MuleSoft

Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1

NEW QUESTION 1
An Integration Mule application is being designed to synchronize customer data between two systems. One system is an IBM Mainframe and the other system is a Salesforce Marketing Cloud (CRM) instance. Both systems have been deployed in their typical configurations, and are to be invoked using the native protocols provided by Salesforce and IBM.
What interface technologies are the most straightforward and appropriate to use in this Mute application to interact with these systems, assuming that Anypoint Connectors exist that implement these interface technologies?

A. IBM: DB access CRM: gRPC
B. IBM: REST CRM:REST
C. IBM: Active MQ CRM: REST
D. IBM: CICS CRM: SOAP

Answer: D

Explanation:
Correct answer is IBM: CICS CRM: SOAP
* Within Anypoint Exchange, MuleSoft offers the IBM CICS connector. Anypoint Connector for IBM CICS Transaction Gateway (IBM CTG Connector) provides integration with back-end CICS apps using the CICS Transaction Gateway.
* Anypoint Connector for Salesforce Marketing Cloud (Marketing Cloud Connector) enables you to connect to the Marketing Cloud API web services (now known as the Marketing Cloud API), which is also known as the Salesforce Marketing Cloud. This connector exposes convenient operations via SOAP for exploiting the capabilities of Salesforce Marketing Cloud.

NEW QUESTION 2
A system API EmployeeSAPI is used to fetch employee's data from an underlying SQL database.
The architect must design a caching strategy to query the database only when there is an update to the employees stable or else return a cached response in order to minimize the number of redundant transactions being handled by the database.
What must the architect do to achieve the caching objective?

A. Use an On Table Row on employees table and call invalidate cache Use an object store caching strategy and expiration interval to empty
B. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and expiration interval to empty
C. Use a Scheduler with a fixed frequency every hour triggering an invalidate cache flow Use an object store caching strategy and set expiration interval to 1-hour
D. Use an on table rule on employees table call invalidate cache and said new employees data to cache Use an object store caching strategy and set expiration interval to 1-hour

Answer: A

NEW QUESTION 3
Mule application is deployed to Customer Hosted Runtime. Asynchronous logging was implemented to improved throughput of the system. But it was observed over the period of time that few of the important exception log messages which were used to rollback transactions are not working as expected causing huge loss to the Organization. Organization wants to avoid these losses. Application also has constraints due to which they cant compromise on throughput much. What is the possible option in this case?

A. Logging needs to be changed from asynchronous to synchronous
B. External log appender needs to be used in this case
C. Persistent memory storage should be used in such scenarios
D. Mixed configuration of asynchronous or synchronous loggers should be used to log exceptions via synchronous way

Answer: D

Explanation:
Correct approach is to use Mixed configuration of asynchronous or synchronous loggers shoud be used to log exceptions via synchronous way Asynchronous logging poses a performance-reliability trade-off. You may lose some messages if Mule crashes before the logging buffers flush to the disk. In this case, consider that you can have a mixed configuration of asynchronous or synchronous loggers in your app. Best practice is to use asynchronous logging over synchronous with a minimum logging level of WARN for a production application. In some cases, enable INFO logging level when you need to confirm events such as successful policy installation or to perform troubleshooting. Configure your logging strategy by editing your application's src/main/resources/log4j2.xml file

NEW QUESTION 4
An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to custom hosted Mule runtimes (on-premises) while others execute in the MuleSoft-hosted runtime plane (CloudHub). To perform the Integra functionality, these Mule applications connect to various backend systems, with multiple applications typically needing to access the backend systems.
How can the organization most effectively avoid creating duplicates in each Mule application of the credentials required to access the backend systems?

A. Create a Mule domain project that maintains the credentials as Mule domain-shared resources Deploy the Mule applications to the Mule domain, so the credentials are available to the Mule applications
B. Store the credentials in properties files in a shared folder within the organization's data center Have the Mule applications load properties files from this shared location at startup
C. Segregate the credentials for each backend system into environment-specific properties files Package these properties files in each Mule application, from where they are loaded at startup
D. Configure or create a credentials service that returns the credentials for each backend system, and that is accessible from customer-hosted and MuleSoft-hosted Mule runtimes Have the Mule applications toad the properties at startup by invoking that credentials service

Answer: D

Explanation:
* "Create a Mule domain project that maintains the credentials as Mule domain-shared resources" is wrong as domain project is not supported in Cloudhub * We should Avoid Creating duplicates in each Mule application but below two options cause duplication of credentials - Store the credentials in properties files in a shared folder within the organization's data center. Have the Mule applications load properties files from this shared location at startup - Segregate the credentials for each backend system into environment-specific properties files. Package these properties files in each Mule application, from where they are loaded at startup

So these are also wrong choices * Credentials service is the best approach in this scenario. Mule domain projects are not supported on CloudHub. Also its is not recommended to have multiple copies of configuration values as this makes difficult to maintain Use the Mule Credentials Vault to encrypt data in a .properties file. (In the context of this document, we refer to the .properties file simply as the properties file.) The properties file in Mule stores data as key-value pairs which may contain information such as usernames, first and last names, and credit card numbers. A Mule application may access this data as it processes messages, for example, to acquire login credentials for an external Web service. However, though this sensitive, private data must be stored in a properties file for Mule to access, it must also be protected against unauthorized – and potentially malicious – use by anyone with access to the Mule application

NEW QUESTION 5
In Anypoint Platform, a company wants to configure multiple identity providers(Idps) for various lines of business (LOBs) Multiple business groups and environments have been defined for the these LOBs. What Anypoint Platform feature can use multiple Idps access the company's business groups and environment?

A. User management
B. Roles and permissions
C. Dedicated load balancers
D. Client Management

Answer: D

Explanation:
Correct answer is Client Management
* Anypoint Platform acts as a client provider by default, but you can also configure external client providers to authorize client applications.
* As an API owner, you can apply an OAuth 2.0 policy to authorize client applications that try to access your API. You need an OAuth 2.0 provider to use an OAuth 2.0 policy.
* You can configure more than one client provider and associate the client providers with different environments. If you configure multiple client providers after you have already created environments, you can associate the new client providers with the environment.
* You should review the existing client configuration before reassigning client providers to avoid any downtime with existing assets or APIs.
* When you delete a client provider from your master organization, the client provider is no longer available in environments that used it.
* Also, assets or APIs that used the client provider can no longer authorize users who want to access them.
--MuleSoft
Reference: https://docs.mulesoft.com/access-management/managing-api-clients
https://www.folkstalk.com/2019/11/mulesoft-integration-and-platform.html

NEW QUESTION 6
A trading company handles millions of requests a day. Due to nature of its business, it requires excellent performance and reliability within its application.
For this purpose, company uses a number of event-based API's hosted on various mule clusters that communicate across a shared message queue sitting within its network.
Which method should be used to meet the company's requirement for its system?

A. XA transactions and XA connected components
B. JMS transactions
C. JMS manual acknowledgements with a reliability pattern
D. VM queues with reliability pattern

Answer: C

NEW QUESTION 7
A REST API is being designed to implement a Mule application.
What standard interface definition language can be used to define REST APIs?

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

Answer: B

NEW QUESTION 8
A finance giant is planning to migrate all its Mule applications to Runtime fabric (RTF). Currently all Mule applications are deployed cloud hub using automated CI/CD scripts.
As an integration architect, which of the below step would you suggest to ensure that the applications from cloudhub are migrated properly to Runtime Fabric (RTF) with an assumption that organization is keen on keeping the same deployment strategy.

A. No changes need to be made to POM.xml file and CI/CD script should be modified as per the RTF configurations
B. runtimeFabric dependency should be added as a mule plug-in to POM.xml file and CI/CD script should be modified as per the RTF configurations
C. runtimeFabric deployment should be added to POM.xml file in allthe mule applications and CI/CD script should be modified as per the RTF configurations
D. runtimeFabric profile should be added mule configuration files in the mule applications and CI/CD script should be modified as per the RTF configurations

Answer: C

NEW QUESTION 9
A company is using Mulesoft to develop API's and deploy them to Cloudhub and on premises targets. Recently it has decided to enable Runtime Fabric deployment option as well and infrastructure is set up for this option.
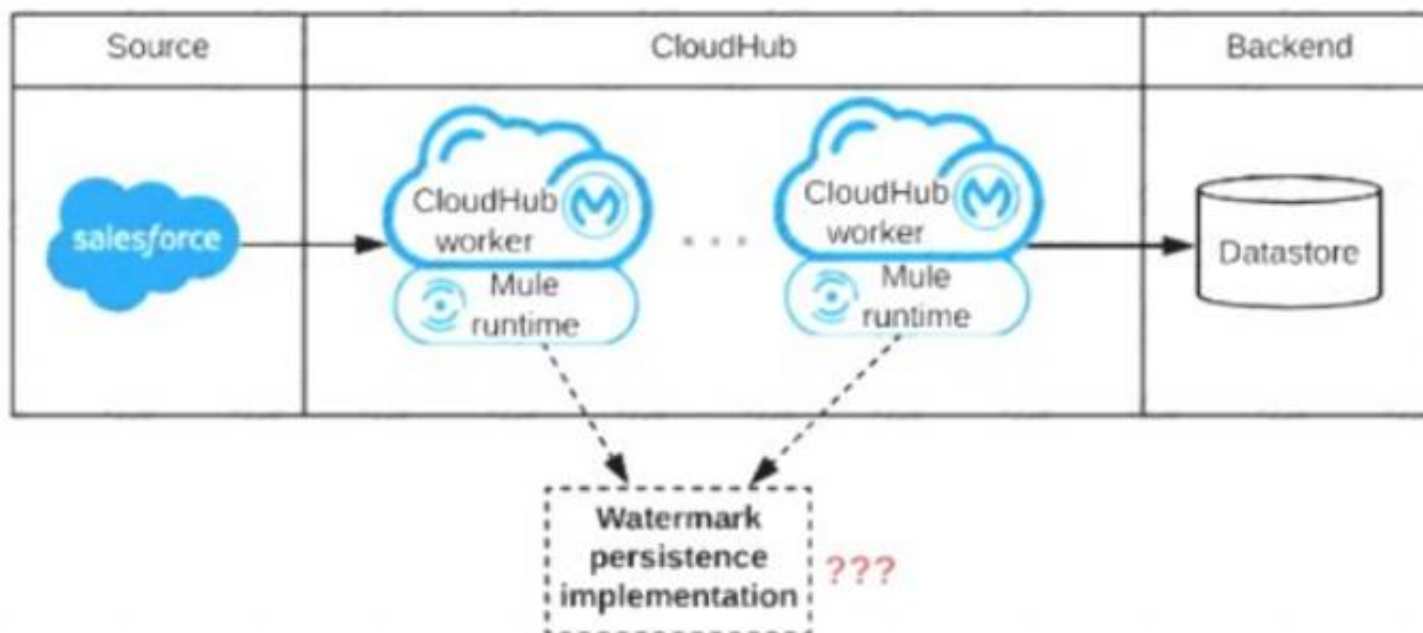What can be used to deploy Runtime Fabric?

A. AnypointCLI
B. Anypoint platform REST API's
C. Directly uploading ajar file from the Runtime manager

D. Mule maven plug-in

Answer: D

NEW QUESTION 10
Refer to the exhibit.

A Mule application is being designed to be deployed to several CIoudHub workers. The Mule application's integration logic is to replicate changed Accounts from Satesforce to a backend system every 5 minutes.
A watermark will be used to only retrieve those Satesforce Accounts that have been modified since the last time the integration logic ran.
What is the most appropriate way to implement persistence for the watermark in order to support the required data replication integration logic?

A. Persistent Anypoint MQ Queue
B. Persistent Object Store
C. Persistent Cache Scope
D. Persistent VM Queue

Answer: B

Explanation:
* An object store is a facility for storing objects in or across Mule applications. Mule uses object stores to persist data for eventual retrieval.
* Mule provides two types of object stores:
1) In-memory store – stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime.
2) Persistent store – Mule persists data when an object store is explicitly configured to be persistent.
In a standalone Mule runtime, Mule creates a default persistent store in the file system. If you do not specify an object store, the default persistent object store is used.
MuleSoft Reference: https://docs.mulesoft.com/mule-runtime/3.9/mule-object-stores

NEW QUESTION 10
A mule application is being designed to perform product orchestration. The Mule application needs to join together the responses from an inventory API and a Product Sales History API with the least latency.
To minimize the overall latency. What is the most idiomatic (used for its intended purpose) design to call each API request in the Mule application?

A. Call each API request in a separate lookup call from Dataweave reduce operator
B. Call each API request in a separate route of a Scatter-Gather
C. Call each API request in a separate route of a Parallel For Each scope
D. Call each API request in a separate Async scope

Answer: B

Explanation:
Scatter-Gather sends a request message to multiple targets concurrently. It collects the responses from all routes, and aggregates them into a single message.

NEW QUESTION 15
What is true about the network connections when a Mule application uses a JMS connector to interact with a JMS provider (message broker)?

A. To complete sending a JMS message, the JMS connector must establish a network connection with the JMS message recipient
B. To receive messages into the Mule application, the JMS provider initiates a network connection to the JMS connector and pushes messages along this connection
C. The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider
D. The AMQP protocol can be used by the JMS connector to portably establish connections to various types of JMS providers

Answer: C

Explanation:
* To send message or receive JMS (Java Message Service) message no separate network connection need to be established. So option A, C and D are ruled out.
Correct Answer The JMS connector supports both sending and receiving of JMS messages over the protocol determined by the JMS provider.
* JMS Connector enables sending and receiving messages to queues and topics for any message service that implements the JMS specification.
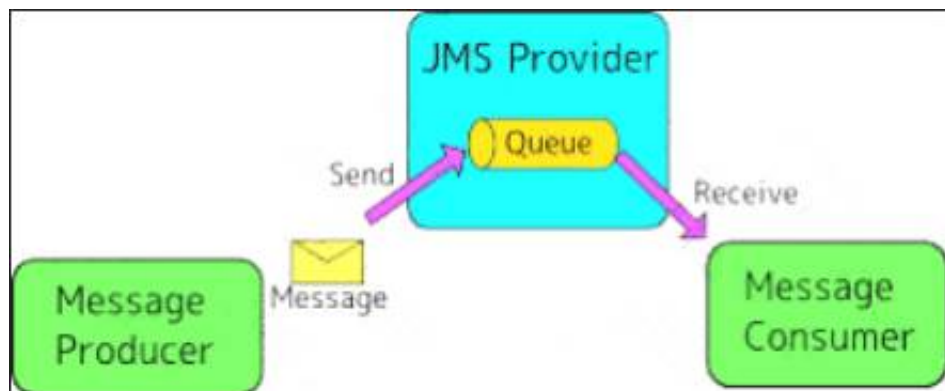* JMS is a widely used API for message-oriented middleware.
* It enables the communication between different components of a distributed application to be loosely coupled, reliable, and asynchronous.
MuleSoft Doc Reference: https://docs.mulesoft.com/jms-connector/1.7/
Diagram, text Description automatically generated

NEW QUESTION 17

An insurance company is implementing a MuleSoft API to get inventory details from the two vendors. Due to network issues, the invocations to vendor applications are getting timed-out intermittently. But the transactions are successful upon reprocessing
What is the most performant way of implementing this requirement?

A. Implement a scatter-gather scope to invoke the two vendor applications on two different routeUse the Until-Successful scope to implement the retry mechanism for timeout errors on each route
B. Implement a Choice scope to invoke the two vendor applications on two different route Use the try-catch scope to implement the retry mechanism for timeout errors on each route
C. Implement a For-Each scope to invoke the two vendor applicationsUse until successful scope to implement the retry mechanism for the timeout errors
D. Implement Round-Robin scope to invoke the two vendor applications on two different routes Use the Try-Catch scope to implement retry mechanism for timeout errors on each route

Answer: A

NEW QUESTION 19

An organization's security requirements mandate centralized control at all times over authentication and authorization of external applications when invoking web APIs managed on Anypoint Platform.
What Anypoint Platform feature is most idiomatic (used for its intended purpose), straightforward, and maintainable to use to meet this requirement?

A. Client management configured in access management
B. Identity management configured in access management
C. Enterprise Security module coded in Mule applications
D. External access configured in API Manager

Answer: B

NEW QUESTION 22

A company is planning to migrate its deployment environment from on-premises cluster to a Runtime Fabric (RTF) cluster. It also has a requirement to enable Mule applications deployed to a Mule runtime instance to store and share data across application replicas and restarts.
How can these requirements be met?

