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
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 3
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 4
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 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 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 8
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 9
Which APIs can be used with DataGraph to create a unified schema?

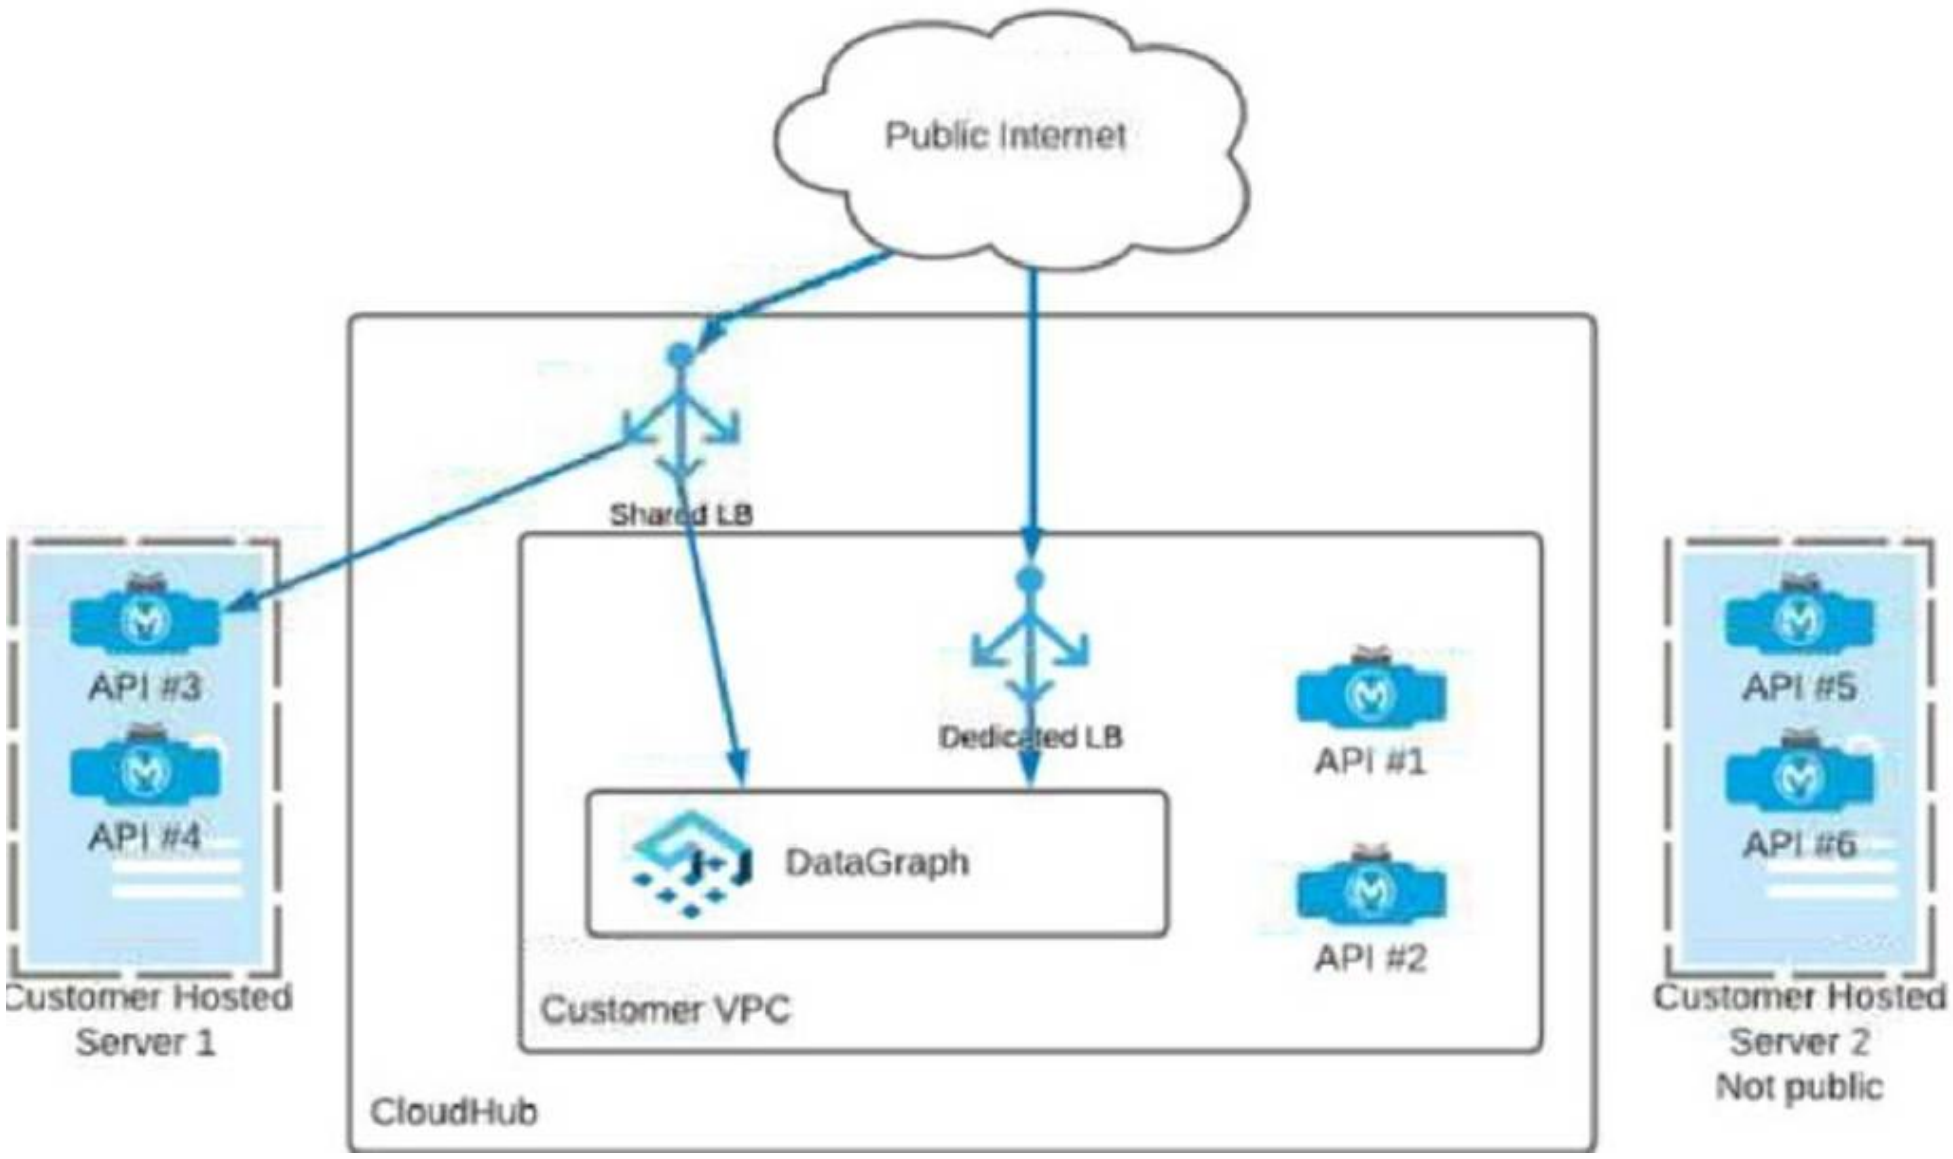

A. APIs 1, 3, 5
B. APIs 2, 4 ,6
C. APIs 1, 2, s5, 6
D. APIs 1, 2, 3, 4

Answer: D

NEW QUESTION 10
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 10
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 12
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 17
Refer to the exhibit.

