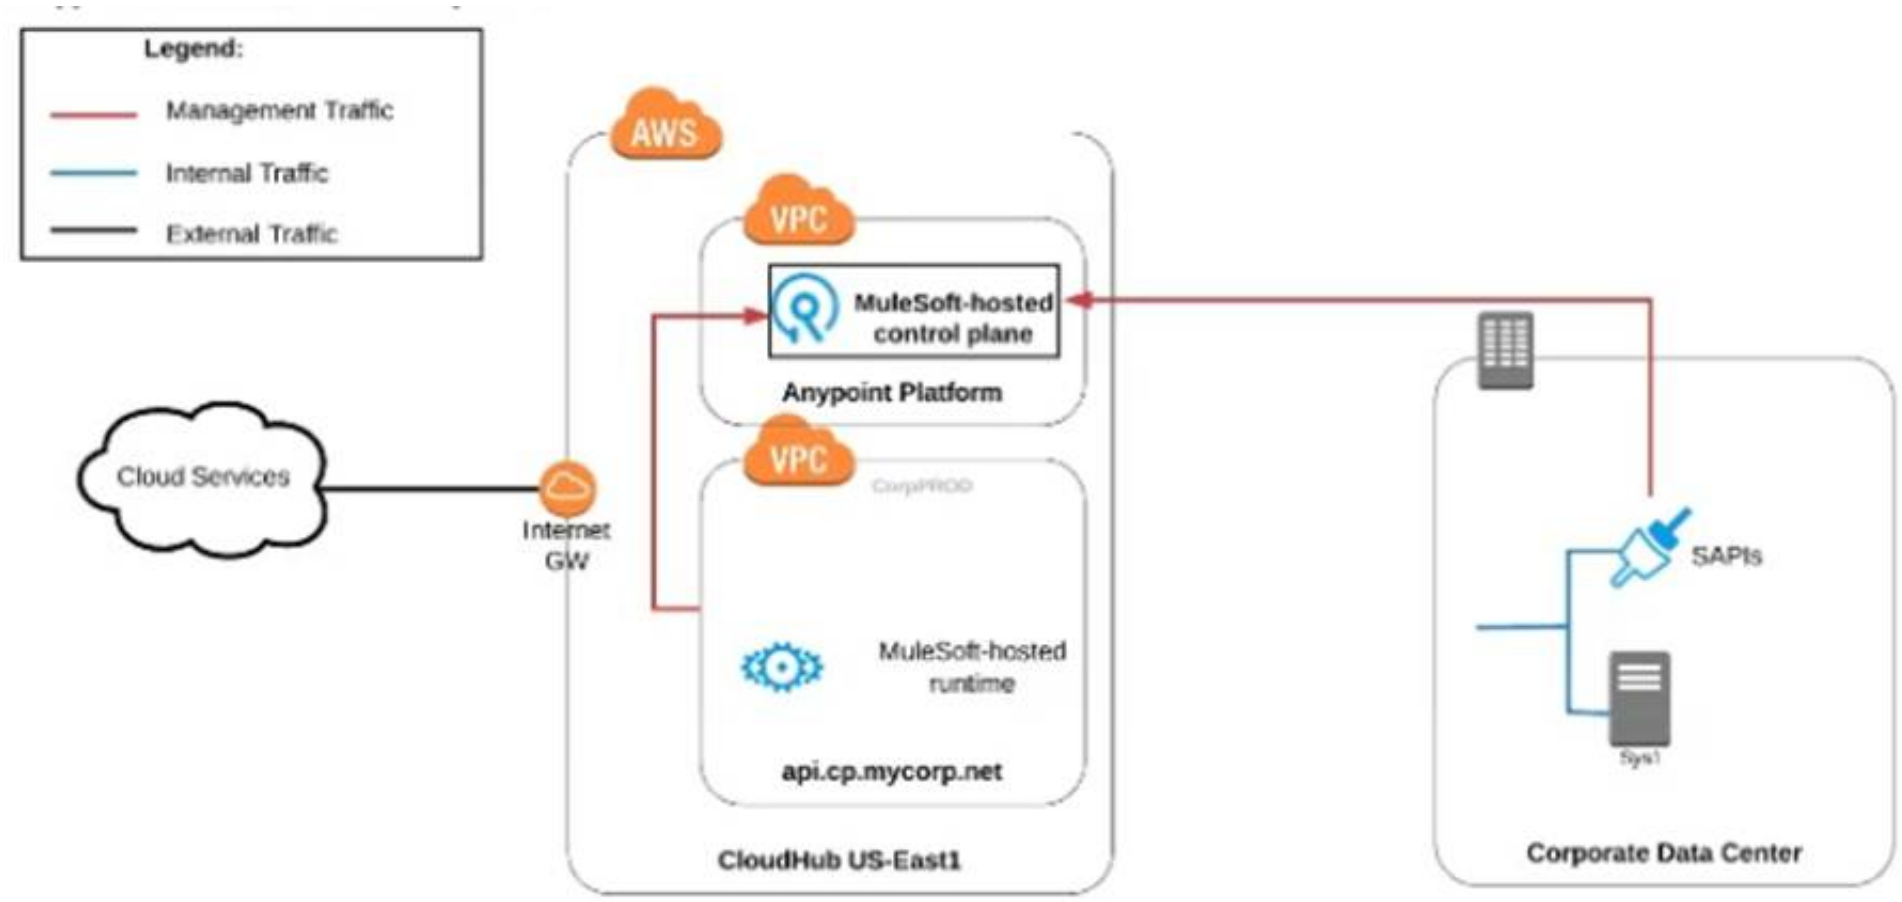

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

NEW QUESTION 49
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 54
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

NEW QUESTION 58
The application network is recomposable: it is built for change because it "bends but does not break"

A. TRUE
B. FALSE

Answer: A

NEW QUESTION 61
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 62
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 67
A TemperatureSensors API instance is defined in API Manager in the PROD environment of the CAR_FACTORY business group. An AcmelemperatureSensors Mule application implements this API instance and is deployed from Runtime Manager to the PROD environment of the CAR_FACTORY business group. A policy that requires a valid client ID and client secret is applied in API Manager to the API instance. Where can an API consumer obtain a valid client ID and client secret to call the
AcmeTemperatureSensors Mule application?

A. In secrets manager, request access to the Shared Secret static username/password
B. In API Manager, from the PROD environment of the CAR_FACTORY business group
C. In access management, from the PROD environment of the CAR_FACTORY business group
D. In Anypoint Exchange, from an API client application that has been approved for the TemperatureSensors API instance

Answer: D

NEW QUESTION 70
An enterprise is embarking on the API-led digital transformation journey, and the central IT team has started to define System APIs. Currently there is no Enterprise Data Model being defined within the enterprise, and the definition of a clean Bounded Context Data Model requires too much effort.
According to MuleSoft's recommended guidelines, how should the System API data model be defined?

A. If there are misspellings of the data fields in the back-end system, Systerm APIs should not correct it, and expose it as-is to mirror the back-end systems