A. Anypoint object store V2 to share data between replicas in the RTF cluster
B. Install the object store pod on one of the cluster nodes
C. Configure Persistence Gateway in any of the servers using Mule Object Store
D. Configure Persistent Gateway at the RTF

Answer: D

NEW QUESTION 23

A Mule application is being designed To receive nightly a CSV file containing millions of records from an external vendor over SFTP, The records from the file need to be validated, transformed. And then written to a database. Records can be inserted into the database in any order.
In this use case, what combination of Mule components provides the most effective and performant way to write these records to the database?

A. Use a Parallel for Each scope to Insert records one by one into the database
B. Use a Scatter-Gather to bulk insert records into the database
C. Use a Batch job scope to bulk insert records into the database.
D. Use a DataWeave map operation and an Async scope to insert records one by one into the database.

Answer: C

Explanation:
Correct answer is Use a Batch job scope to bulk insert records into the database
* Batch Job is most efficient way to manage millions of records. A few points to note here are as follows :
Reliability: If you want reliabilty while processing the records, i.e should the processing survive a runtime crash or other unhappy scenarios, and when restarted process all the remaining records, if yes then go for batch as it uses persistent queues.
Error Handling: In Parallel for each an error in a particular route will stop processing the remaining records in that route and in such case you'd need to handle it using on error continue, batch process does not stop during such error instead you can have a step for failures and have a dedicated handling in it.
Memory footprint: Since question said that there are millions of records to process, parallel for each will aggregate all the processed records at the end and can possibly cause Out Of Memory.
Batch job instead provides a BatchResult in the on complete phase where you can get the count of failures and success. For huge file processing if order is not a concern definitely go ahead with Batch Job

NEW QUESTION 24

Organization wants to achieve high availability goal for Mule applications in customer hosted runtime plane. Due to the complexity involved, data cannot be shared among of different instances of same Mule application. What option best suits to this requirement considering high availability is very much critical to the organization?

A. The cluster can be configured
B. Use third party product to implement load balancer
C. High availability can be achieved only in CloudHub
D. Use persistent object store

Answer: B

Explanation:
High availability is about up-time of your application
A) High availability can be achieved only in CloudHub isn't correct statement. It can be achieved in customer hosted runtime planes as well
B) An object store is a facility for storing objects in or across Mule applications. Mule runtime engine (Mule) uses object stores to persist data for eventual retrieval. It can be used for disaster recovery but not for High Availability. Using object store can't guarantee that all instances won't go down at once. So not an appropriate choice.

NEW QUESTION 28
An organization is designing a Mule application to periodically poll an SFTP location for new files containing sales order records and then process those sales orders. Each sales order must be processed exactly once.
To support this requirement, the Mule application must identify and filter duplicate sales orders on the basis of a unique ID contained in each sales order record and then only send the new sales orders to the downstream system.
What is the most idiomatic (used for its intended purpose) Anypoint connector, validator, or scope that can be configured in the Mule application to filter duplicate sales orders on the basis of the unique ID field contained in each sales order record?

A. Configure a Cache scope to filter and store each record from the received file by the order ID
B. Configure a Database connector to filter and store each record by the order ID
C. Configure an Idempotent Message Validator component to filter each record by the order ID
D. Configure a watermark In an On New or Updated File event source to filter unique records by the order ID

Answer: C

NEW QUESTION 32
An organization uses a set of customer-hosted Mule runtimes that are managed using the Mulesoft-hosted control plane. What is a condition that can be alerted on from Anypoint Runtime Manager without any custom components or custom coding?

A. When a Mule runtime on a given customer-hosted server is experiencing high memory consumption during certain periods
B. When an SSL certificate used by one of the deployed Mule applications is about to expire
C. When the Mute runtime license installed on a Mule runtime is about to expire
D. When a Mule runtime's customer-hosted server is about to run out of disk space

Answer: A

Explanation:
Correct answer is When a Mule runtime on a given customer-hosted server is experiencing high memory consumption during certain periods Using Anypoint Monitoring, you can configure two different types of alerts: Basic alerts for servers and Mule apps Limit per organization: Up to 50 basic alerts for users who do not have a Titanium subscription to Anypoint Platform You can set up basic alerts to trigger email notifications when a metric you are measuring passes a specified threshold. You can create basic alerts for the following metrics for servers or Mule apps: For on-premises servers and CloudHub apps: * CPU utilization * Memory utilization * Thread count Advanced alerts for graphs in custom dashboards in Anypoint Monitoring. You must have a Titanium subscription to use this feature. Limit per organization: Up to 20 advanced alerts

NEW QUESTION 36
An Organization has previously provisioned its own AWS VPC hosting various servers. The organization now needs to use Cloudhub to host a Mule application that will implement a REST API once deployed to Cloudhub, this Mule application must be able to communicate securely with the customer-provisioned AWS VPC resources within the same region, without being interceptable on the public internet.
What Anypoint Platform features should be used to meet these network communication requirements between Cloudhub and the existing customer-provisioned AWS VPC?

A. Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC
B. Configure an external identity provider (IDP) in Anypoint Platform with certificates from the customer provisioned AWS VPC
C. Add a default API Whitelisting policy to API Manager to automatically whitelist the customer provisioned AWS VPC IP ranges needed by the Mule applicaton
D. Use VM queues in the Mule application to allow any non-mule assets within the customer provisioned AWS VPC to subscribed to and receive messages

Answer: A

Explanation:
Correct answer is: Add a Mulesoft hosted Anypoint VPC configured and with VPC Peering to the AWS VPC
* Connecting to your Anypoint VPC extends your corporate network and allows CloudHub workers to access resources behind your corporate firewall.
* You can connect on-premises data centers through a secured VPN tunnel, or a private AWS VPC through VPC peering, or by using AWS Direct Connect.
MuleSoft Doc Reference : https://docs.mulesoft.com/runtime-manager/virtual-private-cloud

NEW QUESTION 41
An organization is using Mulesoft cloudhub and develops API's in the latest version. As a part of requirements for one of the API's, third party API needs to be called. The security team has made it clear that calling any external API needs to have include listing
As an integration architect please suggest the best way to accomplish the design plan to support these requirements?

A. Implement includelist IP on the cloudhub VPC firewall to allow the traffic
B. Implement the validation of includelisted IP operation

C. Implement the Any point filter processor to implement the include list IP
D. Implement a proxy for the third party API and enforce the IPinclude list policy and call this proxy from the flow of the API

Answer: D

NEW QUESTION 46
As a part of project requirement, client will send a stream of data to mule application. Payload size can vary between 10mb to 5GB. Mule application is required to transform the data and send across multiple sftp servers. Due to the cost cuttings in the organization, mule application can only be allocated one worker with size of 0.2 vCore.
As an integration architect , which streaming strategy you would suggest to handle this scenario?

A. In-memory non repeatable stream
B. File based non-repeatable stream
C. In-memory repeatable stream
D. File based repeatable storage

Answer: D

Explanation:
As the question says that data needs to be sent across multiple sftp serves , we cannot use non-repeatable streams. The non-repeatable strategy disables repeatable streams, which enables you to read an input stream only once.
You cant use in memory storage because with 0.2 vcore you will get only 1 GB of heap memory. Hence application will error out for file more than 1 GB.
Hence the correct option is file base repeatable stream

NEW QUESTION 48
A mule application uses an HTTP request operation to involve an external API. The external API follows the HTTP specification for proper status code usage.
What is possible cause when a 3xx status code is returned to the HTTP Request operation from the external API?

A. The request was not accepted by the external API
B. The request was Redirected to a different URL by the external API
C. The request was NOT RECEIVED by the external API
D. The request was ACCEPTED by the external API

Answer: B

Explanation:
3xx HTTP status codes indicate a redirection that the user agent (a web browser or a crawler) needs to take further action when trying to access a particular resource.

NEW QUESTION 53
As an enterprise architect, what are the two reasons for which you would use a canonical data model in the new integration project using Mulesoft Anypoint platform (choose two answers)

A. To have consistent data structure aligned in processes
B. To isolate areas within a bounded context
C. To incorporate industry standard data formats
D. There are multiple canonical definitions of each data type
E. Because the model isolates the back and systems and support mule applications from change

Answer: AB

NEW QUESTION 54
A mule application designed to fulfil two requirements
a) Processing files are synchronously from an FTPS server to a back-end database using VM intermediary
queues for load balancing VM events
b) Processing a medium rate of records from a source to a target system using batch job scope
Considering the processing reliability requirements for FTPS files, how should VM queues be configured for processing files as well as for the batch job scope if the application is deployed to Cloudhub workers?

A. Use Cloud hub persistent queues for FTPS files processingThere is no need to configure VM queues for the batch jobs scope as it uses by default the worker' s disc for VM queueing
B. Use Cloud hub persistent VM queue for FTPS file processingThere is no need to configure VM queues for the batch jobs scope as it uses by default the worker' s JVM memory for VM queueing
C. Use Cloud hub persistent VM queues for FTPS file processing Disable VM queue for the batch job scope
D. Use VM connector persistent queues for FTPS file processing Disable VM queue for the batch job scope

Answer: C

NEW QUESTION 56
Which Salesforce API is invoked to deploy, retrieve, create or delete customization information such as custom object definitions using a Mule Salesforce connector in a Mule application?

A. Metadata API
B. REST API
C. SOAP API
D. Bulk API

Answer: B

NEW QUESTION 57
What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

A. Only a DLB allows the configuration of a custom TLS server certificate
B. Only the SLB can forward HTTP traffic to the VPC-internal ports of the CloudHub workers
C. Both a DLB and the SLB allow the configuration of access control via IP whitelists
D. Both a DLB and the SLB implement load balancing by sending HTTP requests to workers with the lowest workloads

Answer: A

Explanation:
* Shared load balancers don't allow you to configure custom SSL certificates or proxy rules
* Dedicated Load Balancer are optional but you need to purchase them additionally if needed.
* TLS is a cryptographic protocol that provides communications security for your Mule app. TLS offers many different ways of exchanging keys for authentication, encrypting data, and guaranteeing message integrity.
* The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
* Only a DLB allows the configuration of a custom TLS server certificate
* DLB enables you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
* To use a DLB in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.
* MuleSoft Reference: https://docs.mulesoft.com/runtime-manager/dedicated-load-balancer-tutorial Additional Info on SLB Vs DLB:
Table Description automatically generated

	Shared Load Balancer	Dedicated Load Balancer
VPC	Shared VPC (Mulesoft)	VPC (Customer)
Default Load Balancer	Cloudhub provides Deault Shared Load Balancer available in All Environment	Need to Purchase
Organization Use	Multiple Orapnization	Specific to Organization
Certificate	Mulesoft Certificate	Organization Certificate
TLS Support	Yes	Yes
URL Mapping	Fixed URL Mapping	Customer URL Mapping
Timeout	30 Sec Session Timeout	Custom Timeout
Ports	Public Port [80 : 8081, 443 : 8082]	Private Port [80 : 8091, 443 : 8092]
Fashion	Round Robin	Round Robin
Supports HTTPS Protocol	Yes	Yes
Worker Assignment	No	Yes
IP Blacklisting/ Whitelisting	No https://docs.mulesoft.com/runtime-manager/ib-whitelists	Yes
Configure Custom Domain	No	Yes
Custom Certificate	No	Yes
Rate Limit	Lower Rate Limit and applied According to Region	Higher Rate Limit Threshold
VPC	Anypoint VPC optional	Can't Use DLB without Anypoint VPC

NEW QUESTION 59
Refer to the exhibit.
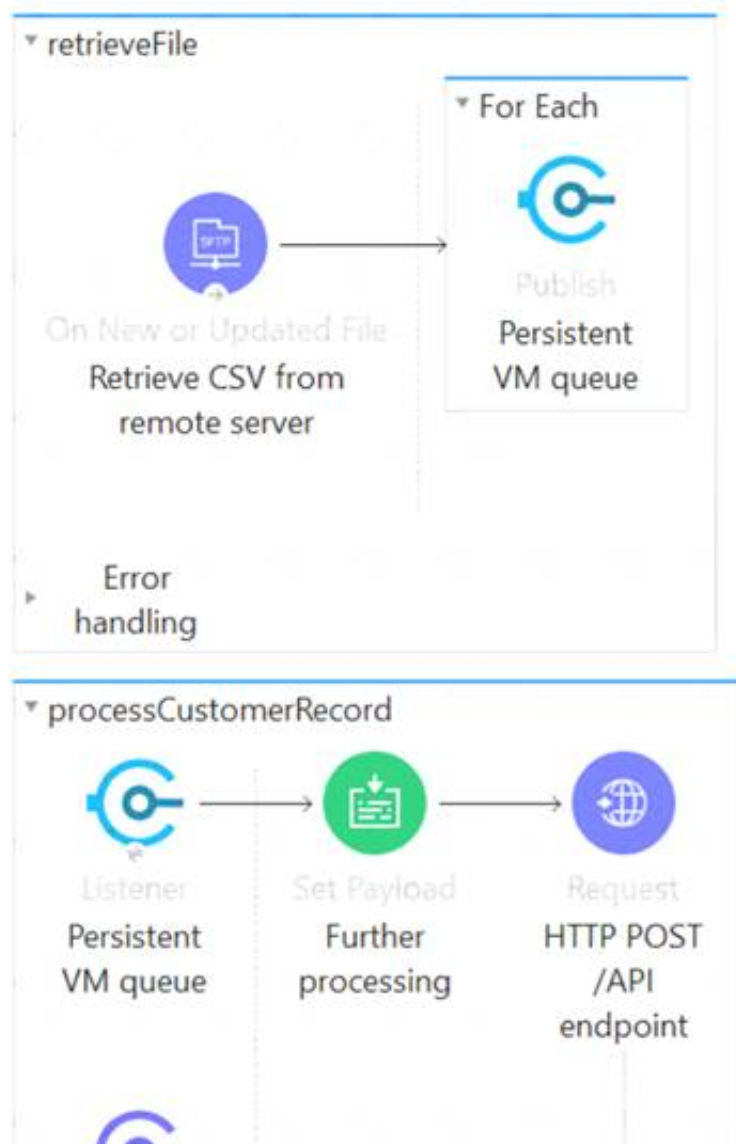
This Mule application is deployed to multiple Cloudhub workers with persistent queue enabled. The retrievefile flow event source reads a CSV file from a remote SFTP server and then publishes each record in the CSV file to a VM queue. The processCustomerRecords flow's VM Listner receives messages from the same VM queue and then processes each message separately.
How are messages routed to the cloudhub workers as messages are received by the VM Listener?

A. Each message is routed to ONE of the Cloudhub workers in a DETERMINSTIC round robin fashion thereby EXACTLY BALANCING messages among the cloudhub workers
B. Each messages routes to ONE of the available Clouhub workers in a NON- DETERMINSTIC non round-robin fashion thereby APPROXIMATELY BALANCING messages among the cloudhub workers
C. Each message is routed to the SAME Cloudhub worker that retrieved the file, thereby BINDING ALLmessages to ONLY that ONE Cloudhub worker
D. Each message is duplicated to ALL of the Cloudhub workers, thereby SHARING EACH message with ALL the Cloudhub workers.