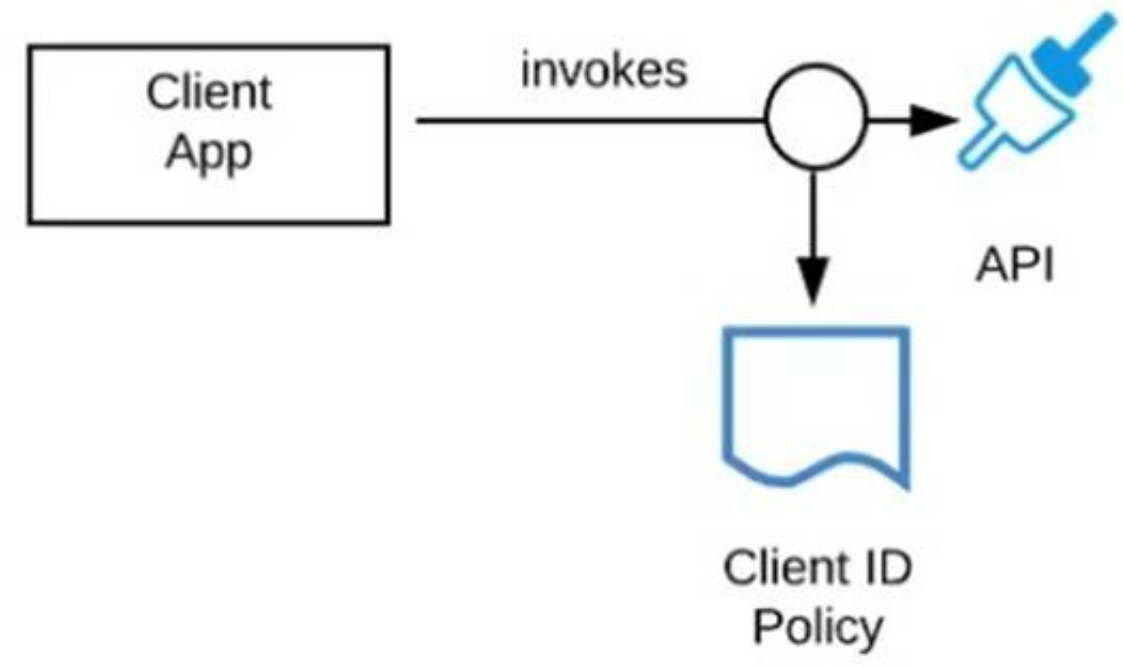

A developer is building a client application to invoke an API deployed to the STAGING environment that is governed by a client ID enforcement policy.
What is required to successfully invoke the API?

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 22
What is true about the technology architecture of Anypoint VPCs?

A. The private IP address range of an Anypoint VPC is automatically chosen by CloudHub
B. Traffic between Mule applications deployed to an Anypoint VPC and on-premises systems can stay within a private network
C. Each CloudHub environment requires a separate Anypoint VPC
D. VPC peering can be used to link the underlying AWS VPC to an on-premises (non AWS) private network

Answer: B

Explanation:
Correct Answer: Traffic between Mule applications deployed to an Anypoint VPC and onpremises systems can stay within a private network
****
>> The private IP address range of an Anypoint VPC is NOT automatically chosen by CloudHub. It is chosen by us at the time of creating VPC using thr CIDR blocks.
CIDR Block: The size of the Anypoint VPC in Classless Inter-Domain Routing (CIDR) notation.
For example, if you set it to 10.111.0.0/24, the Anypoint VPC is granted 256 IP addresses from 10.111.0.0 to 10.111.0.255.
Ideally, the CIDR Blocks you choose for the Anypoint VPC come from a private IP space, and should not overlap with any other Anypoint VPC??s CIDR Blocks, or any CIDR Blocks in use in your corporate network.