B. The data model of the System APIs should make use of data types that approximately mirror those from the back-end systems
C. The data model should define its own naming convention, and not follow the same naming as the back-end systems
D. The System APIs should expose all back-end system fields

Answer: B

NEW QUESTION 71
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 73
What is typically NOT a function of the APIs created within the framework called API-led connectivity?

A. They provide an additional layer of resilience on top of the underlying backend system, thereby insulating clients from extended failure of these systems.
B. They allow for innovation at the user Interface level by consuming the underlying assets without being aware of how data Is being extracted from backend systems.
C. They reduce the dependency on the underlying backend systems by helping unlock data from backend systems In a reusable and consumable way.
D. They can compose data from various sources and combine them with orchestration logic to create higher level value.

Answer: A

NEW QUESTION 78
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 81
What is a key performance indicator (KPI) that measures the success of a typical C4E that is immediately apparent in responses from the Anypoint Platform APIs?

A. The number of production outage incidents reported in the last 24 hours
B. The number of API implementations that have a publicly accessible HTTP endpoint and are being managed by Anypoint Platform
C. The fraction of API implementations deployed manually relative to those deployed using a CI/CD tool
D. The number of API specifications in RAML or OAS format published to Anypoint Exchange

Answer: D

NEW QUESTION 83
True or False. We should always make sure that the APIs being designed and developed are self-servable even if it needs more man-day effort and resources.