Answer: B

NEW QUESTION 62
As a part of project requirement, Java Invoke static connector in a mule 4 application needs to invoke a static method in a dependency jar file. What are two ways to add the dependency to be visible by the connectors class loader?
(Choose two answers)

A. In the Java Invoke static connector configuration, configure a path and name of the dependency jar file
B. Add the dependency jar file to the java classpath by setting the JVM parameters
C. Use Maven command to include the dependency jar file when packaging the application
D. Configure the dependency as a shared library in the project POM
E. Update mule-artefact.json to export the Java package

Answer: BD

NEW QUESTION 67
Refer to the exhibit.
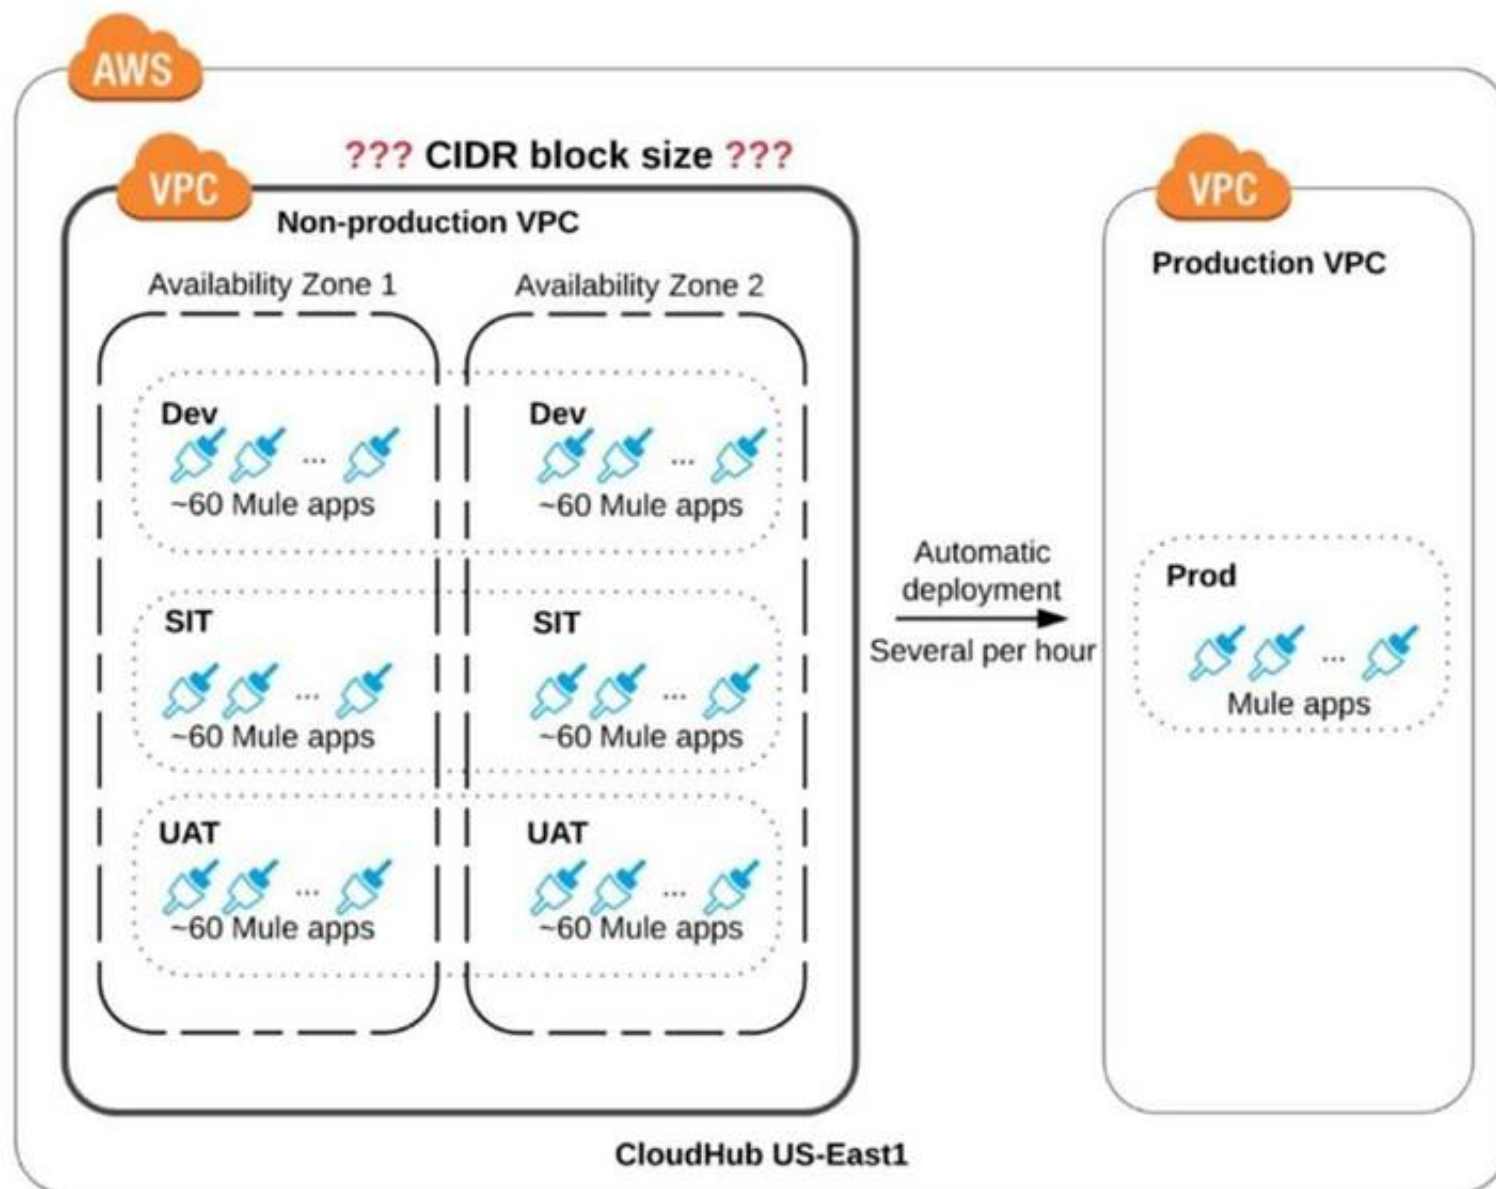
An organization is sizing an Anypoint VPC for the non-production deployments of those Mule applications that connect to the organization's on-premises systems. This applies to approx. 60 Mule applications. Each application is deployed to two CloudHub i workers. The organization currently has three non-production environments (DEV, SIT and UAT) that share this VPC. The AWS region of the VPC has two AZs.

The organization has a very mature DevOps approach which automatically progresses each application through all non-production environments before automatically deploying to production. This process results in several Mule application deployments per hour, using CloudHub's normal zero-downtime deployment feature.

What is a CIDR block for this VPC that results in the smallest usable private IP address range?

A. 10.0.0.0/26 (64 IPS)
B. 10.0.0.0/25 (128 IPs)
C. 10.0.0.0/24 (256 IPs)
D. 10.0.0.0/22 (1024 IPs)

Answer: D

Explanation:
Mule applications are deployed in CloudHub workers and each worker is assigned with a dedicated IP • For zero downtime deployment, each worker in CloudHub needs additional IP addresses • A few IPs in a VPC are reserved for infrastructure (generally 2 IPs) • The IP addresses are usually in a private range with a subnet block specifier, such as 10.0.0.1/24 • The smallest CIDR network subnet block you can assign for your VPC is /24 (256 IP addresses) (60*3 env * 2 worker per application) + 50% of (total) for zero downtime = 540 In this case correct answer is 10.0.0.0/22 as this provided 1024 IP's .
Other IP's are insufficient.

NEW QUESTION 68
An insurance company is using a CIoudHub runtime plane. As a part of requirement, email alert should be sent to internal operations team every time of policy applied to an API instance is deleted As an
integration architect suggest on how this requirement be met?

A. Use audit logs in Anypoint platform to detect a policy deletion and configure the Audit logs alert feature to send an email to the operations team
B. Use Anypoint monitoring to configure an alert that sends an email to the operations team every time a policy is deleted in API manager
C. Create a custom connector to be triggered every time of policy is deleted in API manager
D. Implement a new application that uses the Audit log REST API todetect the policy deletion and send an email to operations team the SMTP connector

Answer: D

NEW QUESTION 71
An organization is designing an integration Mule application to process orders by submitting them to a back-end system for offline processing. Each order will be received by the Mule application through an
HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a back-end system. Orders that cannot be successfully submitted due to rejections from the back-end system will need to be processed manually (outside the back-end system).
The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed. The ActiveMQ broker is located inside the organization's firewall. The back-end system has a track record of unreliability due to both minor network connectivity issues and longer outages.
What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the back-end system while
supporting but minimizing manual order processing?

A. An Until Successful scope to call the back-end system One or more ActiveMQ long-retry queuesOne or more ActiveMQ dead-letter queues for manual

processing
B. One or more On Error scopes to assist calling the back-end systemAn Until Successful scope containing VM components for long retries A persistent dead-letter VM queue configured in CloudHub
C. One or more On Error scopes to assist calling the back-end system One or more ActiveMQ long-retry queues A persistent dead-letter object store configured in the CloudHub Object Store service
D. A Batch Job scope to call the back-end systemAn Until Successful scope containing Object Store components for long retries A dead-letter object store configured in the Mule application

Answer: A

NEW QUESTION 74
Refer to the exhibit.
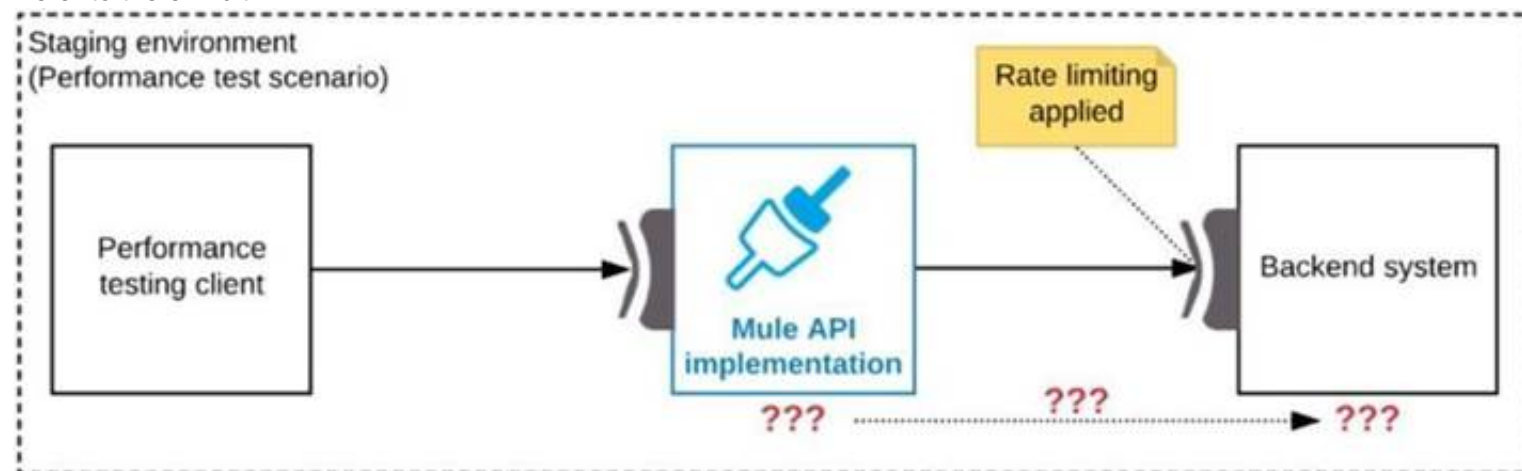
One of the backend systems invoked by an API implementation enforces rate limits on the number of requests a particular client can make. Both the backend system and the API implementation are deployed to several non-production environments in addition to production.
Rate limiting of the backend system applies to all non-production environments. The production environment, however, does NOT have any rate limiting.
What is the most effective approach to conduct performance tests of the API implementation in a staging (non-production) environment?

A. Create a mocking service that replicates the backend system's production performance characteristics.Then configure the API implementation to use the mocking service and conduct the performance tests
B. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
C. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situatio
D. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
E. Conduct scaled-down performance tests in the staging environment against the rate limited backend system then upscale performance results to full production scale

Answer: A

Explanation:
Correct answer is Create a mocking service that replicates the backend system's production performance characteristics. Then configure the API implementation to use the mocking service and conduct the performance tests
* MUnit is for only Unit and integration testing for APIs and Mule apps. Not for performance Testing, even if it has the ability to Mock the backend.
* Bypassing the backend invocation defeats the whole purpose of performance testing. Hence it is not a valid answer.
* Scaled down performance tests cant be relied upon as performance of API's is not linear against load.

NEW QUESTION 76
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?
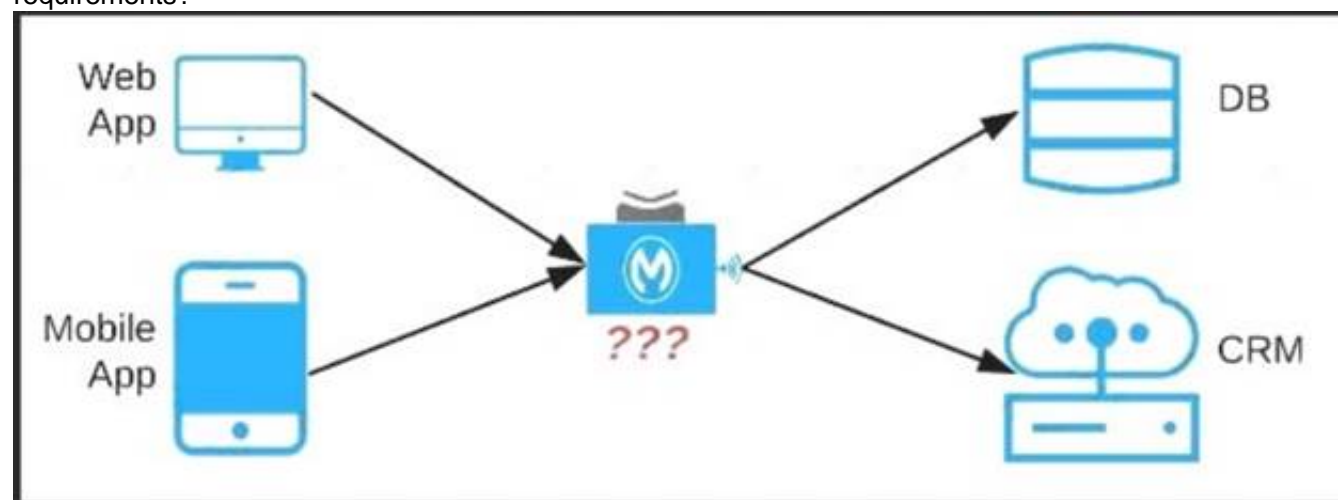
A. A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
B. One set of APIs (Experience API, Process API, and System API) for the web app, and another set forthe 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.

Answer: C

Explanation:
Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot

be fulfilled by single experience layer API

Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.

Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.

Hence the correct answer is: 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

A screenshot of a computer Description automatically generated with low confidence

Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web applications will have different set of requirements which cannot be fulfilled by single experience layer API
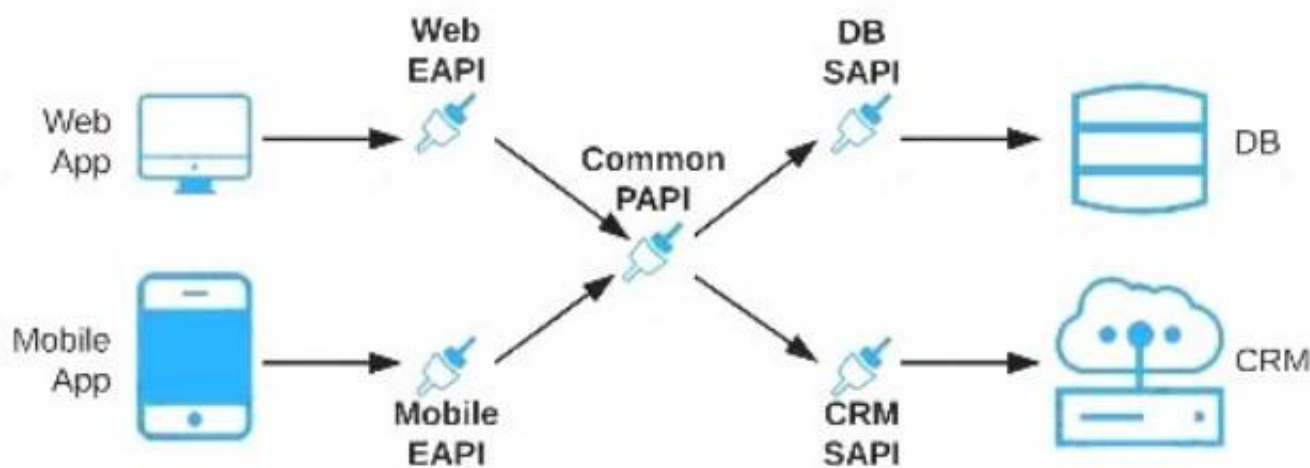
Option : Common Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.

Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates duplicity of code.