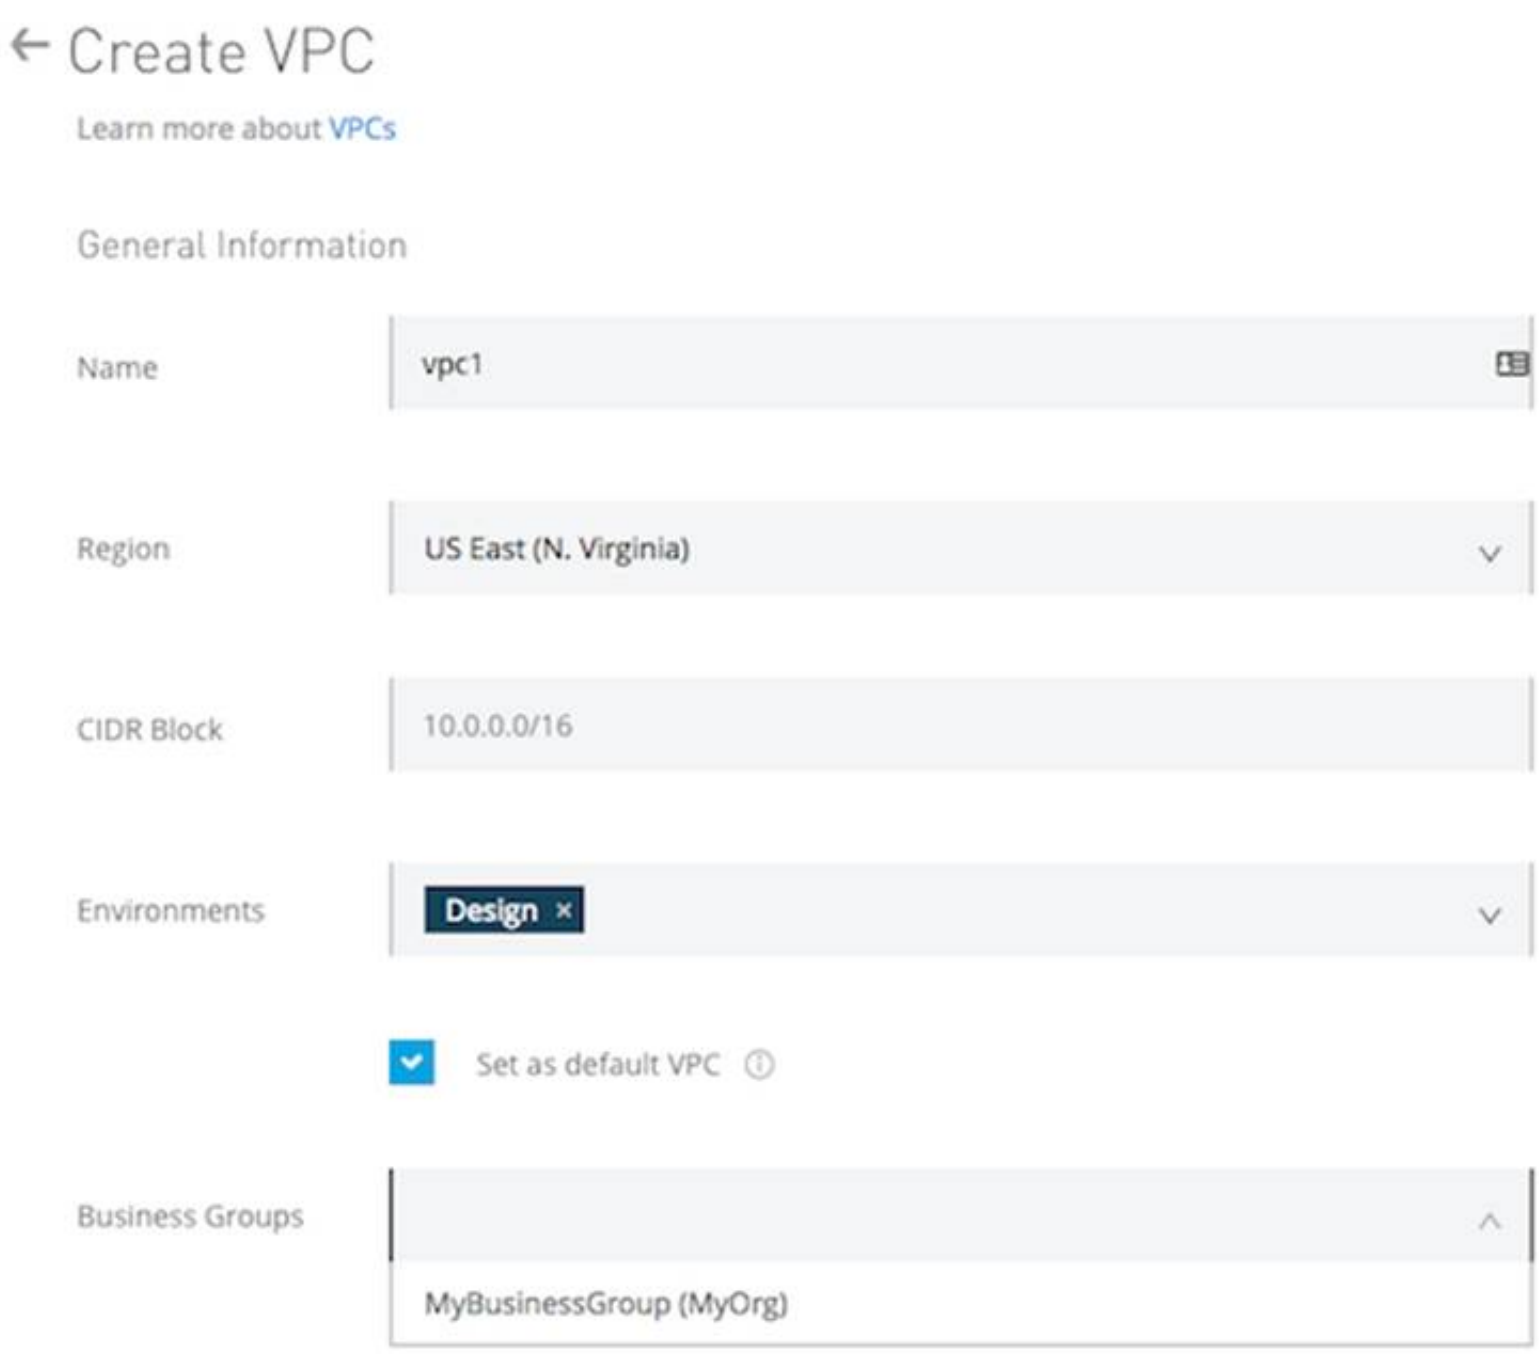

that each CloudHub environment requires a separate Anypoint VPC. Once an Anypoint VPC is created, we can choose a same VPC by multiple environments. However, it is generally a best and recommended practice to always have seperate Anypoint VPCs for Non-Prod and Prod environments.
>> We use Anypoint VPN to link the underlying AWS VPC to an on-premises (non AWS) private network. NOT VPC Peering.
Reference: https://docs.mulesoft.com/runtime-manager/vpn-about
Only true statement in the given choices is that the traffic between Mule applications deployed to an Anypoint VPC and on-premises systems can stay within a private network.
https://docs.mulesoft.com/runtime-manager/vpc-connectivity-methods-concept

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 29
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 31
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 33
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 is using an on-prem cluster in the data center as a runtime plane and MuleSoft- hosted control plane.
How can the company monitor the detailed performance metrics on the Mule applications deployed to the cluster from the control plane?

A. The settings of the Monitoring section in the control plane must be updated to enable detailed logging on the metrics to be captured
B. Monitoring Agent must be installed on each node in the cluster
C. Due to the potential performance impact on the runtime nodes, the Monitoring agent should be installed on a separate server
D. There is no action needed as the on-prem runtime automatically sends the performance data to the control plane

Answer: B

NEW QUESTION 38
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 39
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 40
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 44
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:

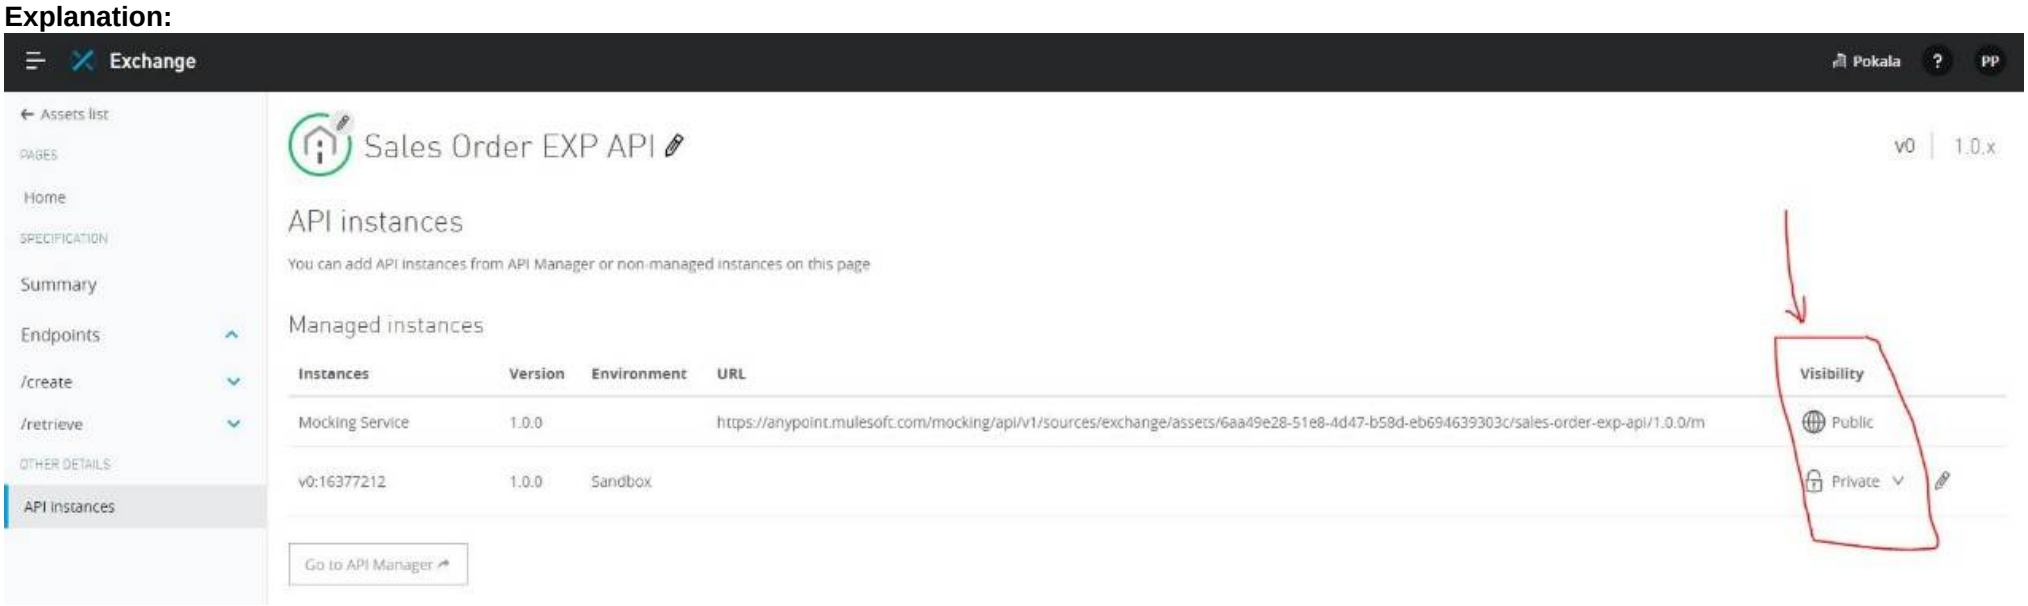

Correct Answer: The visibility level of the API instances of that API that need to be publicly accessible should be set to public visibility.
****
Reference: https://docs.mulesoft.com/exchange/to-share-api-asset-to-portal
https://docs.mulesoft.com/exchange/to-share-api-asset-to-portal

NEW QUESTION 48
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 49
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 54
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 58
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 61
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 65
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 67
An organization has several APIs that accept JSON data over HTTP POST. The APIs are all publicly available and are associated with several mobile applications and web applications.
The organization does NOT want to use any authentication or compliance policies for these APIs, but at the same time, is worried that some bad actor could send payloads that could somehow compromise the applications or servers running the API implementations.
What out-of-the-box Anypoint Platform policy can address exposure to this threat?

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

Answer: D

NEW QUESTION 71
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 74
Refer to the exhibits.