A. FALSE
B. TRUE

Answer: B

NEW QUESTION 85
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 Design and Development and API Runtime Execution and Hosting
>> 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

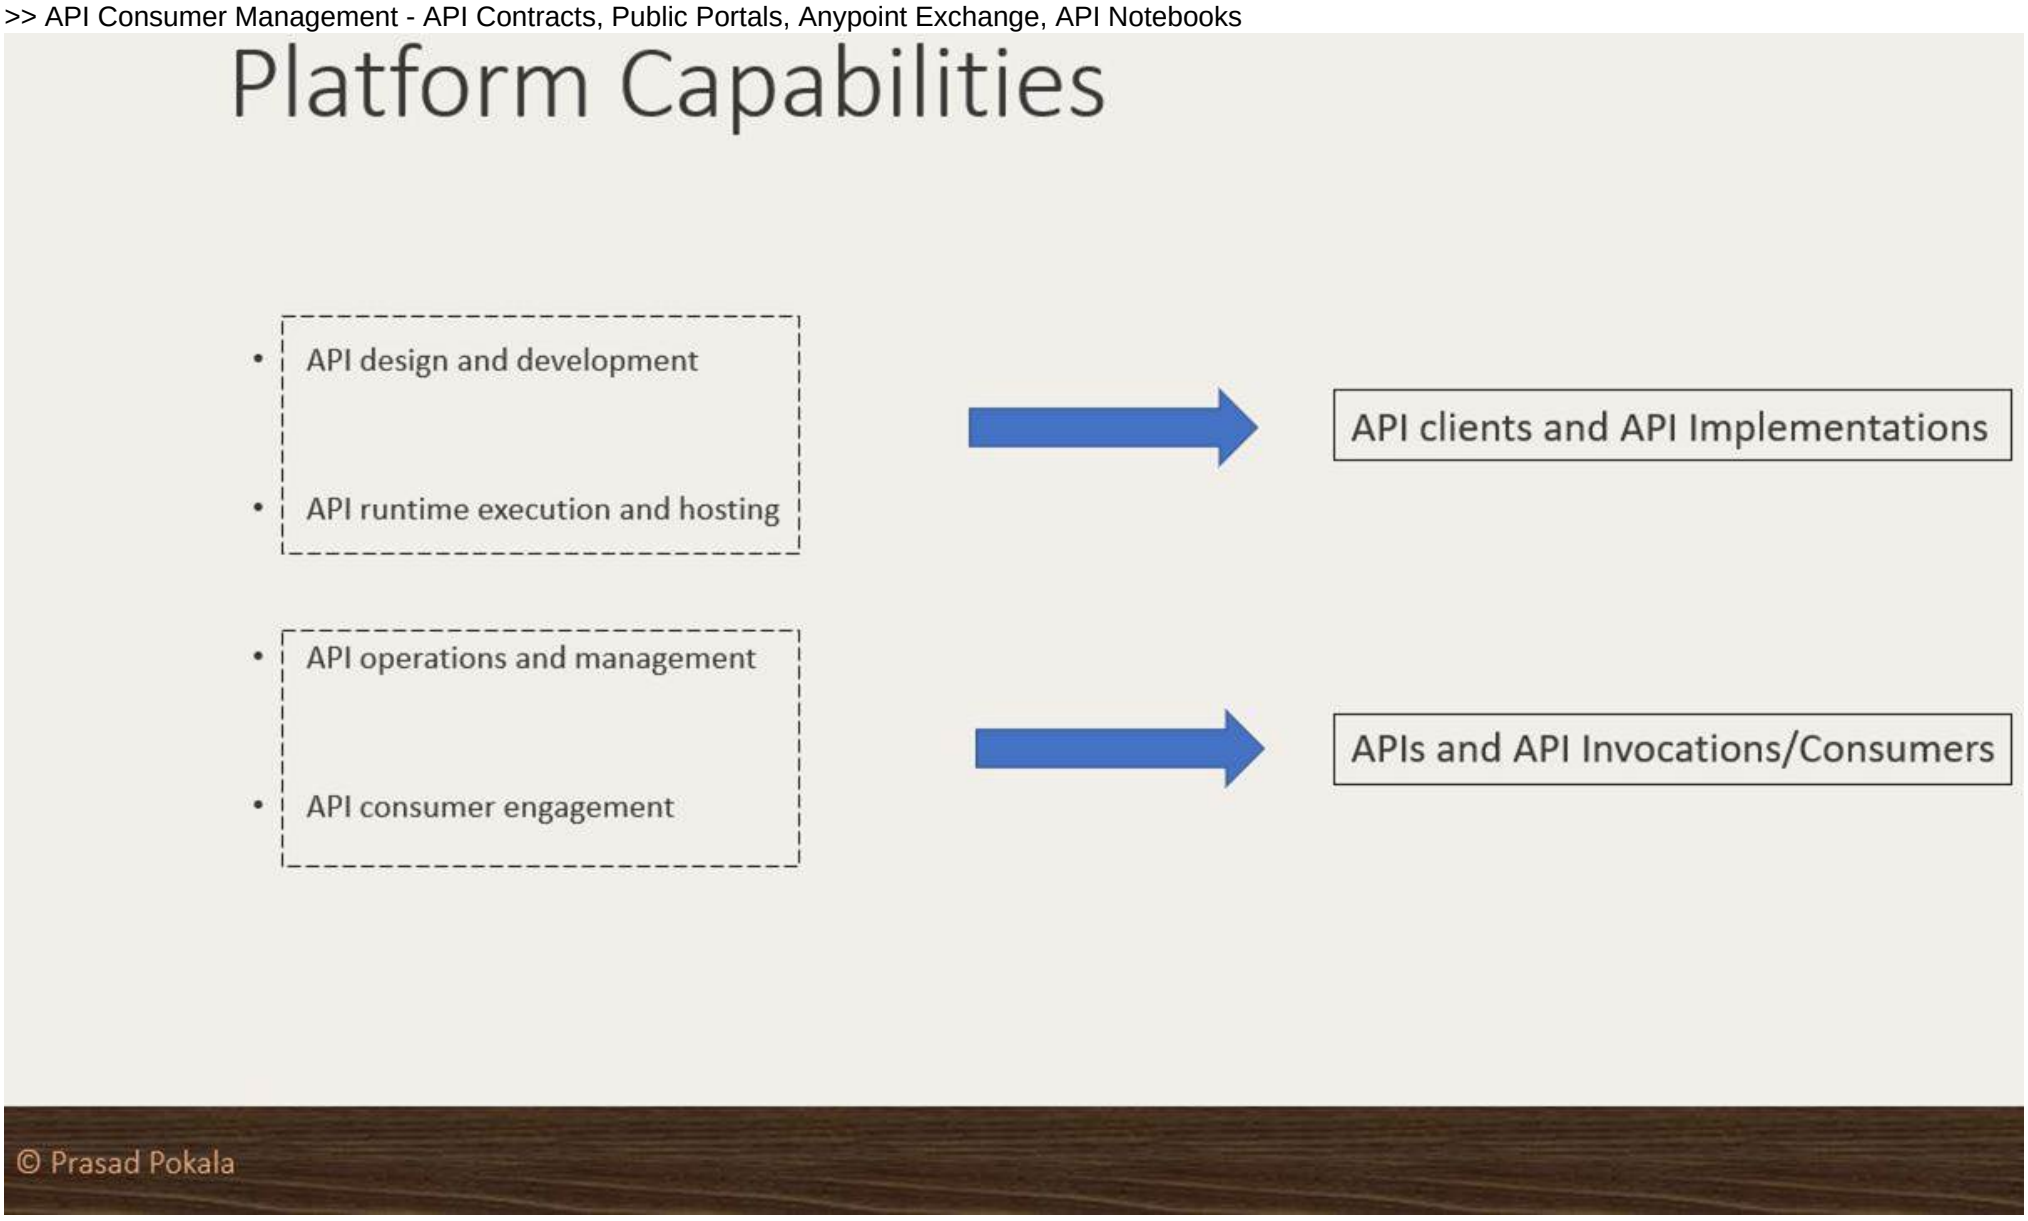

NEW QUESTION 86
Which statement is true about Spike Control policy and Rate Limiting policy?

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 87
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

NEW QUESTION 91
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 92
What condition requires using a CloudHub Dedicated Load Balancer?

A. When cross-region load balancing is required between separate deployments of the same Mule application
B. When custom DNS names are required for API implementations deployed to customer- hosted Mule runtimes
C. When API invocations across multiple CloudHub workers must be load balanced
D. When server-side load-balanced TLS mutual authentication is required between API implementations and API clients

Answer: D

NEW QUESTION 97
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 101
An Anypoint Platform organization has been configured with an external identity provider (IdP) for identity management and client management. What credentials or token must be provided to Anypoint CLI to execute commands against the Anypoint Platform APIs?