Hence the correct answer is: 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
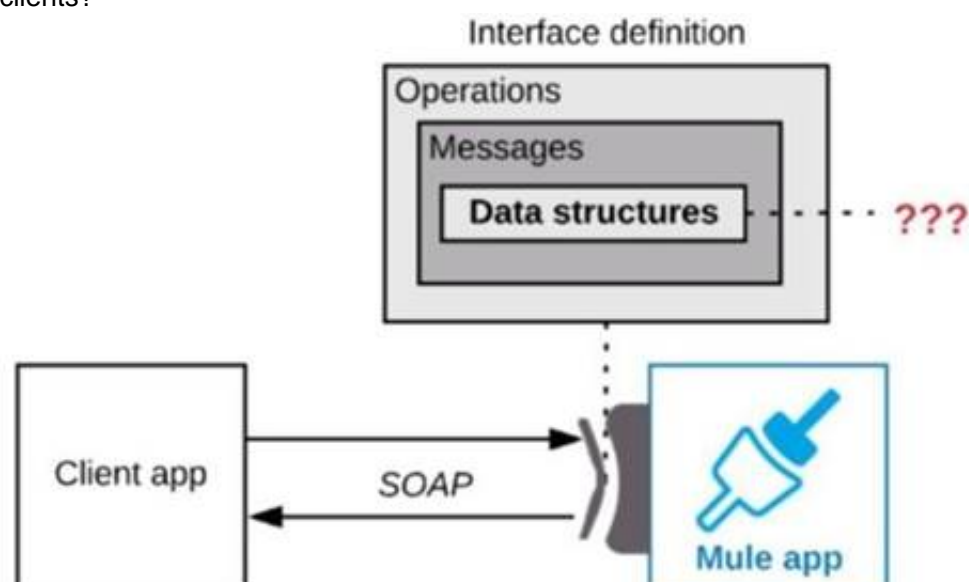
Diagram Description automatically generated
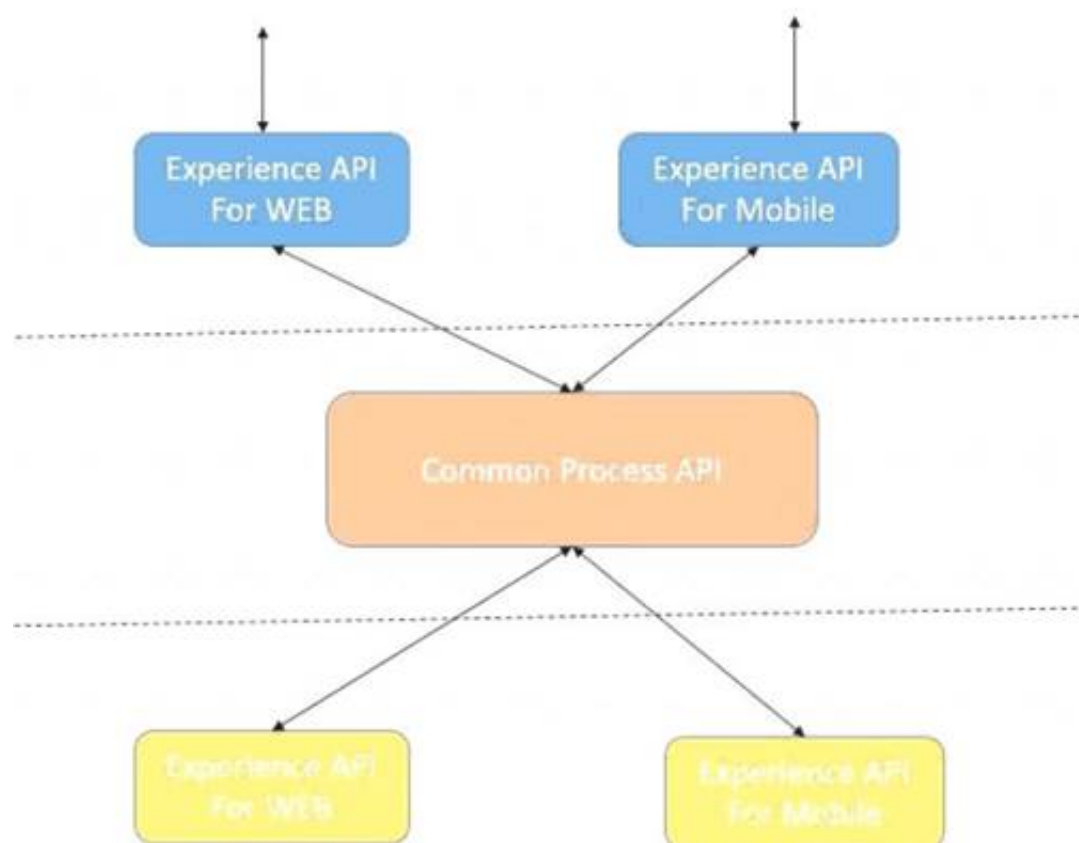
NEW QUESTION 79

Refer to the exhibit.

A Mule application is being designed to expose a SOAP web service to its clients.

What language is typically used inside the web service's interface definition to define the data structures that the web service Is expected to exchange with its clients?

A. WSDL
B. XSD
C. JSON Schema
D. RAMI

Answer: B

Explanation:
Correct Answer XSD In this approach to developing a web service, you begin with an XML schema (XSD file) that defines XML data structures to be used as parameters and return types in the web service operations.
-- Reference:
https://www.w3schools.com/xml/schema_intro.asp

NEW QUESTION 80
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. Apply a Header injection and removal policy that detects the malicious data before it is used
B. Apply an IP blacklist policy to all APIs; the blacklist will Include all bad actors
C. Shut out bad actors by using HTTPS mutual authentication for all API invocations
D. Apply a JSON threat protection policy to all APIs to detect potential threat vectors

Answer: D

Explanation:
We need to note few things about the scenario which will help us in reaching the correct solution.
Point 1 : The APIs are all publicly available and are associated with several mobile applications and web applications. This means Apply an IP blacklist policy is not viable option. as blacklisting IPs is limited to partial web traffic. It can't be useful for traffic from mobile application
Point 2 : The organization does NOT want to use any authentication or compliance policies for these APIs. This means we can not apply HTTPS mutual authentication scheme.
Header injection or removal will not help the purpose.
By its nature, JSON is vulnerable to JavaScript injection. When you parse the JSON object, the malicious code inflicts its damages. An inordinate increase in the size and depth of the JSON payload can indicate injection. Applying the JSON threat protection policy can limit the size of your JSON payload and thwart recursive additions to the JSON hierarchy.
Hence correct answer is Apply a JSON threat protection policy to all APIs to detect potential threat vectors

NEW QUESTION 84
A company is planning to extend its Mule APIs to the Europe region. Currently all new applications are deployed to Cloudhub in the US region following this naming convention
{API name}-{environment}. for example, Orders-SAPI-dev, Orders-SAPI-prod etc.
Considering there is no network restriction to block communications between API's, what strategy should be implemented in order to apply the same new API's running in the EU region of CloudHub as well to minimize latency between API's and target users and systems in Europe?

A. Set region property to Europe (eu-de) in API manager for all the mule application No need to change the naming convention
B. Set region property to Europe (eu-de) in API manager for all the mule applicationChange the naming convention to {API name}-{environment}-{region} and communicate this change to the consuming applications and users
C. Set region property to Europe (eu-de) in runtime manager for all the mule application No need to change the naming convention
D. Set region property to Europe (eu-de) in runtime manager for all the mule applicationChange the naming convention to {API name}-{environment}-{region} and communicate this change to the consuming applications and users

Answer: D

NEW QUESTION 86
An organization currently uses a multi-node Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications. The organization is planning to transition to a deployment model based on Docker containers in a Kubernetes cluster. The organization has already created a standard Docker image containing a Mule runtime and all required dependencies (including a JVM), but excluding the Mule application itself.
What is an expected outcome of this transition to container-based Mule application deployments?

A. Required redesign of Mule applications to follow microservice architecture principles
B. Required migration to the Docker and Kubernetes-based Anypoint Platform - Private Cloud Edition
C. Required change to the URL endpoints used by clients to send requests to the Mule applications
D. Guaranteed consistency of execution environments across all deployments of a Mule application

Answer: A

Explanation:
* Organization can continue using existing load balancer even if backend application changes are there. So option A is ruled out.
* As Mule runtime is within their datacenter, this model is RTF and not PCE. So option C is ruled out.
Mule runtime deployment model within their datacenter, so each Mule runtime hosts several Mule applications
-- This mean PCE or Hybird not RTF - Also mentioned in Question is that - Mule runtime is hosting several Mule Application, so that also rules out RTF and as for hosting multiple Application it will have Domain project which need redesign to make it microservice architecture
----
Correct Answer Required redesign of Mule applications to follow microservice architecture principles

NEW QUESTION 90
A leading bank implementing new mule API.
The purpose of API to fetch the customer account balances from the backend application and display them on the online platform the online banking platform. The

online banking platform will send an array of accounts to Mule API get the account balances.
As a part of the processing the Mule API needs to insert the data into the database for auditing purposes and this process should not have any performance related implications on the account balance retrieval flow
How should this requirement be implemented to achieve better throughput?

A. Implement the Async scope fetch the data from the backend application and to insert records in the Audit database
B. Implement a for each scope to fetch the data from the back-end application and to insert records into the Audit database
C. Implement a try-catch scope to fetch the data from the back-end application and use the Async scope to insert records into the Audit database
D. Implement parallel for each scope to fetch the data from the backend application and use Async scope to insert the records into the Audit database

Answer: D

NEW QUESTION 93
A Mule application name Pub uses a persistence object store. The Pub Mule application is deployed to Cloudhub and it configured to use Object Store v2.
Another Mule application name sub is being developed to retrieve values from the Pub Mule application persistence object Store and will also be deployed to cloudhub.
What is the most direct way for the Sub Mule application to retrieve values from the Pub Mule application persistence object store with the least latency?

A. Use an object store connector configured to access the Pub Mule application persistence object store
B. Use a VM connector configured to directly access the persistence queue of the Pub Mule application persistence object store.
C. Use an Anypoint MQ connector configured to directly access the Pub Mule application persistence object store
D. Use the Object store v2 REST API configured to access the Pub Mule application persistence object store.

Answer: D

Explanation:
* The Object Store V2 API enables API access to Anypoint Platform Object Store v2.
* You can configure a Mule app to use the Object Store REST API to store and retrieve values from an object store in another Mule app. However, Object Store v2 is not designed for app-to-app communication. To share data between two Mule4 apps, use a queue in Anypoint MQ.
* The Object Store v2 APIs enable you to use REST to perform the following:
- Retrieve a list of object stores and keys associated with an application.
- Store and retrieve key-value pairs in an object store.
- Delete key-value pairs from an object store.
- Retrieve Object Store usage statistics for your organization.
- Object Store provides these APIs: Object Store API
Object Store Stats API

NEW QUESTION 96
When using Anypoint Platform across various lines of business with their own Anypoint Platform business groups, what configuration of Anypoint Platform is always performed at the organization level as opposed to at the business group level?

A. Environment setup
B. Identity management setup
C. Role and permission setup
D. Dedicated Load Balancer setup

Answer: B

Explanation:
* Roles are business group specific. Configure identity management in the Anypoint Platform master organization. As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). * Roles and permissions can be set up at business group and organization level also. But Identity Management setup is only done at Organization level * Business groups are self-contained resource groups that contain Anypoint Platform resources such as applications and APIs. Business groups provide a way to separate and control access to Anypoint Platform resources because users have access only to the busine

NEW QUESTION 101
An integration Mute application consumes and processes a list of rows from a CSV file. Each row must be read from the CSV file, validated, and the row data sent to a JMS queue, in the exact order as in the CSV file.
If any processing step for a row falls, then a log entry must be written for that row, but processing of other rows must not be affected.
What combination of Mute components is most idiomatic (used according to their intended purpose) when Implementing the above requirements?

A. Scatter-Gather component On Error Continue scope
B. VM connector first Successful scope On Error Propagate scope
C. For Each scope On Error Continue scope
D. Async scope On Error Propagate scope

Answer: C

Explanation:
* On Error Propagate halts execution and sends error to the client. In this scenario it's mentioned that "processing of other rows must not be affected" so Option B and C are ruled out.
* Scatter gather is used to club multiple responses together before processing. In this scenario, we need sequential processing. So option A is out of choice.
* Correct answer is For Each scope & On Error Continue scope Below requirement can be fulfilled in the below way
1) Using For Each scope , which will send each row from csv file sequentially. each row needs to be sent sequentially as requirement is to send the message in exactly the same way as it is mentioned in the csv file
2) Also other part of requirement is if any processing step for a row fails then it should log an error but should not affect other record processing . This can be achieved using On error Continue scope on these set of activities. so that error will not halt the processing. Also logger needs to be added in error handling section so that it can be logged.
* Attaching diagram for reference. Here it's try scope, but similar would be the case with For Each loop. Diagram Description automatically generated
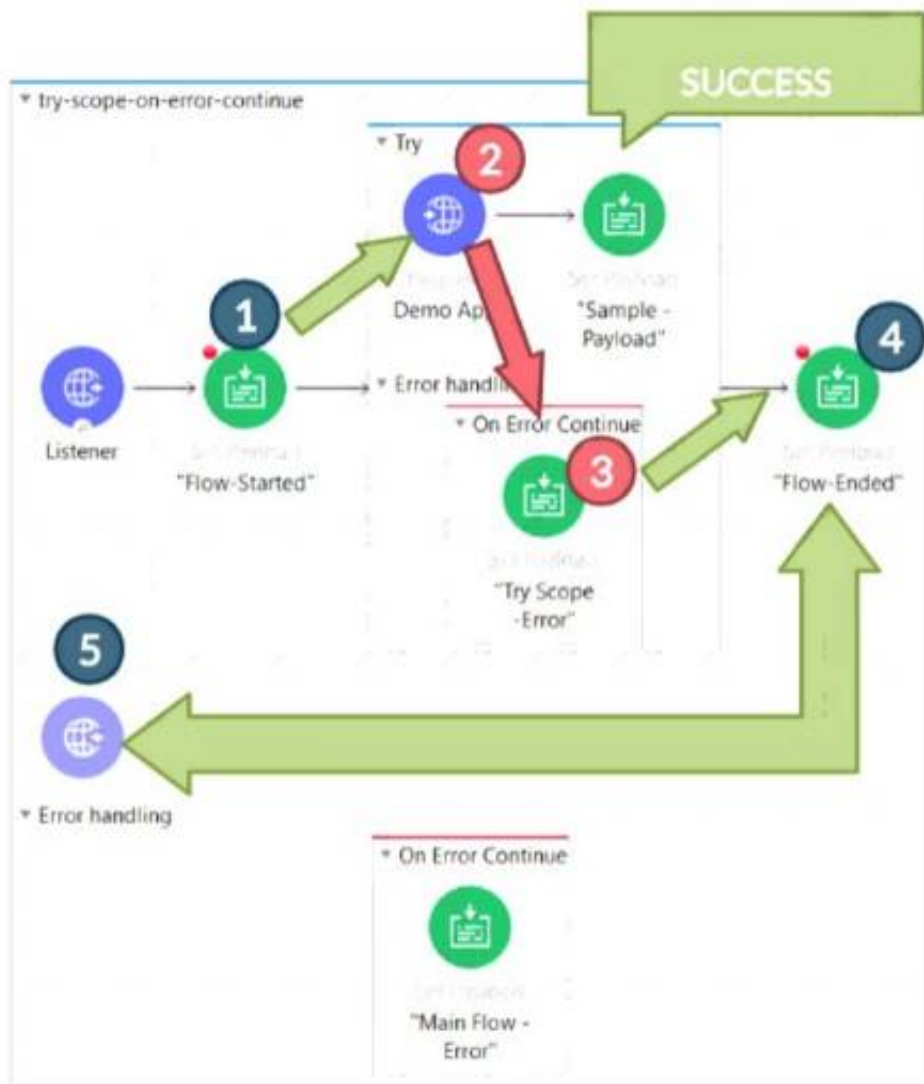
NEW QUESTION 105

An Order microservice and a Fulfillment microservice are being designed to communicate with their dients through message-based integration (and NOT through API invocations).

The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mute application, the Fulfillment microservice.

The Fulfilment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilted message (a kind of event message). Each OrderFulfilted message can be consumed by any interested Mule application, and the Order microservice is one such Mute application.