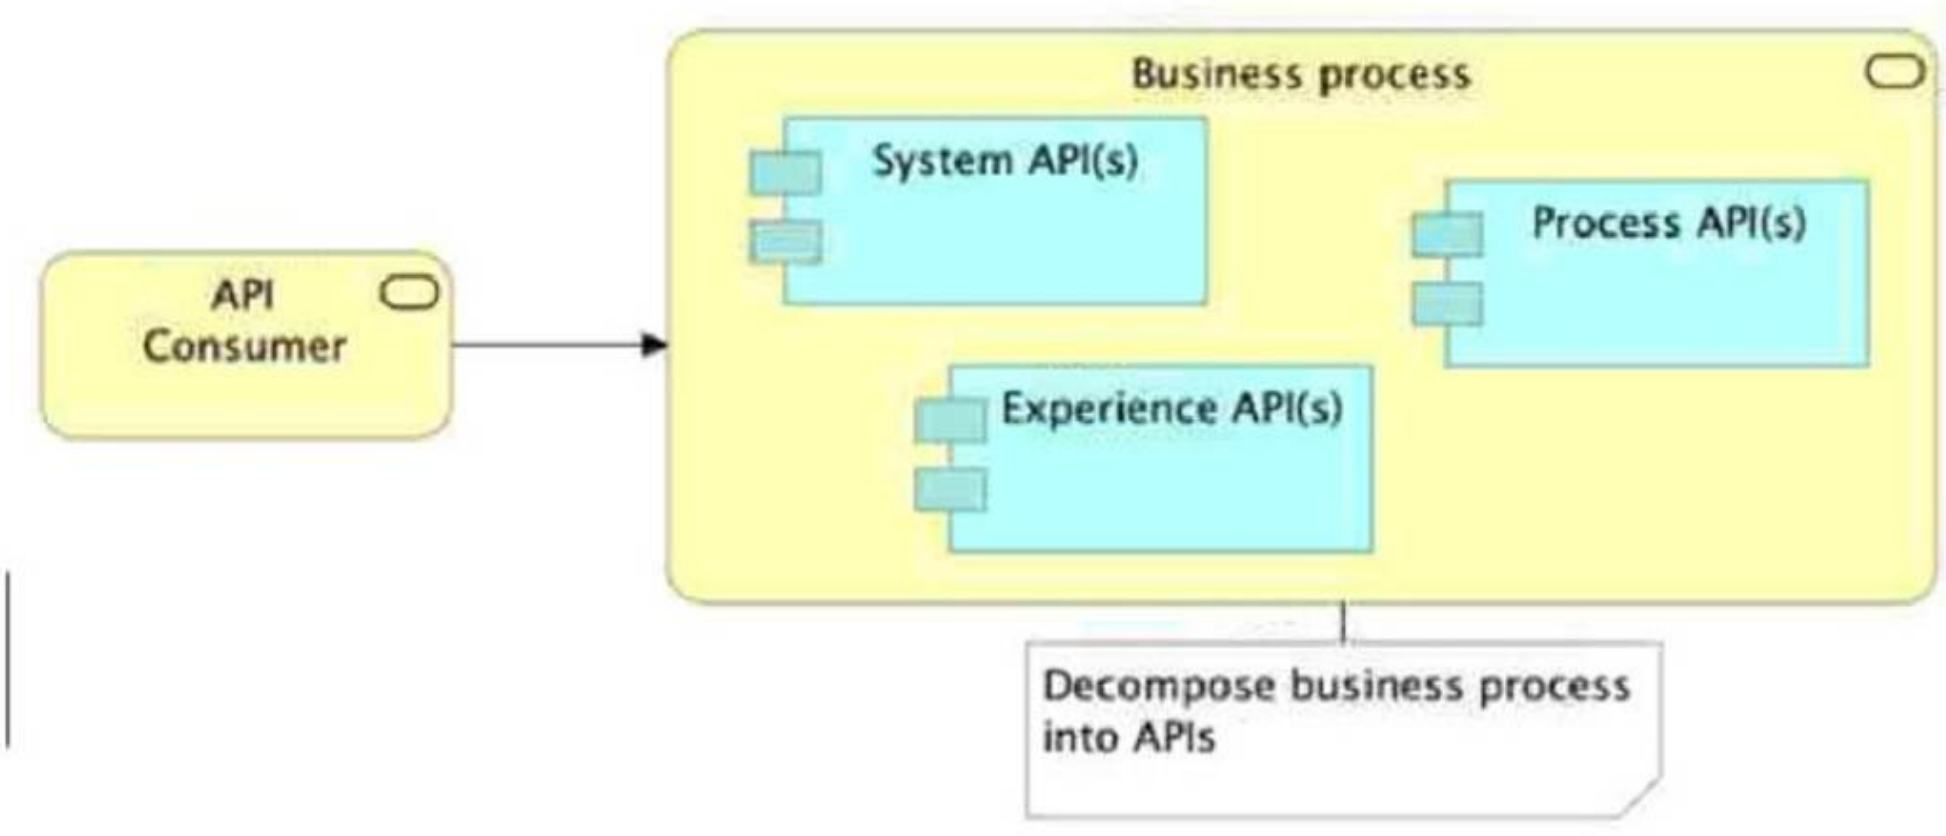

Which architectural constraint is compatible with the API-led connectivity architectural style?

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

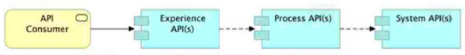

B. Use a Process API to-orchestrate calls to multiple System APIs but not to other Process APIs:
C. Allow System APIs to return data that is not currently required by the identified Process or Experience APIs