A. The credentials provided by the IdP for identity management
B. The credentials provided by the IdP for client management
C. An OAuth 2.0 token generated using the credentials provided by the IdP for client management
D. An OAuth 2.0 token generated using the credentials provided by the IdP for identity management

Answer: A

NEW QUESTION 103
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

NEW QUESTION 106
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 111
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 114
A company stores financial transaction data in two legacy systems. For each legacy system, a separate, dedicated System API (SAPI) exposes data for that legacy system. A Process API (PAPI) merges the data retrieved from ail of the System APIs into a common format. Several API clients call the PAPI through its public domain name.
The company now wants to expose a subset of financial data to a newly developed mobile application that uses a different Bounded Context Data Model. The company wants to follow MuleSoft's best practices for building out an effective application network.
Following MuleSoft's best practices, how can the company expose financial data needed by the mobile application in a way that minimizes the impact on the currently running API clients, API implementations, and support asset reuse?

A. Add two new Experience APIs (EAPI-i and EAPI-2}.Add Mobile PAPI-2 to expose the Intended subset of financial data as requeste
B. Both PAPIs access the Legacy Systems via SAPI-1 and SAP]-2.
C. Add two new Experience APIs (EAPI-i and EAPI-2}.Add Mobile PAPI-2 to expose the Intended subset of financial data as requeste
D. Both PAPIs access the Legacy Systems via SAPI-1 and SAP]-2.
E. Create a new mobile Experince API (EAPI) chat exposes that subset of PAPI endpoint
F. Add transformtion login to the mobile Experince API implementation to make mobile data compatible with the required PAPIs.

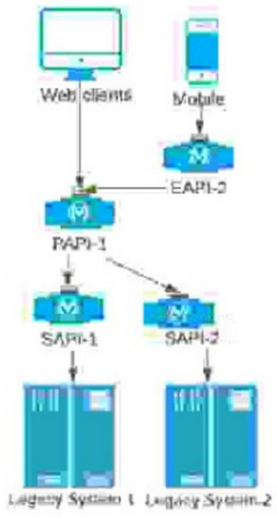

G. Develop and deploy is new PAPI implementation with data transformation and ... login to support this required endpoints of both mobile and web clients.Deploy an API Proxy with an endpoint from API Manager that redirect the existing PAPI endpoints to the new PAPI.

Answer: A

NEW QUESTION 115
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 117
An eCommerce company is adding a new Product Details feature to their website, A customer will launch the product catalog page, a new Product Details link will appear by product where they can click to retrieve the product detail description. Product detail data is updated with product update releases, once or twice a year, Presently the database response time has been very slow due to high volume.
What action retrieves the product details with the lowest response time, fault tolerant, and consistent data?

A. Select the product details from a database in a Cache scope and return them within the API response
B. Select the product details from a database and put them in Anypoint MQ; the Anypoint MO subseriber will receive the product details and return them within the API response
C. Use an object store to store and retrieve the product details originally read from a database and return them within the API response
D. Select the product details from a database and return them within the API response

Answer: C

NEW QUESTION 121
An existing Quoting API is defined in RAML and used by REST clients for interacting with the quoting engine. Currently there is a resource defined in the RAML that allows the creation of quotes; however, a new requirement was just received to allow for the updating of existing quotes.
Which two actions need to be taken to facilitate this change so it can be processed? Choose 2 answers

A. Update the API implementation to accommodate the new update request
B. Remove the old client applications and create new client applications to account for the changes
C. Update the RAML with new method details for the update request
D. Deprecate existing versions of the API in Exchange
E. Add a new API policy to API Manager to allow access to the updated endpoint

Answer: AC

NEW QUESTION 125
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 128
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 130
An API experiences a high rate of client requests (TPS) vwth small message paytoads. How can usage limits be imposed on the API based on the type of client application?

A. Use an SLA-based rate limiting policy and assign a client application to a matching SLA tier based on its type
B. Use a spike control policy that limits the number of requests for each client application type
C. Use a cross-origin resource sharing (CORS) policy to limit resource sharing between client applications, configured by the client application type
D. Use a rate limiting policy and a client ID enforcement policy, each configured by the client application type

Answer: A

NEW QUESTION 132
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 135
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 140
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 142
What Anypoint Connectors support transactions?

A. Database, JMS, VM
B. Database, 3MS, HTTP
C. Database, JMS, VM, SFTP
D. Database, VM, File

Answer: A

NEW QUESTION 146
How can the application of a rate limiting API policy be accurately reflected in the RAML definition of an API?