What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

A. Order messages are sent to an Anypoint MQ exchange OrderFulfilled messages are sent to an Anypoint MQ queue Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices

B. Order messages are sent to a JMS queu

C. OrderFulfilled messages are sent to a JMS topic Both microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices

D. Order messages are sent directly to the Fulfillment microservice

E. OrderFulfilled messages are sent directly to the Order microservice The Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load of each microservice

F. Order messages are sent to a JMS queu

G. OrderFulfilled messages are sent to a JMS topic The Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice

Answer: B

Explanation:

* If you need to scale a JMS provider/ message broker, - add nodes to scale it horizontally or - add memory to scale it vertically * Cons of adding another JMS provider/ message broker: - adds cost. - adds complexity to use two JMS brokers - adds Operational overhead if we use two brokers, say, ActiveMQ and IBM MQ * So Two options that mention to use two brokers are not best choice. * It's mentioned that "The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message. Each OrderFulfilled message can be consumed by any interested Mule application." - When you publish a message on a topic, it goes to all the subscribers who are interested - so zero to many subscribers will receive a copy of the message. - When you send a message on a queue, it will be received by exactly one consumer. * As we need multiple consumers to consume the message below option is not valid choice: "Order messages are sent to an Anypoint MQ exchange. OrderFulfilled messages are sent to an Anypoint MQ queue. Both microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices" * Order messages are only consumed by one Mule application, the Fulfillment microservice, so we will publish it on queue and OrderFulfilled message can be consumed by any interested Mule application so it need to be published on Topic using same broker. * Correct Answer Best choice in this scenario is: "Order messages are sent to a JMS queue. OrderFulfilled messages are sent to a JMS topic. Both microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices" Tried to depict scenario in diagram:
Diagram Description automatically generated
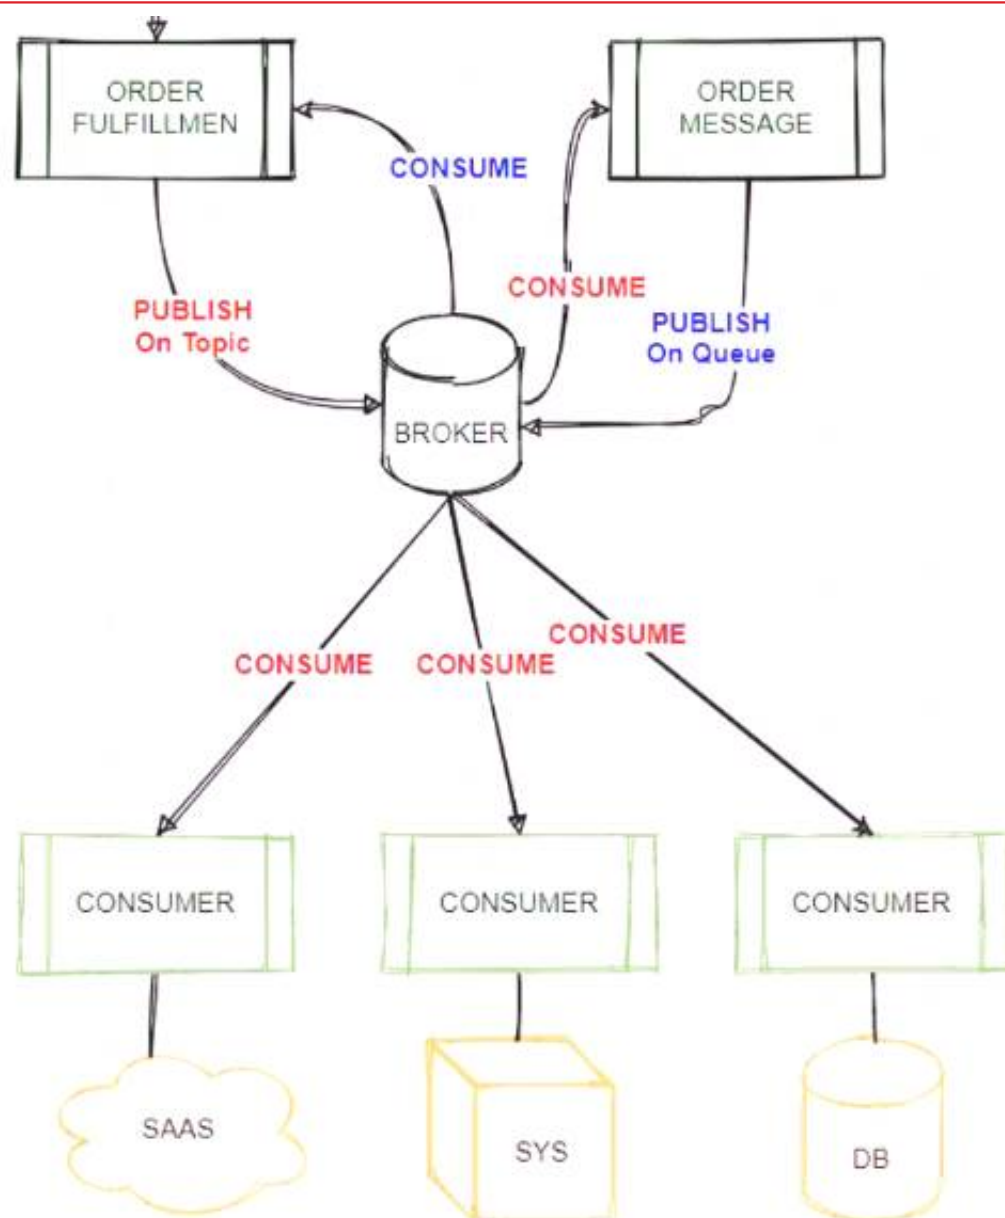
NEW QUESTION 108
An organization if struggling frequent plugin version upgrades and external plugin project dependencies. The team wants to minimize the impact on applications by creating best practices that will define a set of default dependencies across all new and in progress projects.
How can these best practices be achieved with the applications having the least amount of responsibility?

A. Create a Mule plugin project with all the dependencies and add it as a dependency in each application's POM.xml file
B. Create a mule domain project with all the dependencies define in its POM.xml file and add each application to the domain Project
C. Add all dependencies in each application's POM.xml file
D. Create a parent POM of all the required dependencies and reference each in each application's POM.xml file

Answer: D

NEW QUESTION 109
Mule application A receives a request Anypoint MQ message REQU with a payload containing a
variable-length list of request objects. Application A uses the For Each scope to split the list into individual objects and sends each object as a message to an Anypoint MQ queue.
Service S listens on that queue, processes each message independently of all other messages, and sends a response message to a response queue.
Application A listens on that response queue and must in turn create and publish a response Anypoint MQ message RESP with a payload containing the list of responses sent by service S in the same order as the request objects originally sent in REQU.
Assume successful response messages are returned by service S for all request messages.
What is required so that application A can ensure that the length and order of the list of objects in RESP and REQU match, while at the same time maximizing message throughput?

A. Use a Scatter-Gather within the For Each scope to ensure response message order Configure the Scatter-Gather with a persistent object store
B. Perform all communication involving service S synchronously from within the For Each scope, so objects in RESP are in the exact same order as request objects in REQU
C. Use an Async scope within the For Each scope and collect response messages in a second For Each scope in the order In which they arrive, then send RESP using this list of responses
D. Keep track of the list length and all object indices in REQU, both in the For Each scope and in all communication involving service Use persistent storage when creating RESP

Answer: D

Explanation:
: Using Anypoint MQ, you can create two types of queues: Standard queue These queues don't guarantee a specific message order. Standard queues are the best fit for applications in which messages must be delivered quickly. FIFO (first in, first out) queue These queues ensure that your messages arrive in order. FIFO queues are the best fit for applications requiring strict message ordering and exactly-once delivery, but in which message
delivery speed is of less importance Use of FIFO queue is no where in the option and also it decreased throughput. Similarly persistent object store is not the preferred solution approach when you maximizing message throughput. This rules out one of the options. Scatter Gather does not support ObjectStore. This rules out one of the options. Standard Anypoint MQ queues don't guarantee a specific message order hence using another for each block to collect response wont work as requirement here is to ensure the order. Hence considering all the above factors the feasible approach is Perform all communication involving service S synchronously from within the For Each scope, so objects in RESP are in the exact same order as request objects in REQU

NEW QUESTION 110
An auto mobile company want to share inventory updates with dealers DI and D2 asynchronously and concurrently via queues Q1 and Q2. Dealer DI must consume the message from the queue Q1 and dealer D2 to must consume a message from the queue Q2.
Dealer D1 has implemented a retry mechanism to reprocess the transaction in case of any errors while processing the inventers updates. Dealer D2 has not implemented any retry mechanism.
How should the dealers acknowledge the message to avoid message loss and minimize impact on the current implementation?

A. Dealer D1 must use auto acknowledgement and dealer D2 can use manual acknowledgement and acknowledge the message after successful processing
B. Dealer D1 can use auto acknowledgement and dealer D2 can use IMMEDIATE acknowledgement and acknowledge the message of successful processing
C. Dealer D1 and dealer D2 must use AUTO acknowledgement and acknowledge the message after successful processing
D. Dealer D1 can use AUTO acknowledgement and dealer D2 must use manual acknowledgement and acknowledge the message after successful processing

Answer: D

NEW QUESTION 112
Which of the below requirements prevent the usage of Anypoint MQ in a company's network? (Choose two answers)

A. single message payload can be up to 15 MB
B. payloads must be encrypted
C. the message broker must be hosted on premises
D. support for point-to-point messaging
E. ability for a third party outside the company's network to consume events from the queue

Answer: CD

NEW QUESTION 114
A Mule application is synchronizing customer data between two different database systems.
What is the main benefit of using eXtended Architecture (XA) transactions over local transactions to synchronize these two different database systems?

A. An XA transaction synchronizes the database systems with the least amount of Mule configuration or coding
B. An XA transaction handles the largest number of requests in the shortest time
C. An XA transaction automatically rolls back operations against both database systems if any operation falls
D. An XA transaction writes to both database systems as fast as possible

Answer: B

NEW QUESTION 116
An organization has decided on a cloud migration strategy to minimize the organization's own IT resources. Currently the organization has all of its new applications running on its own premises and uses an on-premises load balancer that exposes all APIs under the base URL (https://api.rutujar.com).
As part of migration strategy, the organization is planning to migrate all of its new applications and load balancer CloudHub.
What is the most straightforward and cost-effective approach to Mule application deployment and load balancing that preserves the public URL's?

A. Deploy the Mule application to CloudhubCreate a CNAME record for base URL(httpsr://api.rutujar.com) in the Cloudhub shared load balancer that points to the A record of theon-premises load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications
B. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL (https://api.rutujar.com) in the organization's DNS server to point to the A record of the Cloudhub dedicated load balancerApply mapping rules in DLB to map URLto their corresponding Mule applications
C. Deploy the Mule application to CloudhubUpdate a CNAME record for base URL (https://api.rutujar.com) in the organization's DNS server to point to the A record of the CloudHub shared load balancerApply mapping rules in SLB to map URLto their corresponding Mule applications
D. For each migrated Mule application, deploy an API proxy application to Cloudhub with all traffic to themule applications routed through a Cloud Hub Dedicated load balancer (DLB)Update a CNAME record for base URL (https://api.rutujar.com) in the organization's DNS server to point to the A record of the CloudHub dedicated load balancerApply mapping rules in DLB to map each API proxy application who is responding new application

Answer: B

NEW QUESTION 118
How does timeout attribute help inform design decisions while using JMS connector listening for incoming messages in an extended architecture (XA) transaction?

A. After the timeout is exceeded, stale JMS consumer threads are destroyed and new threads are created
B. The timeout specifies the time allowed to pass between receiving JMS messages on the same JMS connection and then after the timeout new JMS connection is established
C. The time allowed to pass between committing the transaction and the completion of the mule flow and then after the timeout flow processing triggers an error
D. The timeout defines the time that is allowed to pass without the transaction ending explicitly and after the timeout expires, the transaction rolls back

Answer: D

NEW QUESTION 123
An organization is evaluating using the CloudHub shared Load Balancer (SLB) vs creating a CloudHub dedicated load balancer (DLB). They are evaluating how this choice affects the various types of certificates used by CloudHub deplpoyed Mule applications, including MuleSoft-provided, customer-provided, or Mule application-provided certificates.
What type of restrictions exist on the types of certificates that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

A. Only MuleSoft-provided certificates are exposed.
B. Only customer-provided wildcard certificates are exposed.
C. Only customer-provided self-signed certificates are exposed.
D. Only underlying Mule application certificates are exposed (pass-through)

Answer: A

Explanation:
https://docs.mulesoft.com/runtime-manager/dedicated-load-balancer-tutorial

NEW QUESTION 127
An organization's governance process requires project teams to get formal approval from all key stakeholders for all new Integration design specifications. An integration Mule application Is being designed that interacts with various backend systems. The Mule application will be created using Anypoint Design Center or Anypoint Studio and will then be deployed to a customer-hosted runtime.
What key elements should be included in the integration design specification when requesting approval for this Mule application?

A. SLAs and non-functional requirements to access the backend systems
B. Snapshots of the Mule application's flows, including their error handling
C. A list of current and future consumers of the Mule application and their contact details
D. The credentials to access the backend systems and contact details for the administrator of each system

Answer: A

Explanation:
SLAs and non-functional requirements to access the backend systems. Only this option actually speaks to
design parameters and reqs. * Below two are technical implementations and not the part of design: - Snapsho of the Mule application's flows, including their error handling - The credentials to access the backend systems and contact details for the administrator of each system * List of consumers is not relevant to the design

NEW QUESTION 132
An organization has deployed both Mule and non-Mule API implementations to integrate its customer and order management systems. All the APIs are available to REST clients on the public internet.
The organization wants to monitor these APIs by running health checks: for example, to determine if an API can properly accept and process requests. The organization does not have subscriptions to any external monitoring tools and also does not want to extend its IT footprint.
What Anypoint Platform feature provides the most idiomatic (used for its intended purpose) way to monitor the availability of both the Mule and the non-Mule API implementations?

A. API Functional Monitoring
B. Runtime Manager
C. API Manager
D. Anypoint Visualizer

Answer: D

NEW QUESTION 134
Which Mulesoft feature helps users to delegate their access without sharing sensitive credentials or giving full control of accounts to 3rd parties?

A. Secure Scheme
B. client id enforcement policy
C. Connected apps
D. Certificates

Answer: C

Explanation:
Connected Apps
The Connected Apps feature provides a framework that enables an external application to integrate with Anypoint Platform using APIs through OAuth 2.0 and OpenID Connect. Connected apps help users delegate their access without sharing sensitive credentials or giving full control of their accounts to third parties. Actions taken by connected apps are audited, and users can also revoke access at any time. Note that some products do not currently include client IDs in this release of the Connected Apps feature. The Connected Apps feature enables you to use secure authentication protocols and control an app's access to user data. Additionally, end users can authorize the app to access their Anypoint Platform data.
Mule Ref Doc : https://docs.mulesoft.com/access-management/connected-apps-overview

NEW QUESTION 136
An organization uses a four(4) node customer hosted Mule runtime cluster to host one(1) stateless api implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution. Each node in the cluster has been sized to be able to accept four(4) times the current number of requests.
Two(2) nodes in the cluster experience a power outage and are no longer available. The load balancer directs the outage and blocks the two unavailable the nodes from receiving further HTTP requests.
What performance-related consequence is guaranteed to happen to average, assuming the remaining cluster nodes are fully operational?

A. 100% increase in the average response time of the API
B. 50% reduction in the throughput of the API
C. 100% increase in the number of requests received by each remaining node
D. 50% increase in the JVM heap memory consumed by each remaining node

Answer: C

Explanation:
* "100% increase in the throughput of the API" might look correct, as the number of requests processed per second might increase, but is it guaranteed to increase by 100%? Using 4 nodes will definitely increase throughput of system. But it is cant be precisely said if there would be 100% increase in throughput as it depends on many other factors. Also it is nowhere mentioned in the description that all nodes have same CPU/memory assigned. The question is about the guaranteed behavior * Increasing number of nodes will have no impact on response time as we are scaling application horizontally and not vertically. Similarly there is no change in JVM heap memory usage. * So Correct answer is 50% reduction in the number of requests being received by each node This is because of the two reasons. 1) API is mentioned as stateless 2) Load Balancer is used

NEW QUESTION 140

An organization is creating a set of new services that are critical for their business. The project team prefers using REST for all services but is willing to use SOAP with common WS-" standards if a particular service requires it.

What requirement would drive the team to use SOAP/WS-* for a particular service?