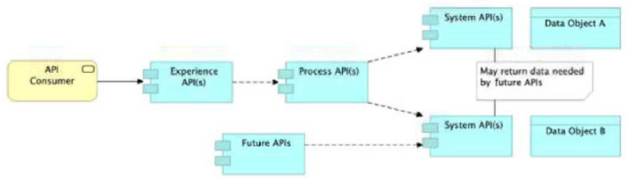

D. Handle customizations for the end-user application at the Process layer rather than at the Experience layer

Answer: B

NEW QUESTION 76
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

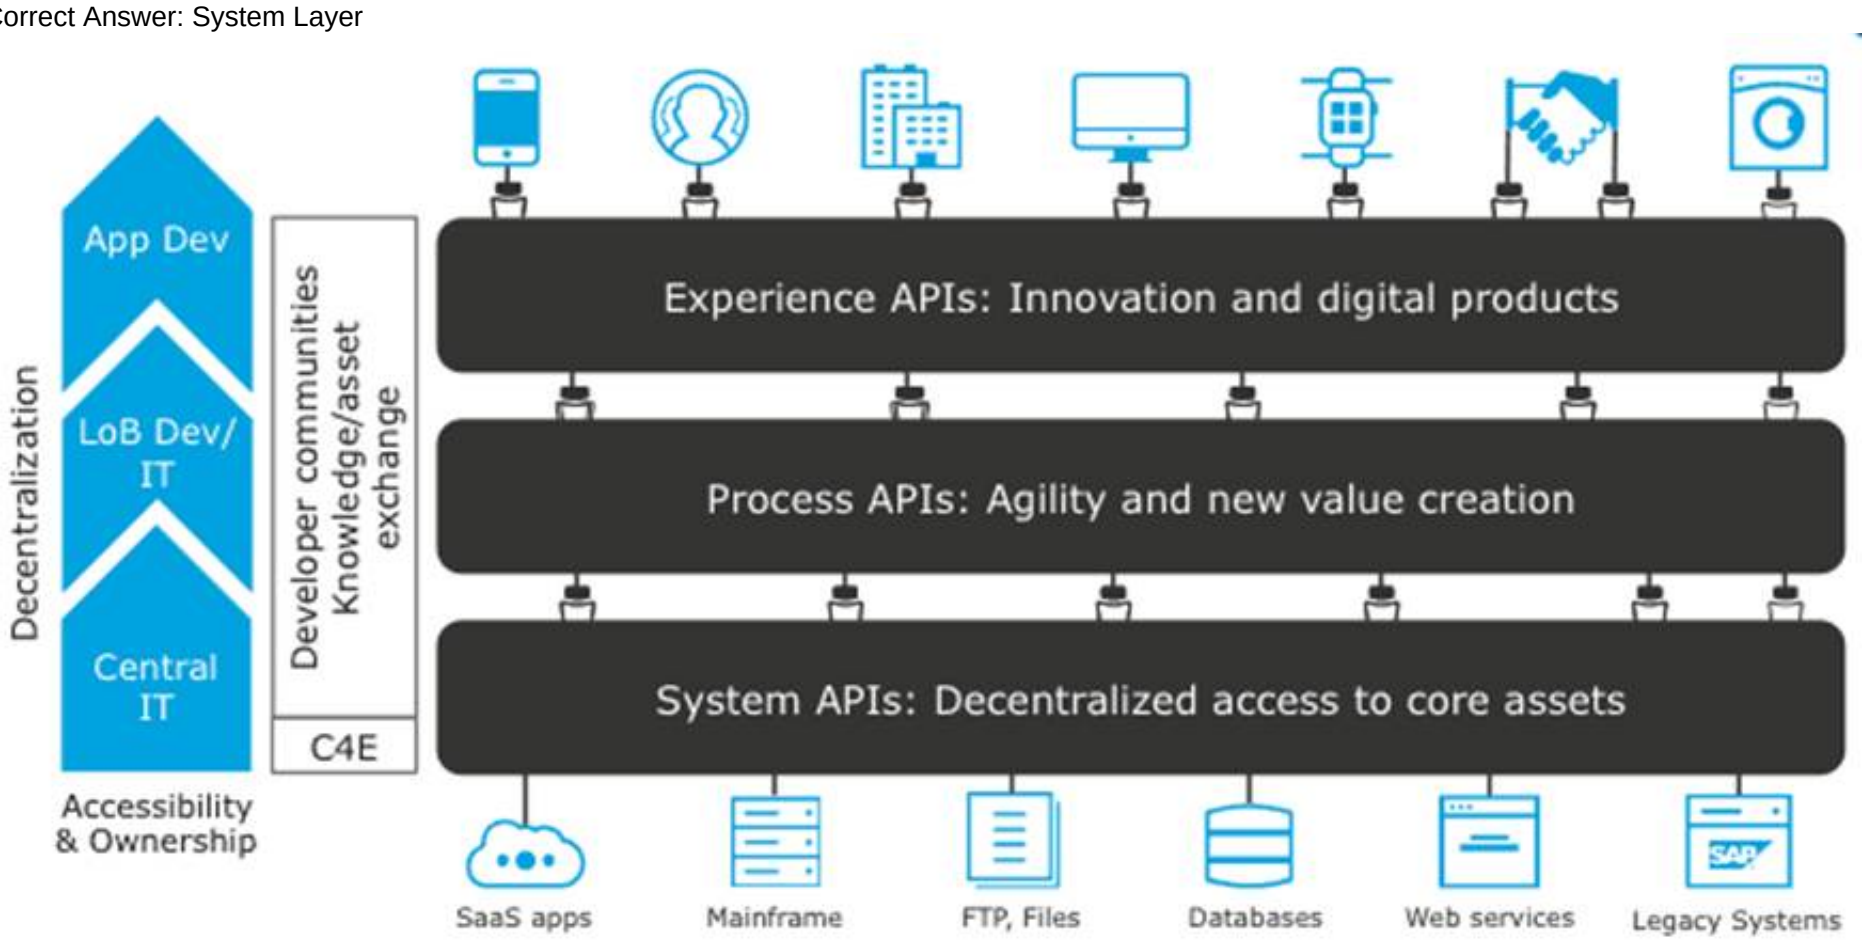

The APIs used in an API-led approach to connectivity fall into three categories:
System APIs – these usually access the core systems of record and provide a means of insulating the user from the complexity or any changes to the underlying systems. Once built, many users, can access data without any need 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 77
When must an API implementation be deployed to an Anypoint VPC?

A. When the API Implementation must invoke publicly exposed services that are deployed outside of CloudHub in a customer- managed AWS instance
B. When the API implementation must be accessible within a subnet of a restricted customer-hosted network that does not allow public access
C. When the API implementation must be deployed to a production AWS VPC using the Mule Maven plugin
D. When the API Implementation must write to a persistent Object Store

Answer: A

NEW QUESTION 78
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 82
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 86
Refer to the exhibit. An organization is running a Mule standalone runtime and has configured Active Directory as the Anypoint Platform external Identity Provider. The organization does not have budget for other system components.

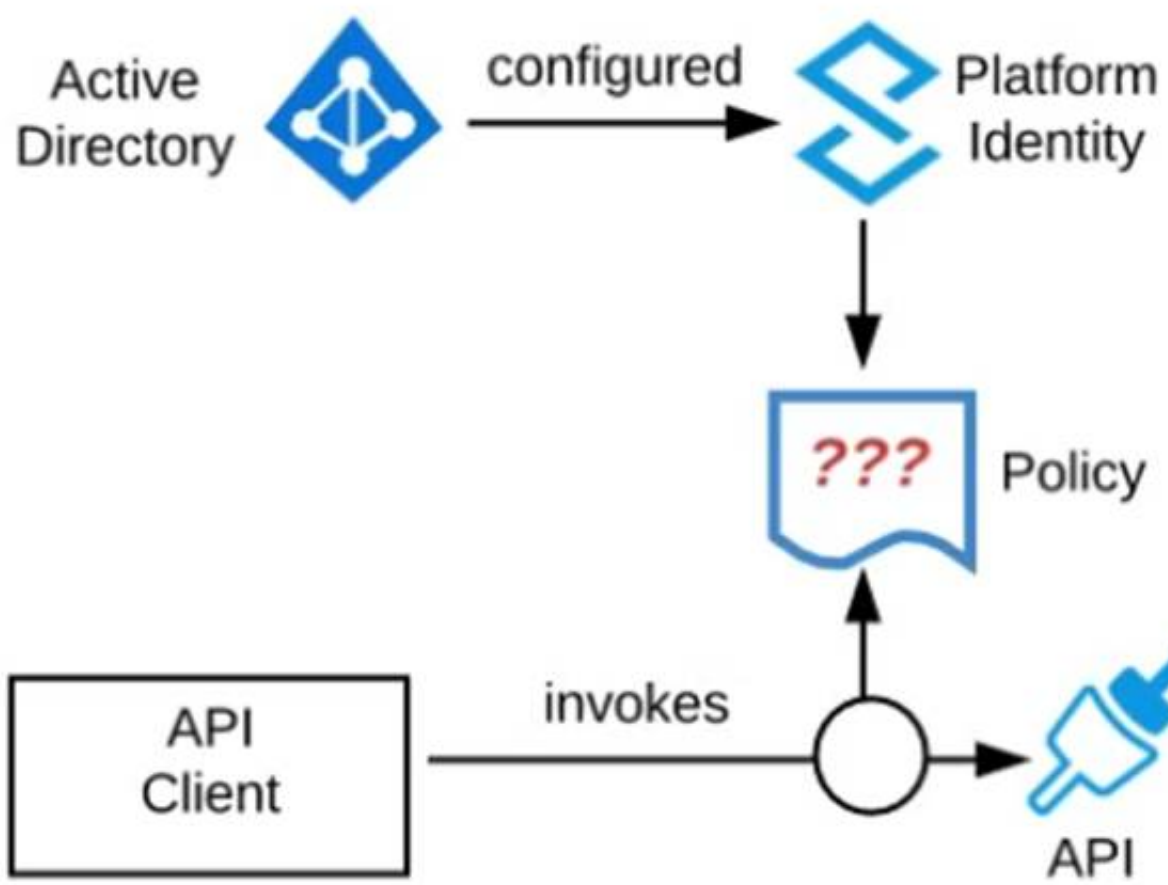

What policy should be applied to all instances of APIs in the organization to most effecuvelyKestrict access to a specific group of internal users?

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

Answer: A

NEW QUESTION 87
A business process is being implemented within an organization's application network. The architecture group proposes using a more coarse-grained application network design with relatively fewer APIs deployed to the application network compared to a more fine-grained design.
Overall, which factor typically increases with a more coarse-grained design for this business process implementation and deployment compared with using a more fine- grained design?

A. The complexity of each API implementation
B. The number of discoverable assets related to APIs deployed in the application network
C. The number of possible connections between API implementations in the application network
D. The usage of network infrastructure resources by the application network

Answer: A

NEW QUESTION 92
A circuit breaker strategy is planned in order to meet the goal of improved response time and demand on a downstream API.
* Circuit Open: More than 10 errors per minute for three minutes
* Circuit Half-Open: One error per minute
* Circuit Closed: Less than one error per minute for five minutes
Out of several proposals from the engineering team, which option will meet this goal?

A. Create a custom policy that implements the circuit breaker and includes policy template expressions for the required settings
B. Create Anypoint Monitoring alerts for Circuit Open/Closed configurations, and then implement a retry strategy for Circuit Half-Open configuration
C. Add the Circuit Breaker policy to the API instance, and configure the required settings
D. Implement the strategy in a Mule application, and provide the settings in the YAMLconfiguration

Answer: C

NEW QUESTION 97
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 98
A large lending company has developed an API to unlock data from a database server and web server. The API has been deployed to Anypoint Virtual Private Cloud (VPC) on CloudHub 1.0.
The database server and web server are in the customer's secure network and are not accessible through the public internet. The database server is in the customer's AWS VPC, whereas the web server is in the customer's on-premises corporate data center. How can access be enabled for the API to connect with the database server and the web server?