A. By refining the resource definitions by adding a description of the rate limiting policy behavior
B. By refining the request definitions by adding a remaining Requests query parameter with description, type, and example
C. By refining the response definitions by adding the out-of-the-box Anypoint Platform rate- limit-enforcement securityScheme with description, type, and example
D. By refining the response definitions by adding the x-ratelimit-* response headers with description, type, and example

Answer: D

Explanation:
Correct Answer: By refining the response definitions by adding the x-ratelimit-* response headers with description, type, and example

Response Headers

The following access-limiting policies return headers having information about the current state of the request:

- X-Ratelimit-Remaining: The amount of available quota.
- X-Ratelimit-Limit: The maximum available requests per window.
- X-Ratelimit-Reset: The remaining time, in milliseconds, until a new window starts.

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.

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

NEW QUESTION 151
An Order API triggers a sequence of other API calls to look up details of an order's items in a back-end inventory database. The Order API calls the OrderItems process API, which calls the Inventory system API. The Inventory system API performs database operations in the back-end inventory database.
The network connection between the Inventory system API and the database is known to
be unreliable and hang at unpredictable times.
Where should a two-second timeout be configured in the API processing sequence so that the Order API never waits more than two seconds for a response from the OrderItems process API?

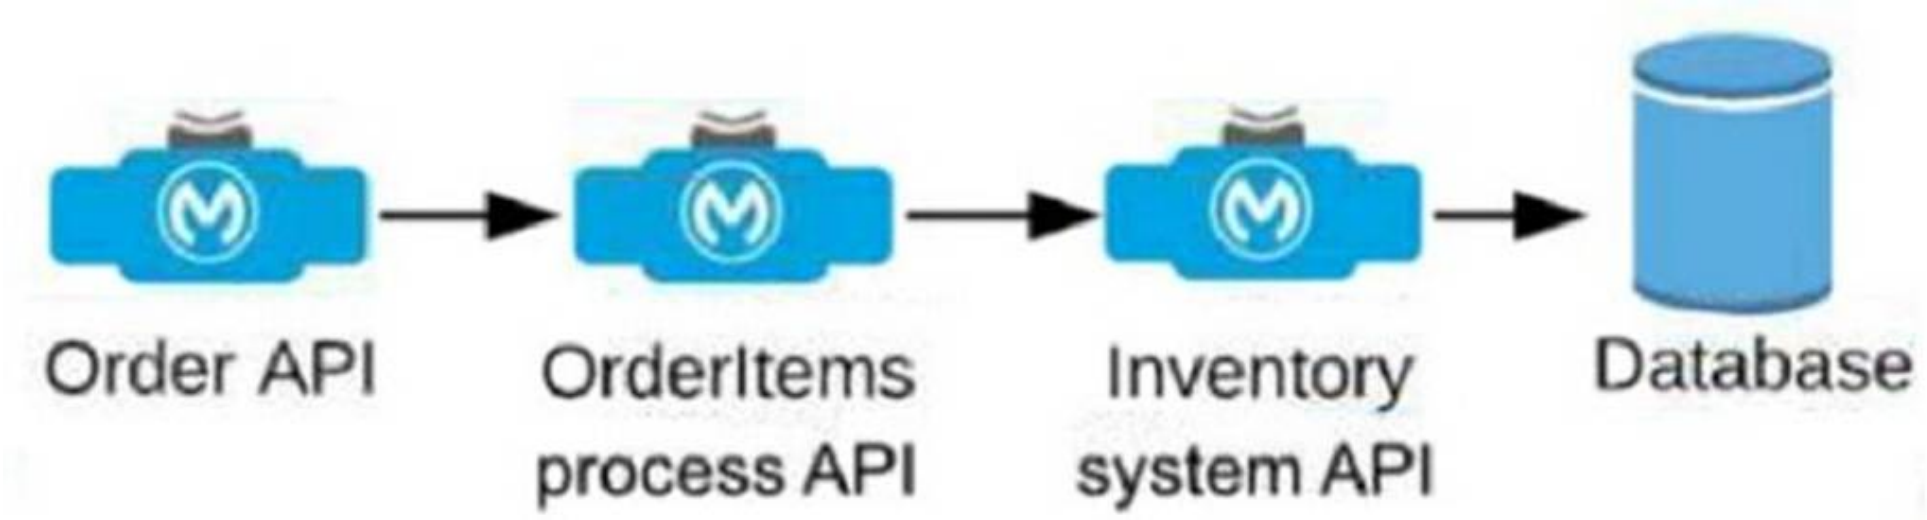

A. In the OrderItems process API implementation
B. In the Order API implementation
C. In the Inventory system API implementation
D. In the inventory database

Answer: A

NEW QUESTION 156
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 160

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 162
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 167
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 172
.....

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