A. Must use XML payloads for the service and ensure that it adheres to a specific schema
B. Must publish and share the service specification (including data formats) with the consumers of the service
C. Must support message acknowledgement and retry as part of the protocol
D. Must secure the service, requiring all consumers to submit a valid SAML token

Answer: D

Explanation:

Security Assertion Markup Language (SAML) is an open standard that allows identity providers (IdP) to pass authorization credentials to service providers (SP). SAML transactions use Extensible Markup Language (XML) for standardized communications between the identity provider and service providers.
SAML is the link between the authentication of a user's identity and the authorization to use a service. WS-Security is the key extension that supports many authentication models including: basic
username/password credentials, SAML, OAuth and more.
A common way that SOAP API's are authenticated is via SAML Single Sign On (SSO). SAML works by facilitating the exchange of authentication and authorization credentials across applications. However, there is no specification that describes how to add SAML to REST web services.
Reference: https://www.oasis-open.org/committees/download.php/16768/wss-v1.1-spec-os-SAMLTokenProfile.pdf

NEW QUESTION 144

An organization is implementing a Quote of the Day API that caches today's quote. What scenario can use the CloudHub Object Store connector to persist the cache's state?

A. When there is one deployment of the API implementation to CloudHub and another one to customer hosted mule runtime 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 CloudHub deployment of the API implementation to three workers that must share the cache state.
D. When there are three CloudHub deployments of the API implementation to three separate CloudHub regions that must share the cache state.

Answer: C

Explanation:

Object Store Connector is a Mule component that allows for simple key-value storage. Although it can serve a wide variety of use cases, it is mainly design for: - Storing synchronization information, such as watermarks. - Storing temporal information such as access tokens. - Storing user information. Additionally, Mule Runtime uses Object Stores to support some of its own components, for example: - The Cache module uses an Object Store to maintain all of the cached data. - The OAuth module (and every OAuth enabled connector) uses Object Stores to store the access and refresh tokens. Object Store data is in the same region as the worker where the app is initially deployed. For example, if you deploy to the Singapore region, the object store persists in the Singapore region. MuleSoft Reference : https://docs.mulesoft.com/object-store-connector/1.1/ Data can be shared between different instances of the Mule application. This is not recommended for Inter Mule app communication. Coming to the question, object store cannot be used to share cached data if it is deployed as separate Mule applications or deployed under separate Business Groups. Hence correct answer is When there is one CloudHub deployment of the API implementation to three workers that must share the cache state.

NEW QUESTION 146

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

Explanation:

Correct answer is When server-side load-balanced TLS mutual authentication is required between API implementations and API clients CloudHub dedicated load balancers (DLBs) are an optional component of Anypoint Platform that enable you to route external HTTP and HTTPS traffic to multiple Mule applications deployed to CloudHub workers in a Virtual Private Cloud (VPC). Dedicated load balancers enable you to: * Handle load balancing among the different CloudHub workers that run your application. * Define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication. * Configure proxy rules that map your applications to custom domains. This enables you to host your applications under a single domain

NEW QUESTION 148

A manufacturing company is planning to deploy Mule applications to its own Azure Kubernetes Service infrastructure.
The organization wants to make the Mule applications more available and robust by deploying each Mule application to an isolated Mule runtime in a Docker container while managing all the Mule applications from the MuleSoft-hosted control plane.
What is the most idiomatic (used for its intended purpose) choice of runtime plane to meet these organizational requirements?

A. Anypoint Platform Private Cloud Edition
B. Anypoint Runtime Fabric
C. CloudHub
D. Anypoint Service Mesh

Answer: B

NEW QUESTION 150

A project uses Jenkins to implement CI/CD process. It was observed that each Mule package contains some of the Jenkins files and folders for configurations of CI/CD jobs.
As these files and folders are not part of the actual package, expectation is that these should not be part of deployed archive.
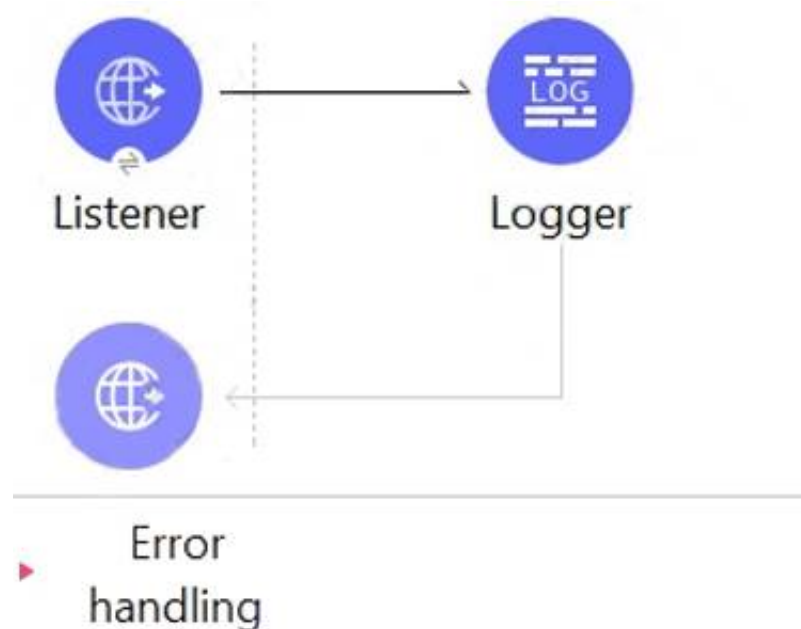
Which file can be used to exclude these files and folders from the deployed archive?

A. muleignore
B. _unTrackMule
C. muleInclude
D. _muleExclude

Answer: D

NEW QUESTION 153
Refer to the exhibit.

The HTTP Listener and the Logger are being handled from which thread pools respectively?

A. CPU_INTENSIVE and Dedicated Selector pool
B. UBER and NONBLOCKING
C. Shared Selector Pool and CPU LITE
D. BLOCKING _IO and UBER

Answer: C

NEW QUESTION 156
A mule application must periodically process a large dataset which varies from 6 GB lo 8 GB from a back-end database and write transform data lo an FTPS server using a properly configured bad job scope.
The performance requirements of an application are approved to run in the cloud hub 0.2 vCore with 8 GB storage capacity and currency requirements are met. How can the high rate of records be effectively managed in this application?

A. Use streaming with a file storage repeatable strategy for reading records from the database and batch aggregator with streaming to write to FTPS
B. Use streaming with an in-memory reputable store strategy for reading records from the database and batch aggregator with streaming to write to FTPS
C. Use streaming with a file store repeatable strategy for reading records from the database and batch aggregator with an optimal size
D. Use streaming with a file store repeatable strategy reading records from the database and batch aggregator without any required configuration

Answer: A

NEW QUESTION 157
A Mule application uses the Database connector.
What condition can the Mule application automatically adjust to or recover from without needing to restart or redeploy the Mule application?

A. One of the stored procedures being called by the Mule application has been renamed
B. The database server was unavailable for four hours due to a major outage but is now fully operational again
C. The credentials for accessing the database have been updated and the previous credentials are no longer valid
D. The database server has been updated and hence the database driver library/JAR needs a minor version upgrade

Answer: B

Explanation:
* Any change in the application will require a restart except when the issue outside the app. For below situations , you would need to redeploy the code after doing necessary changes
-- One of the stored procedures being called by the Mule application has been renamed. In this case, in the Mule application you will have to do changes to accommodate the new stored procedure name.
-- Required redesign of Mule applications to follow microservice architecture principles. As code is changed, deployment is must
-- If the credentials changed and you need to update the connector or the properties.
-- The credentials for accessing the database have been updated and the previous credentials are no longer valid. In this situation you need to restart or redeploy depending on how credentials are configured in Mule application.
* So Correct answer is The database server was unavailable for four hours due to a major outage but is now fully operational again as this is the only external issue to application.

NEW QUESTION 161
The ABC company has an Anypoint Runtime Fabric on VMs/Bare Metal (RTF-VM) appliance installed on its own customer-hosted AWS infrastructure.
Mule applications are deployed to this RTF-VM appliance. As part of the company standards, the Mule application logs must be forwarded to an external log management tool (LMT).

Given the company's current setup and requirements, what is the most idiomatic (used for its intended purpose) way to send Mule application logs to the external LMT?

A. In RTF-VM, install and configure the external LTM's log-forwarding agent
B. In RTF-VM, edit the pod configuration to automatically install and configure an Anypoint Monitoring agent
C. In each Mule application, configure custom Log4j settings
D. In RTF-V
E. configure the out-of-the-box external log forwarder

Answer: A

NEW QUESTION 164
An organization's security policies mandate complete control of the login credentials used to log in to Anypoint Platform. What feature of Anypoint Platform should be used to meet this requirement?

A. Enterprise Security Module
B. Client ID Secret
C. Federated Identity Management
D. Federated Client Management

Answer: C

Explanation:
Correct answer is Federated Identity Management As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO). Configure identity management using one of the following single sign-on standards: OpenID Connect: End user identity verification by an authorization server including SSO SAML 2.0: Web-based authorization including
cross-domain SSO Where as Client Management is where Anypoint Platform acts as a client provider by default, but you can also configure external client providers to authorize client applications. As an API owner, you can apply an OAuth 2.0 policy to authorize client applications that try to access your API. You need an OAuth 2.0 provider to use an OAuth 2.0 policy

NEW QUESTION 166
The AnyAirline organization's passenger reservations center is designing an integration solution that combines invocations of three different System APIs (bookFlight, bookHotel, and bookCar) in a business transaction. Each System API makes calls to a single database.
The entire business transaction must be rolled back when at least one of the APIs fails.
What is the most idiomatic (used for its intended purpose) way to integrate these APIs in near real-time that provides the best balance of consistency, performance, and reliability?

A. Implement eXtended Architecture (XA) transactions between the API implementations Coordinate between the API implementations using a Saga patternImplement caching in each API implementation to improve performance
B. Implement local transactions within each API implementationConfigure each API implementation to also participate in the same eXtended Architecture (XA) transactionImplement caching in each API implementation to improve performance
C. Implement local transactions in each API implementation Coordinate between the API implementations using a Saga patternApply various compensating actions depending on where a failure occurs
D. Implement an eXtended Architecture (XA) transaction manager in a Mule application using a Saga patternConnect each API implementation with the Mule application using XA transactions Apply various compensating actions depending on where a failure occurs

Answer: C

NEW QUESTION 169
What API policy would LEAST likely be applied to a Process API?

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

Answer: D

Explanation:
Key to this question lies in the fact that Process API are not meant to be accessed directly by clients. Lets analyze options one by one. Client ID enforcement : This is applied at process API level generally to ensure that identity of API clients is always known and available for API-based analytics Rate Limiting : This policy is applied on Process Level API to secure API's against degradation of service that can happen in case load received is more than it can handle Custom circuit breaker : This is also quite useful feature on process level API's as it saves the API client the wasted time and effort of invoking a failing API. JSON threat protection : This policy is not required at Process API and rather implemented as Experience API's. This policy is used to safeguard application from malicious attacks by injecting malicious code in JSON object. As ideally Process API's are never called from external world , this policy is never used on Process API's Hence correct answer is
JSON threat protection MuleSoft Documentation Reference : https://docs.mulesoft.com/api-manager/2.x/policy-mule3-json-threat

NEW QUESTION 171
An organization is migrating all its Mule applications to Runtime Fabric (RTF). None of the Mule applications use Mule domain projects.
Currently, all the Mule applications have been manually deployed to a server group among several customer hosted Mule runtimes.
Port conflicts between these Mule application deployments are currently managed by the DevOps team who carefully manage Mule application properties files.
When the Mule applications are migrated from the current customer-hosted server group to Runtime Fabric (RTF), fo the Mule applications need to be rewritten and what DevOps port configuration responsibilities change or stay the same?

A. Yes, the Mule applications Must be rewrittenDevOps No Longer needs to manage port conflicts between the Mule applications
B. Yes, the Mule applications Must be rewritten DevOps Must Still Manage port conflicts.
C. NO, The Mule applications do NOT need to be rewrittenDevOps MUST STILL manage port conflicts
D. NO, the Mule applications do NO need to be rewrittenDevOps NO LONGER needs to manage port conflicts between the Mule applications.

Answer: C

Explanation:
* Anypoint Runtime Fabric is a container service that automates the deployment and orchestration of your Mule applications and gateways.
* Runtime Fabric runs on customer-managed infrastructure on AWS, Azure, virtual machines (VMs) or bare-metal servers.
* As none of the Mule applications use Mule domain projects. applications are not required to be rewritten. Also when applications are deployed on RTF, by default ingress is allowed only on 8081.
* Hence port conflicts are not required to be managed by DevOps team

NEW QUESTION 174
An integration Mute application is being designed to process orders by submitting them to a backend system for offline processing. Each order will be received by the Mute application through an HTTPS POST and must be acknowledged immediately. Once acknowledged, the order will be submitted to a backend system. Orders that cannot be successfully submitted due to rejections from the backend system will need to be processed manually (outside the backend system).
The Mule application will be deployed to a customer-hosted runtime and is able to use an existing ActiveMQ broker if needed.
The backend system has a track record of unreliability both due to minor network connectivity issues and longer outages.
What idiomatic (used for their intended purposes) combination of Mule application components and ActiveMQ queues are required to ensure automatic submission of orders to the backend system, while minimizing manual order processing?

A. An On Error scope Non-persistent VM ActiveMQ Dead Letter Queue for manual processing
B. An On Error scope MuleSoft Object Store ActiveMQ Dead Letter Queue for manual processing
C. Until Successful component MuleSoft Object Store ActiveMQ is NOT needed or used
D. Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing

Answer: D

Explanation:
Correct answer is using below set of activities Until Successful component ActiveMQ long retry Queue ActiveMQ Dead Letter Queue for manual processing We will see why this is correct answer but before that lets understand few of the concepts which we need to know. Until Successful Scope The Until Successful scope processes messages through its processors until the entire operation succeeds. Until Successful repeatedly retries to process a message that is attempting to complete an activity such as: - Dispatching to outbound endpoints, for example, when calling a remote web service that may have availability issues. - Executing a component method, for example, when executing on a Spring bean that may depend on unreliable resources. - A sub-flow execution, to keep re-executing several actions until they all succeed, - Any other message processor execution, to allow more complex scenarios. How this will help requirement : Using Until Successful Scope we can retry sending the order to backend systems in case of error to avoid manual processing later. Retry values can be configured in Until Successful Scope Apache ActiveMQ It is an open source message broker written in Java together with a full Java Message Service client ActiveMQ has the ability to deliver messages with delays thanks to its scheduler. This functionality is the base for the broker redelivery plug-in. The redelivery plug-in can intercept dead letter processing and reschedule the failing messages for redelivery. Rather than being delivered to a DLQ, a failing message is scheduled to go to the tail of the original queue and redelivered to a message consumer. How this will help requirement : If backend application is down for a longer duration where Until Successful Scope wont work, then we can make use of ActiveMQ long retry Queue. The redelivery plug-in can intercept dead letter processing and reschedule the failing messages for redelivery. Mule Reference:
https://docs.mulesoft.com/mule-runtime/4.3/migration-core-until-successful

NEW QUESTION 175
An organization is successfully using API led connectivity, however, as the application network grows, all the manually performed tasks to publish share and discover, register, apply policies to, and deploy an API are becoming repetitive pictures driving the organization to automate this process using efficient CI/'CD pipeline. Considering Anypoint platforms capabilities how should the organization approach automating is API lifecycle?

A. Use runtime manager rest apis for API management and mavenforAPI deployment
B. Use Maven with a custom configuration required for the API lifecycle
C. Use Anypoint CLI or Anypoint Platform REST apis with scripting language such as groovy
D. Use Exchange rest api's for API management and MavenforAPI deployment

Answer: D

NEW QUESTION 177
An external web UI application currently accepts occasional HTTP requests from client web browsers to change (insert, update, or delete) inventory pricing information in an inventory system's database. Each inventory pricing change must be transformed and then synchronized with multiple customer experience systems in near real-time (in under 10 seconds). New customer experience systems are expected to be added in the future.
The database is used heavily and limits the number of SELECT queries that can be made to the database to 10 requests per hour per user.
What is the most scalable, idiomatic (used for its intended purpose), decoupled. reusable, and maintainable integration mechanism available to synchronize each inventory pricing change with the various customer experience systems in near real-time?

A. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the watermark attribute set to an appropriate database columnIn the same now, use a Scatter-Gather to call each customer experience system's REST API with transformed inventory-pricing records
B. Add a trigger to the inventory-pricing database table so that for each change to the inventory pricing database, a stored procedure is called that makes a REST call to a Mule applicationWrite the Mule application to publish each Mule event as a message to an Anypoint MQ exchange Write other Mule applications to subscribe to the Anypoint MQ exchange, transform each receivedmessage, and then update the Mule application's corresponding customer experience system(s)
C. Replace the external web UI application with a Mule application to accept HTTP requests from client web browsersIn the same Mule application, use a Batch Job scope to test if the database request will succeed, aggregate pricing changes within a short time window, and then update both the inventory pricing database and each customer experience system using a Parallel For Each scope
D. Write a Mule application with a Database On Table Row event source configured for the inventory pricing database, with the ID attribute set to an appropriate database columnIn the same flow, use a Batch Job scope to publish transformed Inventory-pricing records to an Anypoint MQ queueWrite other Mule applications to subscribe to the Anypoint MQ queue, transform each received message, and then update the Mule application's corresponding customer experience system(s)

Answer: B

NEW QUESTION 179
An organization has just developed a Mule application that implements a REST API. The mule application will be deployed to a cluster of customer hosted Mule runtimes.

What additional infrastructure component must the customer provide in order to distribute inbound API requests across the Mule runtimes of the cluster?

A. A message broker
B. An HTTP Load Balancer
C. A database
D. An Object Store

Answer: B

Explanation:
Correct answer is An HTTP Load Balancer.
Key thing to note here is that we are deploying application to customer hosted Mule runtime. This means we will need load balancer to route the requests to different instances of the cluster.

NEW QUESTION 180
A marketing organization is designing a Mule application to process campaign data. The Mule application will periodically check for a file in a SFTP location and process the records in the file. The size of the file can vary from 10MB to 5GB. Due to the limited availabitlty of vCores, the Mule application is deployed to a single CloudHub worker configured with vCore size 0.2.
The application must transform and send different formats of this file to three different downstream SFTP locations.
What is the most idiomatic (used for its intended purpose) and performant way to configure the SFTP operations or event sources to process the large files to support these deployment requirements?

A. Use an in-memory repeatable stream
B. Use a file-stored non-repeatable stream
C. Use an in-memory non-repeatable stream
D. Use a file-stored repeatable stream

Answer: A

NEW QUESTION 183
An organization is building a test suite for their applications using m-unit. The integration architect has recommended using test recorder in studio to record the processing flows and then configure unit tests based on the capture events
What are the two considerations that must be kept in mind while using test recorder (Choose two answers)

A. Tests for flows cannot be created with Mule errors raised inside the flow or already existing in the incoming event
B. Recorder supports smoking a message before or inside a ForEach processor
C. The recorder support loops where the structure of the data been tested changes inside the iteration
D. A recorded flow execution ends successfully but the result does not reach its destination because the application is killed
E. Mocking values resulting from parallel processes are possible and will not affect the execution of theprocesses that follow in the test

Answer: AD

NEW QUESTION 188
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

Explanation:
* Semantic Versioning is a 3-component number in the format of X.Y.Z, where : X stands for a major version.
Y stands for a minor version:
Z stands for a patch.
So, SemVer is of the form Major.Minor.Patch Coming to our question , minor version of the API has been changed which is backward compatible. Hence there is no change required on API client end. If they want to make use of new featured that have been added as a part of minor version change they may need to change code at their end. Hence correct answer is The API client code ONLY needs to be changed if it needs to take advantage of new features.
Diagram Description automatically generated

NEW QUESTION 189
What aspects of a CI/CD pipeline for Mule applications can be automated using MuleSoft-provided Maven plugins?

A. Compile, package, unit test, validate unit test coverage, deploy
B. Compile, package, unit test, deploy, integration test (Incorrect)
C. Compile, package, unit test, deploy, create associated API instances in API Manager
D. Import from API designer, compile, package, unit test, deploy, publish to Anypoint Exchange

Answer: A

Explanation:
Correct answer is "Compile, package, unit test, validate unit test coverage, deploy"
Anypoint Platform supports continuous integration and continuous delivery using industry standard tools Mule Maven Plugin The Mule Maven plugin can automate building, packaging and deployment of Mule applications from source projects Using the Mule Maven plugin, you can automate your Mule application deployment to CloudHub, to Anypoint Runtime Fabric, or on-premises, using any of the following deployment strategies • CloudHub deployment • Runtime Fabric deployment • Runtime Manager REST API deployment • Runtime Manager agent deployment MUnit Maven Plugin The MUnit Maven plugin can automate test execution, and ties in with the Mule Maven plugin. It provides a full suite of integration and unit test capabilities, and is fully integrated with Maven and Surefire for integration with your continuous deployment environment. Since MUnit 2.x, the coverage report goal is integrated with the maven reporting section. Coverage Reports are generated during Maven's site lifecycle, during the coverage-report goal. One of the features of MUnit Coverage is to fail the build if a certain coverage level is not reached. MUnit is not used for integration testing Also publishing to Anypoint Exchange or to create associated API instances in API Manager is not a part of CICD pipeline which can ne achieved using mulesoft provided maven plugin
Explanation
Architecture mentioned in the question can be diagrammatically put as below. Persistent Object Store is the correct answer .
* Mule Object Stores: An object store is a facility for storing objects in or across Mule applications. Mule uses object stores to persist data for eventual retrieval. Mule provides two types of object stores:
1) In-memory store – stores objects in local Mule runtime memory. Objects are lost on shutdown of the Mule runtime. So we cant use in memory store in our scenario as we want to share watermark within all cloudhub workers
2) Persistent store – Mule persists data when an object store is explicitly configured to be persistent. Hence this watermark will be available even any of the worker goes down
Diagram Description automatically generated
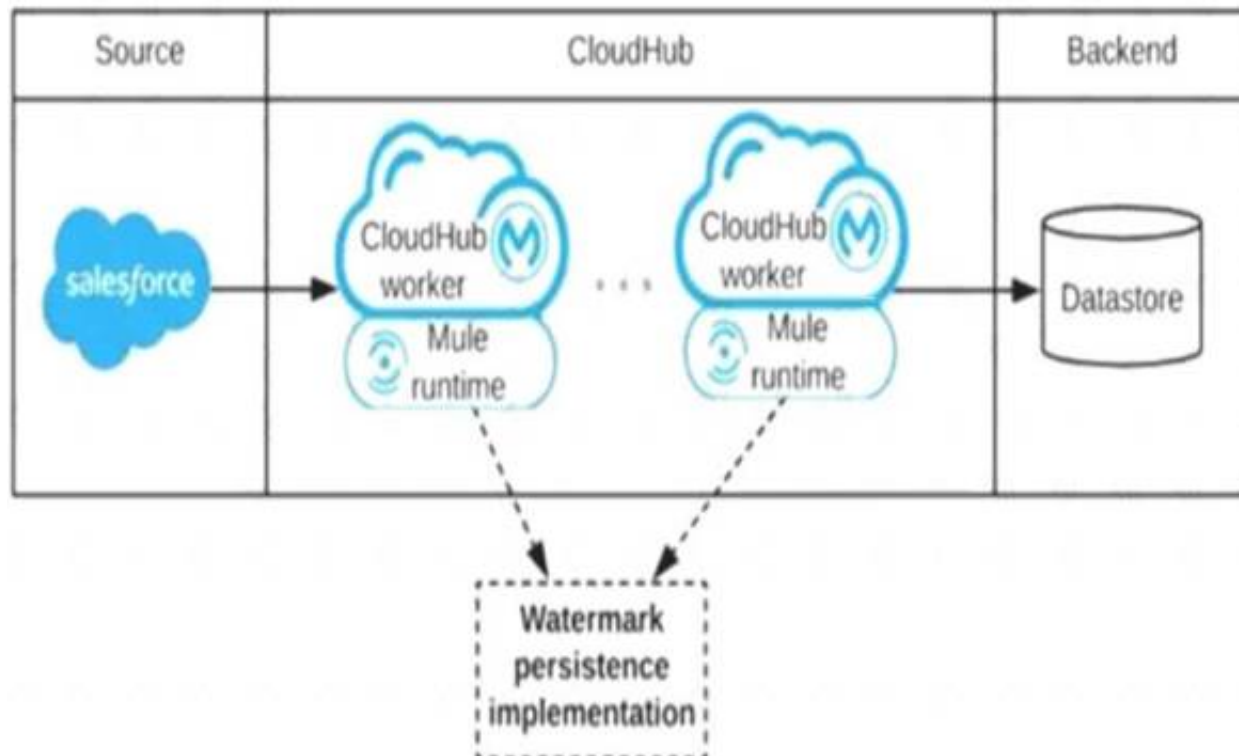
NEW QUESTION 191
A team would like to create a project skeleton that developers can use as a starting point when creating API Implementations with Anypoint Studio. This skeleton should help drive consistent use of best practices within the team.
What type of Anypoint Exchange artifact(s) should be added to Anypoint Exchange to publish the project skeleton?

A. A custom asset with the default API implementation
B. A RAML archetype and reusable trait definitions to be reused across API implementations
C. An example of an API implementation following best practices
D. a Mule application template with the key components and minimal integration logic

Answer: D

Explanation:
* Sharing Mule applications as templates is a great way to share your work with other people who are in your organization in Anypoint Platform. When they need to build a similar application they can create the mule application using the template project from Anypoint studio.
* Anypoint Templates are designed to make it easier and faster to go from a blank canvas to a production application. They're bit for bit Mule applications requiring only Anypoint Studio to build and design, and are deployable both on-premises and in the cloud.
* Anypoint Templates are based on five common data Integration patterns and can be customized and extended to fit your integration needs. So even if your use case involves different endpoints or connectors than those included in the template, they still offer a great starting point.
Some of the best practices while creating the template project: - Define the common error handler as part of template project, either using pom dependency or mule config file - Define common logger/audit framework as part of the template project - Define the env specific properties and secure properties file as per the requirement - Define global.xml for global configuration - Define the config file for connector configuration like Http,Salesforce,File,FTP etc - Create separate folders to create DWL,Properties,SSL certificates etc - Add the dependency and configure the pom.xml as per the business need - Configure the mule-artifact.json as per the business need

NEW QUESTION 195

Anypoint Exchange is required to maintain the source code of some of the assets committed to it, such as Connectors, Templates, and API specifications.
What is the best way to use an organization's source-code management (SCM) system in this context?

A. Organizations should continue to use an SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching, and merging
B. Organizations need to use Anypoint Exchange as the main SCM system to centralize versioning and avoid code duplication
C. Organizations can continue to use an SCM system of their choice for branching and merging, as long as they follow the branching and merging strategy enforced by Anypoint Exchange
D. Organizations need to point Anypoint Exchange to their SCM system so Anypoint Exchange can pull source code when requested by developers and provide it to Anypoint Studio

Answer: B

Explanation:
* Organization should continue to use SCM system of their choice, in addition to keeping source code for these asset types in Anypoint Exchange, thereby enabling parallel development, branching.
* Reason is that Anypoint exchange is not full fledged version repositories like GitHub.
* But at same time it is tightly coupled with Mule assets

NEW QUESTION 198
A travel company wants to publish a well-defined booking service API to be shared with its business partners. These business partners have agreed to ONLY consume SOAP services and they want to get the service contracts in an easily consumable way before they start any development. The travel company will publish the initial design documents to Anypoint Exchange, then share those documents with the business partners. When using an API-led approach, what is the first design document the travel company should deliver to its business partners?

A. Create a WSDL specification using any XML editor
B. Create a RAML API specification using any text editor
C. Create an OAS API specification in Design Center
D. Create a SOAP API specification in Design Center

Answer: A

Explanation:
SOAP API specifications are provided as WSDL. Design center doesn't provide the functionality to create WSDL file. Hence WSDL needs to be created using XML editor

NEW QUESTION 199
An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organizational has all of its Mule applications running on its own premises and uses an premises load balancer that exposes all APIs under the base URL https://api.acme.com
As part of the migration strategy, the organization plans to migrate all of its Mule applications and load balancer to cloudhub
What is the most straight-forward and cost effective approach to the Mule applications deployment and load balancing that preserves the public URLs?

A. Deploy the Mule applications to CloudhubUpdate the CNAME record for an api.acme.com in the organizations DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map URLs to their corresponding Mule applications
B. For each migrated Mule application, deploy an API proxy Mule application to Cloudhub with all applications under the control of a dedicated load balancer(CLB)Update the CNAME record for api.acme.com in the organization DNS server pointing to the A record of a cloudhub dedicated load balancer(DLB)Apply mapping rules in the DLB to map each API proxy application to its corresponding Mule applications
C. Deploy the Mule applications to CloudhubCreate CNAME record for api.acme.com in the Cloudhub Shared load balancer (SLB) pointing to the A record of the on-premise load balancerApply mapping rules in the SLB to map URLs to their corresponding Mule applications
D. Deploy the Mule applications to CloudhubUpdate the CNAME record for api.acme.com in the organization DNS server pointing to the A record of the cloudhub shared load balancer(SLB)Apply mapping rules in the SLB to map URLs to their corresponding Mule applications.

Answer: A

Explanation:
https://help.mulesoft.com/s/feed/0D52T000055pzgsSAA.

NEW QUESTION 201
A company wants its users to log in to Anypoint Platform using the company's own internal user credentials. To achieve this, the company needs to integrate an external identity provider (IdP) with the company's
Anypoint Platform master organization, but SAML 2.0 CANNOT be used. Besides SAML 2.0, what single-sign-on standard can the company use to integrate the IdP with their Anypoint Platform master organization?

A. SAML 1.0
B. OAuth 2.0
C. Basic Authentication
D. OpenID Connect

Answer: D

Explanation:
As the Anypoint Platform organization administrator, you can configure identity management in Anypoint Platform to set up users for single sign-on (SSO).
Configure identity management using one of the following single sign-on standards:
1) OpenID Connect: End user identity verification by an authorization server including SSO
2) SAML 2.0: Web-based authorization including cross-domain SSO

NEW QUESTION 203
An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger of flow for processing a high-volume set of records from the source system and synchronize with the SaaS system using the Batch job scope

After processing 1000 records in a periodic synchronization of 1 lakh records, the replicas in which batch job instance was started went down due to unexpected failure in the runtime fabric environment
What is the consequence of losing the replicas that run the Batch job instance?

A. The remaining 99000 records will be lost and left and processed
B. The second replicas will take over processing the remaining 99000 records
C. A new replacement replica will be available and will be process all 1,00,000 records from scratch leading to duplicate record processing
D. A new placement replica will be available and will take or processing the remaining 99,000 records

Answer: B

NEW QUESTION 208
A customer wants to use the mapped diagnostic context (MDC) and logging variables to enrich its logging and improve tracking by providing more context in the logs.
The customer also wants to improve the throughput and lower the latency of message processing.
As an Mulesoft integration architect can you advise, what should the customer implement to meet these requirements?

A. Use synchronous logging and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
B. Useasync logger at the level greater than INFO and use pattern layout with [%MDC] in the log4j2,xml configuration file and then configure the logging variables
C. Useasync logger at the level equal to DEBUG orTRACEand use patternlayout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables
D. Use synchronous logging at the INFO DEBUG or Trace level and use pattern layout with [%MDC] in the log4j2.xml configuration file and then configure the logging variables

Answer: B

NEW QUESTION 210
What is not true about Mule Domain Project?