A. Set up VPC peering with AWS VPC and a VPN tunnel to the customer's on-premises corporate data center
B. Set up VPC peering with AWS VPC and the customer's on-premises corporate data center
C. Setup a transit gateway to the customer's on-premises corporate data center through AWS VPC
D. Set up VPC peering with the customer's on-premises corporate data center and a VPN tunnel to AWS VPC

Answer: A

NEW QUESTION 99
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 103
In an organization, the InfoSec team is investigating Anypoint Platform related data traffic.
From where does most of the data available to Anypoint Platform for monitoring and alerting originate?

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

D. From the Mule runtime irrespective of the deployment model

Answer: D

NEW QUESTION 104
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?

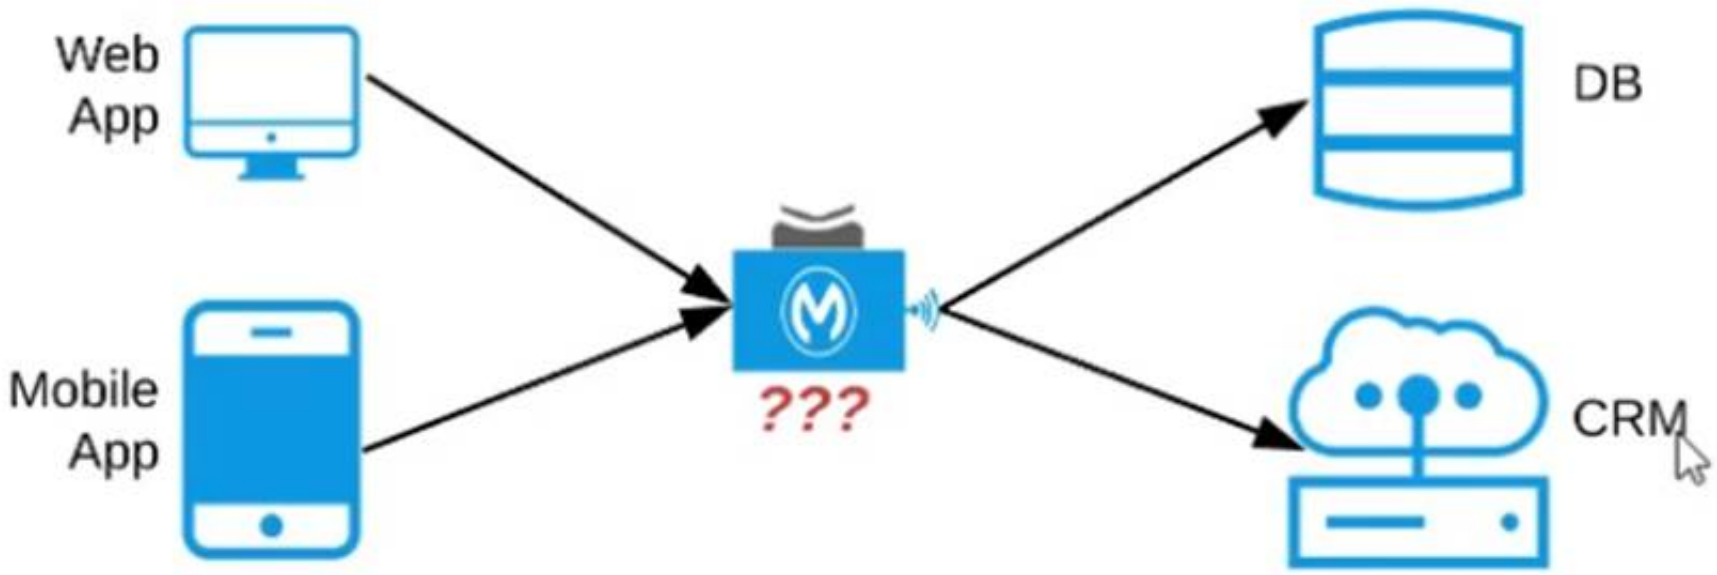

A) A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes

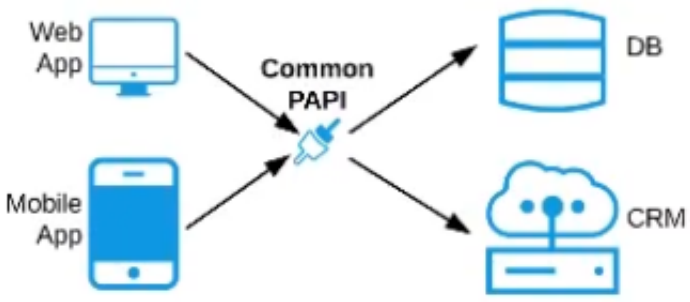

B) One set of APIs (Experience API, Process API, and System API) for the web app, and another set for the mobile app

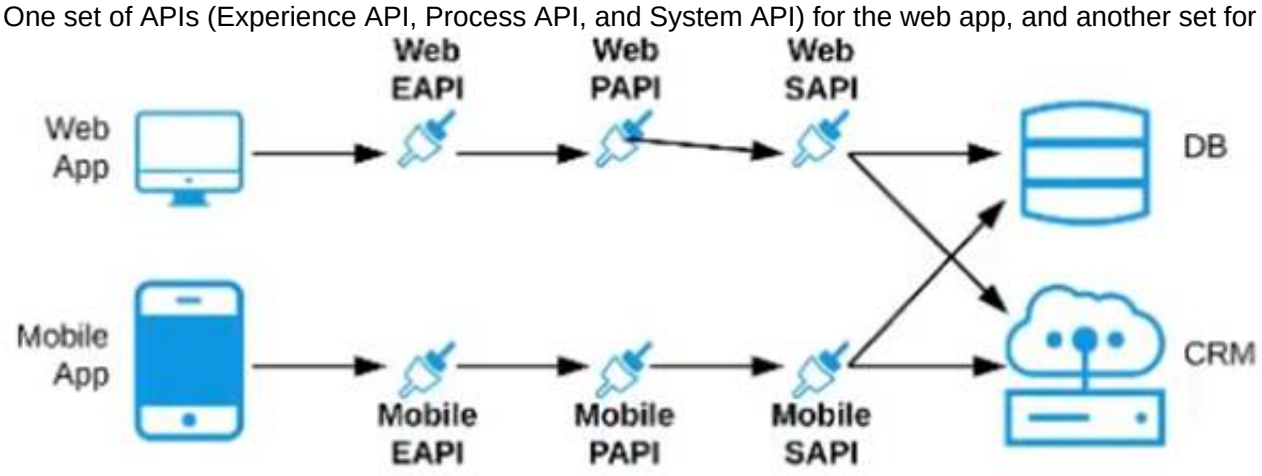

C) Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system

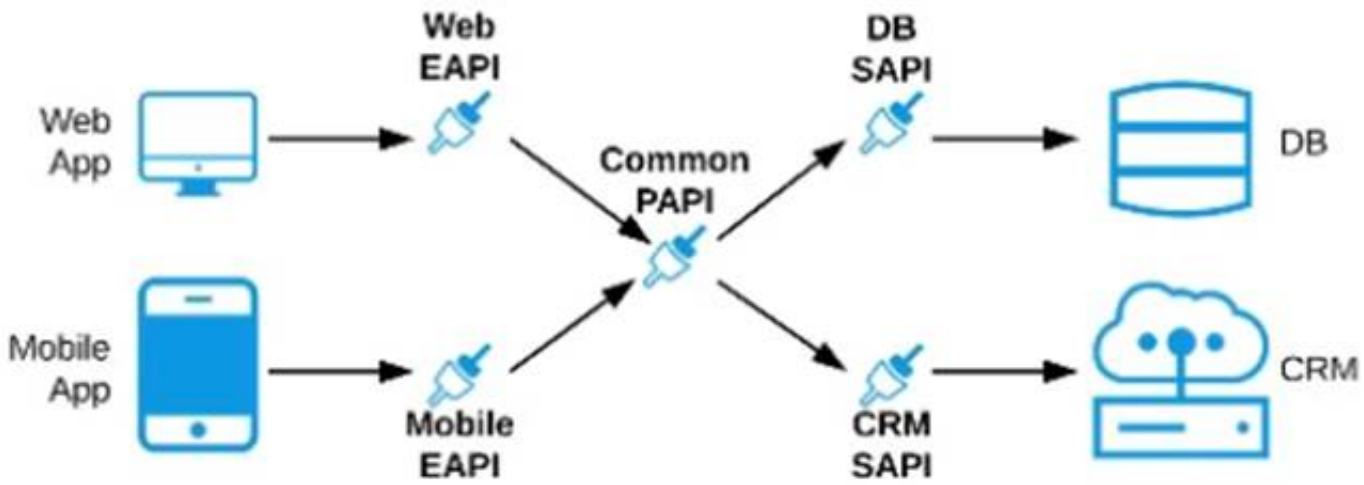

D) A common Experience API used by both the web and mobile apps, but separate Process APIs for the web and mobile apps that interact with the database and the CRM System

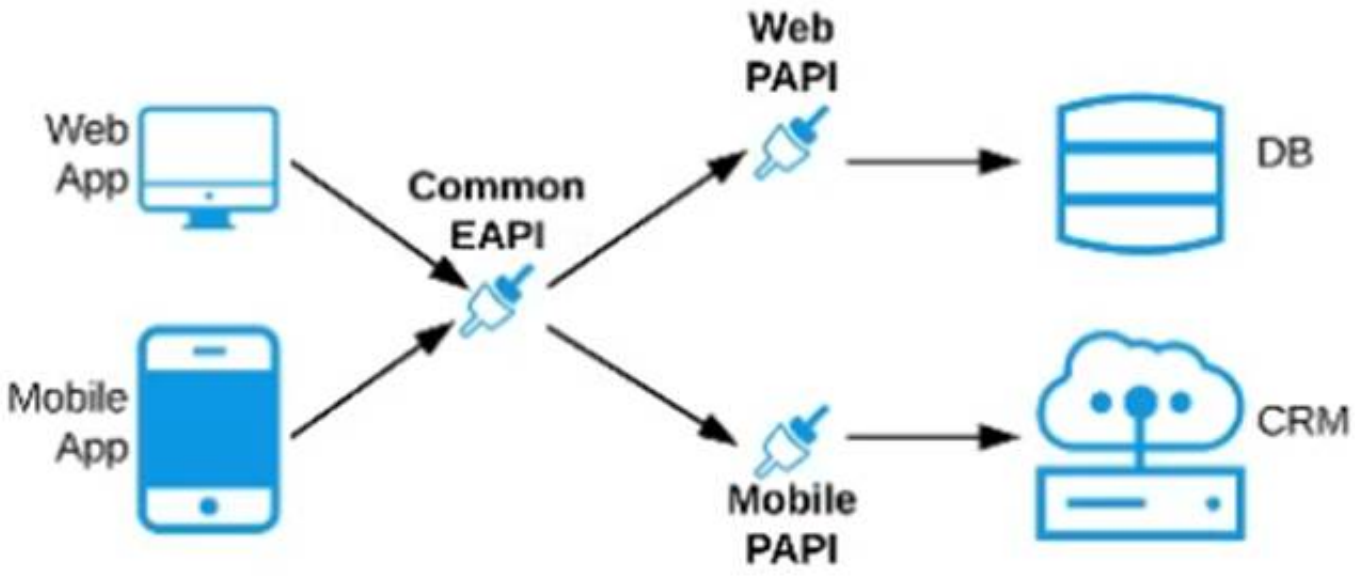

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

NEW QUESTION 107
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 110
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 114
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 115
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