A. This allows Mule applications to share resources
B. Expose multiple services within the Mule domain on the same port
C. Only available Anypoint Runtime Fabric
D. Send events (messages) to other Mule applications using VM queues

Answer: C

Explanation:
* Mule Domain Project is ONLY available for customer-hosted Mule runtimes, but not for Anypoint Runtime
Fabric
* Mule domain project is available for Hybrid and Private Cloud (PCE). Rest all provide application isolation and can't support domain project.
What is Mule Domain Project?
* A Mule Domain Project is implemented to configure the resources that are shared among different projects. These resources can be used by all the projects associated with this domain. Mule applications can be associated with only one domain, but a domain can be associated with multiple projects. Shared resources allow multiple development teams to work in parallel using the same set of reusable connectors. Defining these connectors as shared resources at the domain level allows the team to: - Expose multiple services within the domain through the same port. - Share the connection to persistent storage. - Share services between apps through a well-defined interface. - Ensure consistency between apps upon any changes because the configuration is only set in one place.
* Use domains Project to share the same host and port among multiple projects. You can declare the http connector within a domain project and associate the domain project with other projects. Doing this also allows to control thread settings, keystore configurations, time outs for all the requests made within multiple applications. You may think that one can also achieve this by duplicating the http connector configuration across all the applications. But, doing this may pose a nightmare if you have to make a change and redeploy all the applications.
* If you use connector configuration in the domain and let all the applications use the new domain instead of a default domain, you will maintain only one copy of the http connector configuration. Any changes will require only the domain to the redeployed instead of all the applications.
You can start using domains in only three steps:
1) Create a Mule Domain project
2) Create the global connector configurations which needs to be shared across the applications inside the Mule Domain project
3) Modify the value of domain in mule-deploy.properties file of the applications Graphical user interface Description automatically generated
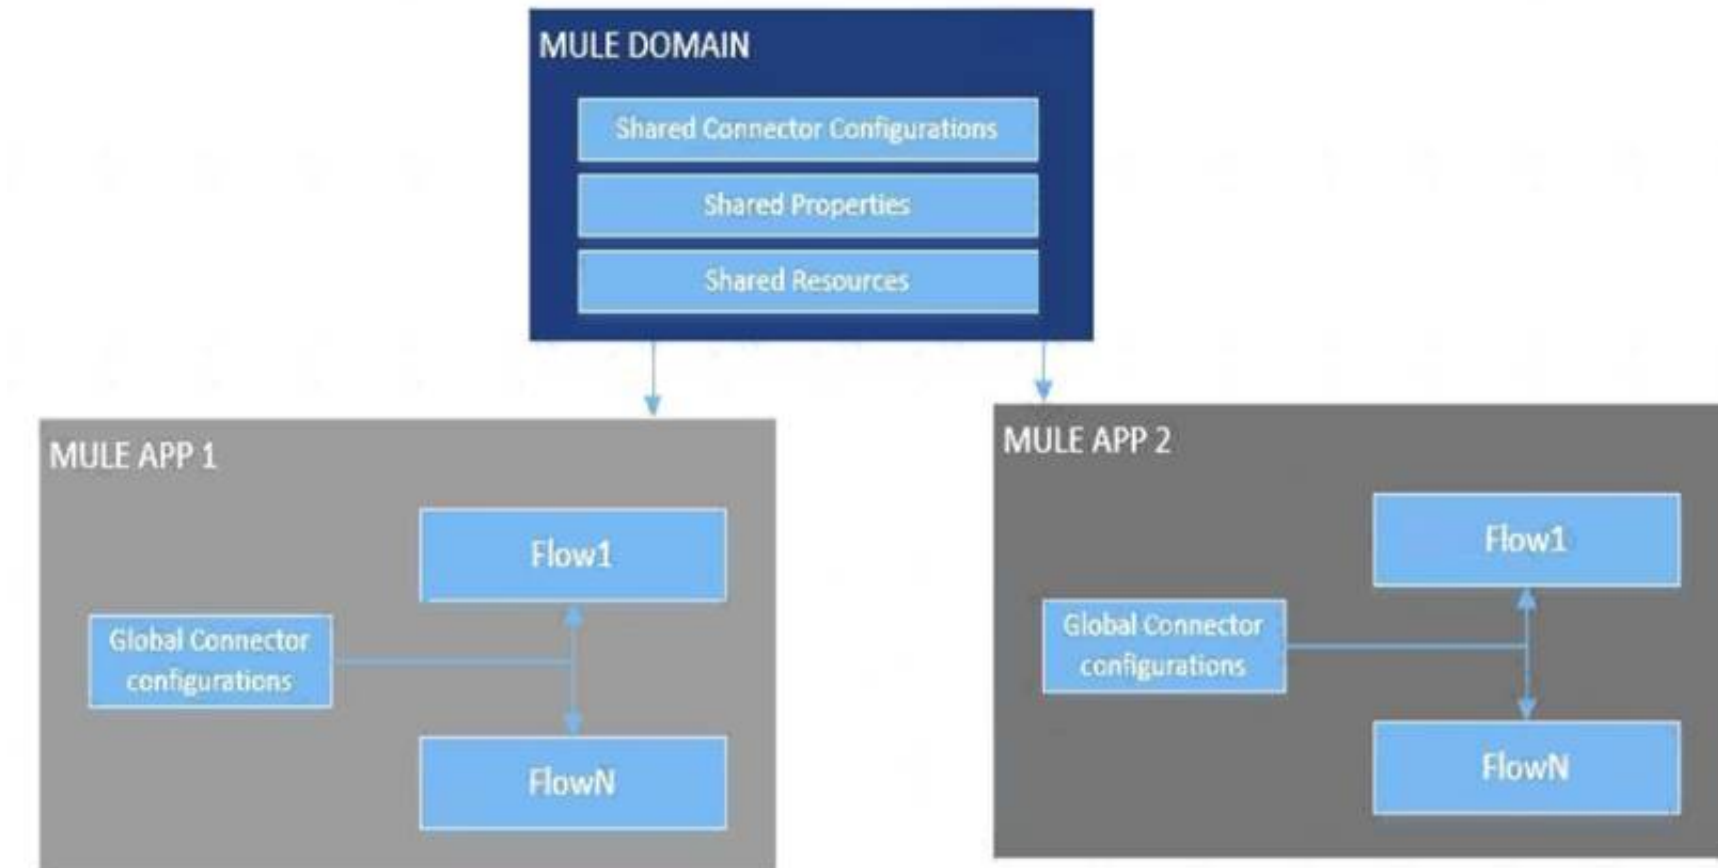
NEW QUESTION 215

An insurance provider is implementing Anypoint platform to manage its application infrastructure and is using the customer hosted runtime for its business due to certain financial requirements it must meet. It has built a number of synchronous API's and is currently hosting these on a mule runtime on one server

These applications make use of a number of components including heavy use of object stores and VM queues. Business has grown rapidly in the last year and the insurance provider is starting to receive reports of reliability issues from its applications.

The DevOps team indicates that the API's are currently handling too many requests and this is over loading the server. The team has also mentioned that there is a significant downtime when the server is down for maintenance.

As an integration architect, which option would you suggest to mitigate these issues?

A. Add a load balancer and add additional servers in a server group configuration
B. Add a load balancer and add additional servers in a cluster configuration
C. Increase physical specifications of server CPU memory and network
D. Change applications by use an event-driven model

Answer: B

NEW QUESTION 220

Refer to the exhibit.
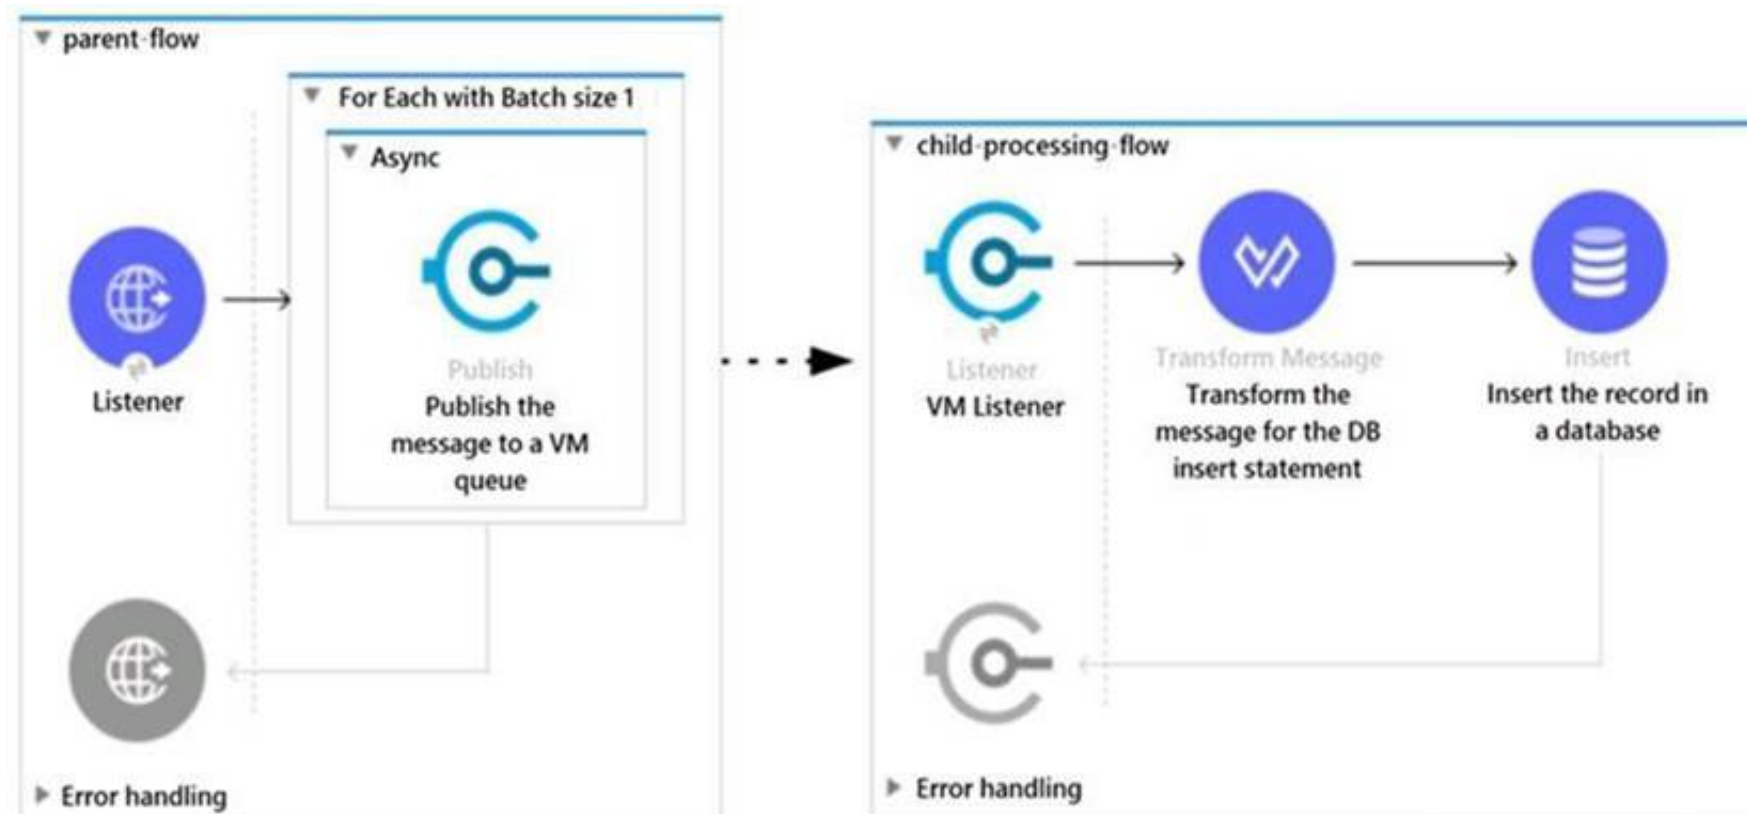
A Mule 4 application has a parent flow that breaks up a JSON array payload into 200 separate items, then sends each item one at a time inside an Async scope to a VM queue.

A second flow to process orders has a VM Listener on the same VM queue. The rest of this flow processes each received item by writing the item to a database. This Mule application is deployed to four CloudHub workers with persistent queues enabled.

What message processing guarantees are provided by the VM queue and the CloudHub workers, and how are VM messages routed among the CloudHub workers for each invocation of the parent flow under normal operating conditions where all the CloudHub workers remain online?

A. EACH item VM message is processed AT MOST ONCE by ONE CloudHub worker, with workers chosen in a deterministic round-robin fashion Each of the four CloudHub workers can be expected to process 1/4 of the Item VM messages (about 50 items)
B. EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker Each of the four CloudHub workers can be expected to process some item VM messages

C. ALL Item VM messages are processed AT LEAST ONCE by the SAME CloudHub worker where the parent flow was invokedThis one CloudHub worker processes ALL 200 item VM messages
D. ALL item VM messages are processed AT MOST ONCE by ONE ARBITRARY CloudHub worker This one CloudHub worker processes ALL 200 item VM messages

Answer: B

Explanation:
Correct answer is EACH item VM message is processed AT LEAST ONCE by ONE ARBITRARY CloudHub worker. Each of the four CloudHub workers can be expected to process some item VM messages In Cloudhub, each persistent VM queue is listened on by every CloudHub worker - But each message is read and processed at least once by only one CloudHub worker and the duplicate processing is possible - If the CloudHub worker fails , the message can be read by another worker to prevent loss of messages and this can lead to duplicate processing - By default , every CloudHub worker's VM Listener receives different messages from VM Queue Referenece: https://dzone.com/articles/deploying-mulesoft-application-on-1-worker-vs-mult

NEW QUESTION 224
A global, high-volume shopping Mule application is being built and will be deployed to CloudHub. To improve performance, the Mule application uses a Cache scope that maintains cache state in a CloudHub object store. Web clients will access the Mule application over HTTP from all around the world, with peak volume coinciding with business hours in the web client's geographic location. To achieve optimal performance, what Anypoint Platform region should be chosen for the CloudHub object store?

A. Choose the same region as to where the Mule application is deployed
B. Choose the US-West region, the only supported region for CloudHub object stores
C. Choose the geographically closest available region for each web client
D. Choose a region that is the traffic-weighted geographic center of all web clients

Answer: A

Explanation:
CloudHub object store should be in same region where the Mule application is deployed. This will give optimal performance.
Before learning about Cache scope and object store in Mule 4 we understand what is in general Caching is and other related things.
WHAT DOES "CACHING" MEAN?
Caching is the process of storing frequently used data in memory, file system or database which saves processing time and load if it would have to be accessed from original source location every time.
In computing, a cache is a high-speed data storage layer which stores a subset of data, so that future requests for that data are served up faster than is possible by accessing the data's primary storage location. Caching allows you to efficiently reuse previously retrieved or computed data.
How does Caching work?
The data in a cache is generally stored in fast access hardware such as RAM (Random-access memory) and may also be used in correlation with a software component. A cache's primary purpose is to increase data retrieval performance by reducing the need to access the underlying slower storage layer.
Caching in MULE 4
In Mule 4 caching can be achieved in mule using cache scope and/or object-store. Cache scope internally uses Object Store to store the data.
What is Object Store
Object Store lets applications store data and states across batch processes, Mule components, and applications, from within an application. If used on cloud hub, the object store is shared between applications deployed on Cluster.
Cache Scope is used in below-mentioned cases:
Need to store the whole response from the outbound processor
Data returned from the outbound processor does not change very frequently
As Cache scope internally handle the cache hit and cache miss scenarios it is more readable Object Store is used in below-mentioned cases:
Need to store custom/intermediary data To store watermarks
Sharing the data/stage across applications, schedulers, batch.
If CloudHub object store is in same region where the Mule application is deployed it will aid in fast access of data and give optimal performance.

NEW QUESTION 226
A Mule application uses APIkit for SOAP to implement a SOAP web service. The Mule application has been deployed to a CloudHub worker in a testing environment.
The integration testing team wants to use a SOAP client to perform Integration testing. To carry out the integration tests, the integration team must obtain the interface definition for the SOAP web service.
What is the most idiomatic (used for its intended purpose) way for the integration testing team to obtain the interface definition for the deployed SOAP web service in order to perform integration testing with the SOAP client?

A. Retrieve the OpenAPI Specification file(s) from API Manager
B. Retrieve the WSDL file(s) from the deployed Mule application
C. Retrieve the RAML file(s) from the deployed Mule application
D. Retrieve the XML file(s) from Runtime Manager

Answer: D

NEW QUESTION 229
An organization has strict unit test requirement that mandate every mule application must have an MUnit test suit with a test case defined for each flow and a minimum test coverage of 80%.
A developer is building Munit test suit for a newly developed mule application that sends API request to an
external rest API.
What is the effective approach for successfully executing the Munit tests of this new application while still achieving the required test coverage for the Munit tests?

A. Invoke the external endpoint of the rest API from the mule floors
B. Mark the rest API invocations in the Munits and then call the mocking service flow that simulates standard responses from the REST API
C. Mock the rest API invocation in the Munits and return a mock response for those invocations
D. Create a mocking service flow to simulate standard responses from the rest API and then configure the mule flows to call the marking service flow

Answer: C

NEW QUESTION 230
......

Relate Links

100% Pass Your MCIA-Level-1 Exam with Exambible Prep Materials

https://www.exambible.com/MCIA-Level-1-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/