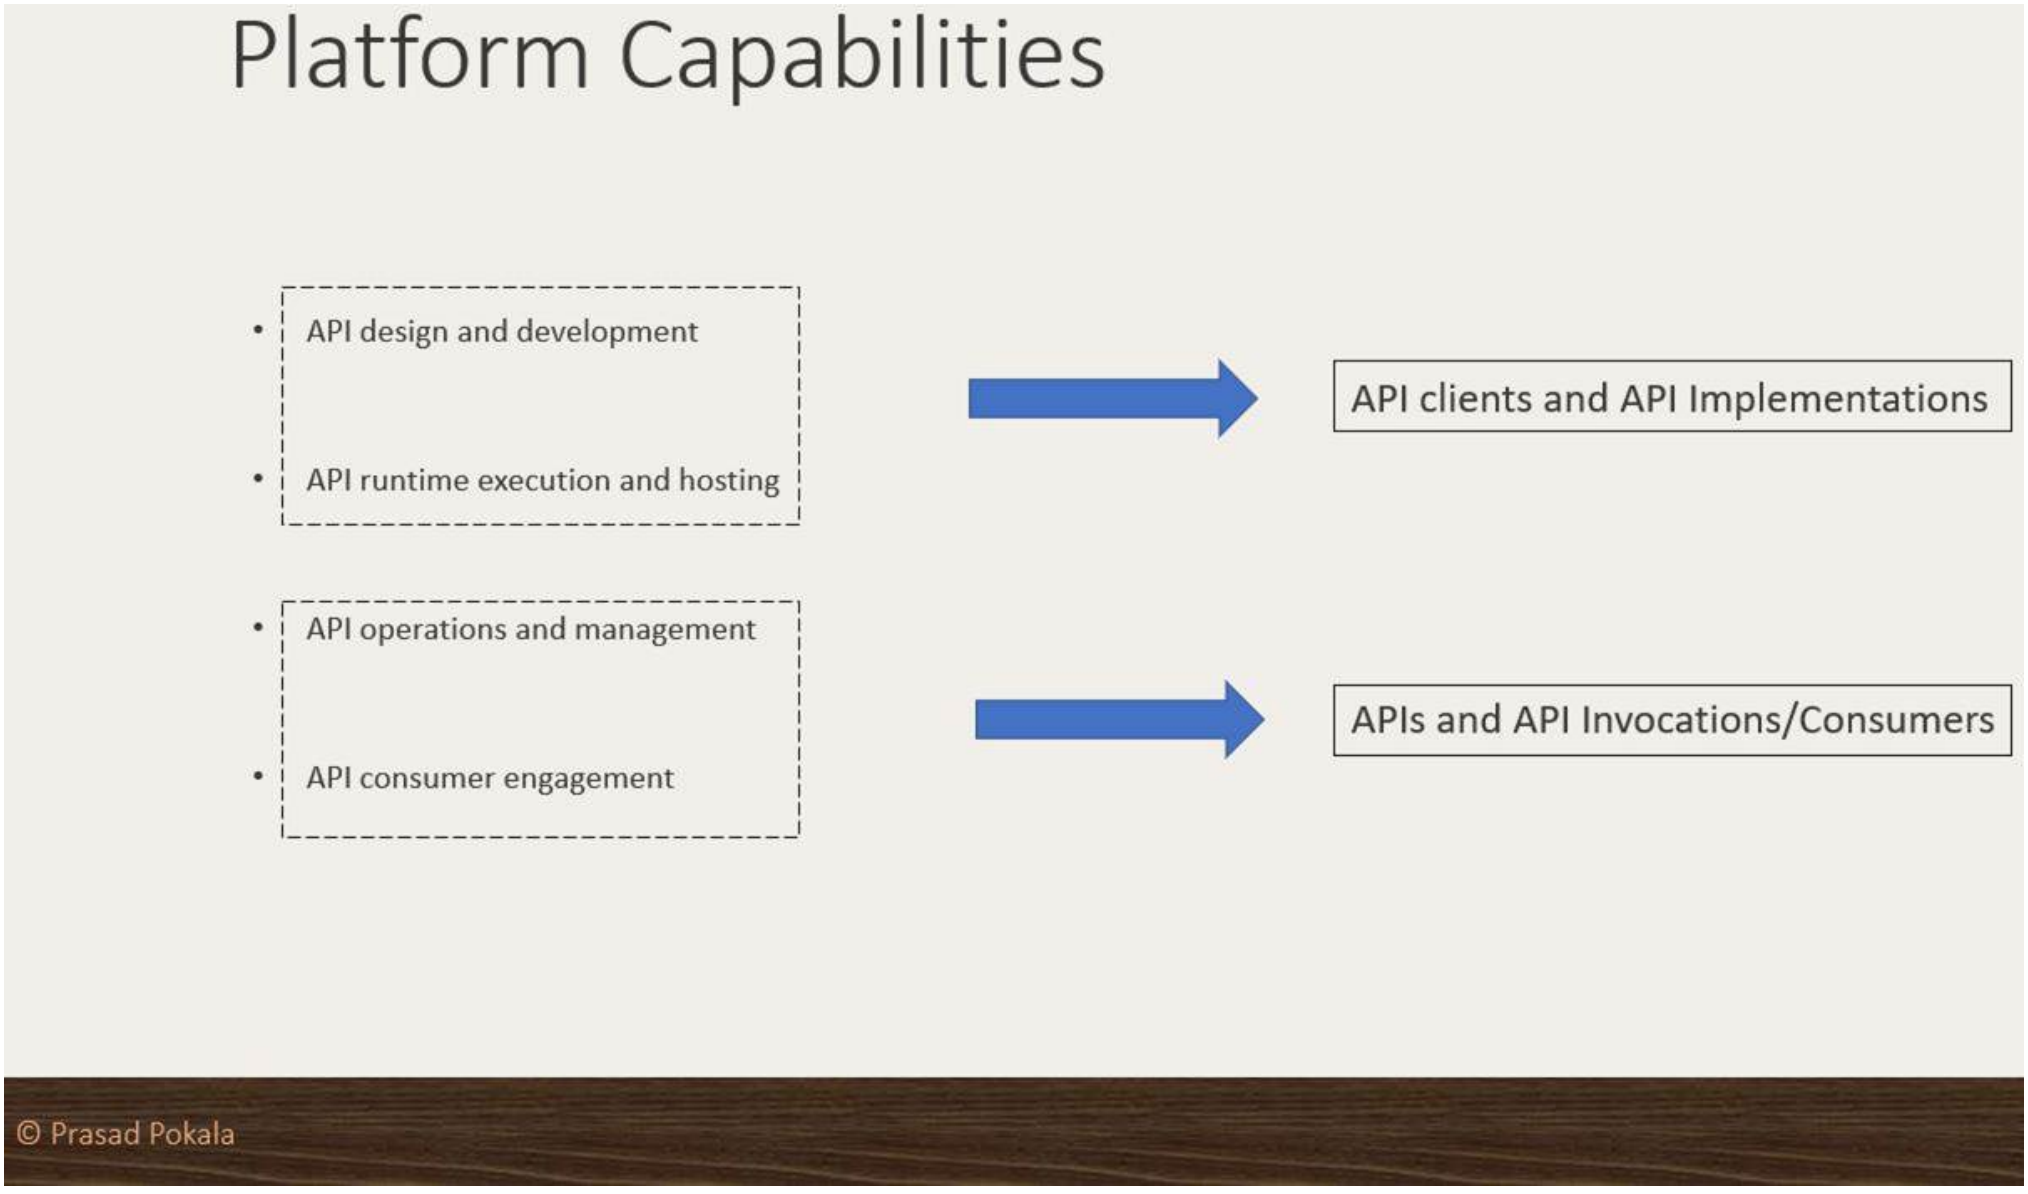

NEW QUESTION 117
The Line of Business (LoB) of an eCommerce company is requesting a process that sends automated notifications via email every time a new order is processed through the customer's mobile application or through the internal company's web application. In the future, multiple notification channels may be added: for example, text messages and push notifications.
What is the most effective API-led connectivity approach for the scenario described above?

A. Create one Experience API for the web application and one for the mobile applicatio
B. Create a Process API to orchestrate and retrieve the email template from = databas
C. Create a System API that sends the email using the Anypoint Connector for Emai
D. Create one Experience API for the web application and one for the mobile applicatio
E. Create a Process API to orchestrate and retrieve the email template from = databas
F. Create a System API that sends the email using the Anypoint Connector for Email.
G. Create one Experience API for the web application and one for the mobile application, Create a Process API to orchestrate, retrieve the email template from a database, and send the email using the Anypoint Connector for Email.

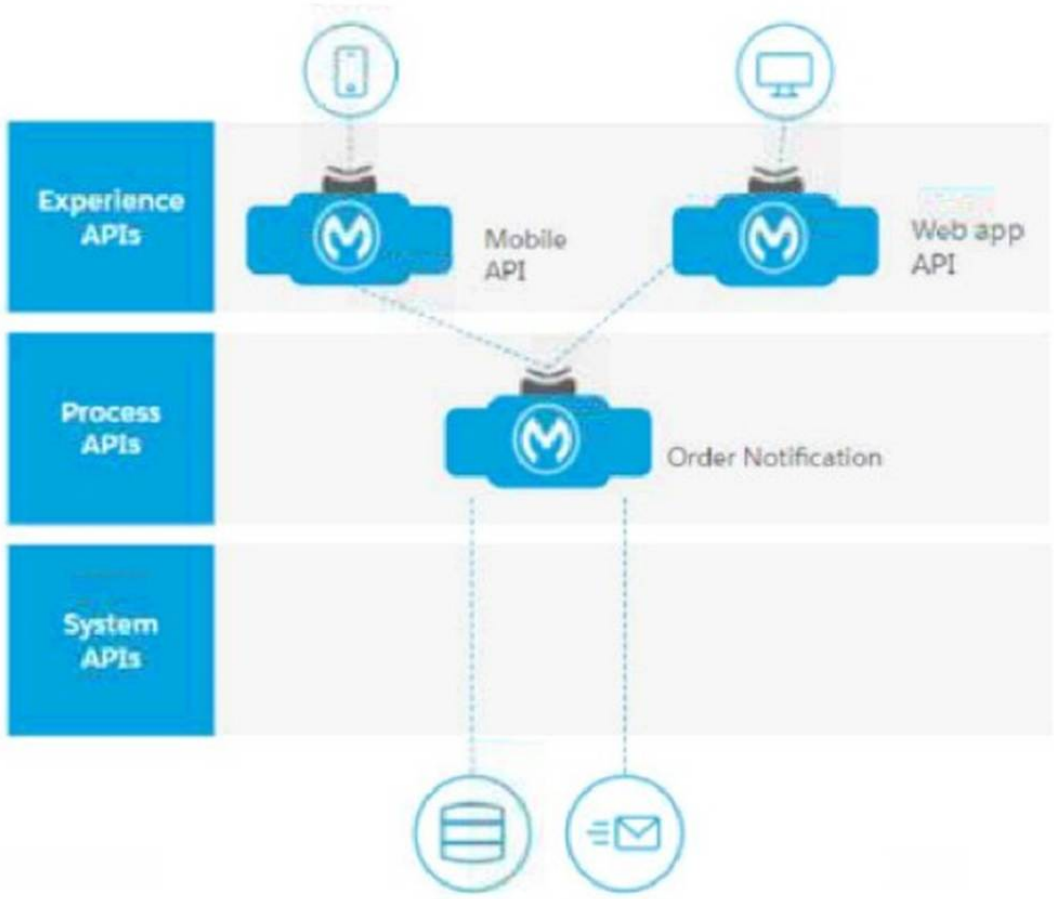

H. Create Experience APIs for both the web application and mobile application.Create a Process API ta orchestrate, retrieve the email template from e database, and send the email using the Anypoint Connector for Email.
I. Create Experience APIs for both the web application and mobile application.(Create 3 Process API to orchestrate and retrieve the email template from 2 databes
J. Create a System API that sends the email using the Anypoint Connector for Email.

Answer: A

NEW QUESTION 121
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 124
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 128
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 131
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 136
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 138
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 143
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 146
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 148
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 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 155
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 157
What is true about automating interactions with Anypoint Platform using tools such as Anypoint Platform REST APIs, Anypoint CU, or the Mule Maven plugin?

A. Access to Anypoint Platform APIs and Anypoint CU can be controlled separately through the roles and permissions in Anypoint Platform, so that specific users can get access to Anypoint CLI white others get access to the platform APIs
B. Anypoint Platform APIs can ONLY automate interactions with CloudHub, while the Mule Maven plugin is required for deployment to customer-hosted Mule runtimes
C. By default, the Anypoint CLI and Mule Maven plugin are NOT included in the Mule runtime, so are NOT available to be used by deployed Mule applications
D. API policies can be applied to the Anypoint Platform APIs so that ONLY certain LOBs have access to specific functions

Answer: C

NEW QUESTION 158
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 160
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 163
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 168
.....

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