

Salesforce

Exam Questions MuleSoft-Integration-Architect-I

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam



NEW QUESTION 1

According to MuleSoft, which system integration term describes the method, format, and protocol used for communication between two system?

- A. Component
- B. interaction
- C. Message
- D. Interface

Answer: D

NEW QUESTION 2

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

- A. Redirect client requests through an HTTP 303 temporary redirect status code to the fallback API whenever the Order API is unavailable
- B. Set an option in the HTTP Requester component that invokes the order API to instead invoke a fallback API whenever an HTTP 4XX or 5XX response status code is received from Order API
- C. Create a separate entry for the order API in API manager and then invoke this API as a fallback API if the primary Order API is unavailable
- D. Search Anypoint Exchange for a suitable existing fallback API and then implement invocations to their fallback API in addition to the Order API

Answer: A

NEW QUESTION 3

What is an advantage that Anypoint Platform offers by providing universal API management and Integration-Platform-as-a-Service (iPaaS) capabilities in a unified platform?

- A. Ability to use a single iPaaS to manage and integrate all API gateways
- B. Ability to use a single connector to manage and integrate all APIs
- C. Ability to use a single control plane for both full-lifecycle API management and integration
- D. Ability to use a single iPaaS to manage all API developer portals

Answer: C

NEW QUESTION 4

Insurance organization is planning to deploy Mule application in MuleSoft Hosted runtime plane. As a part of requirement , application should be scalable . highly available. It also has regulatory requirement which demands logs to be retained for at least 2 years. As an Integration Architect what step you will recommend in order to achieve this?

- A. It is not possible to store logs for 2 years in CloudHub deploymen
- B. External log management system is required.
- C. When deploying an application to CloudHub , logs retention period should be selected as 2 years
- D. When deploying an application to CloudHub, worker size should be sufficient to store 2 years data
- E. Logging strategy should be configured accordingly in log4j file deployed with the application.

Answer: A

NEW QUESTION 5

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

NEW QUESTION 6

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

NEW QUESTION 7

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

NEW QUESTION 8

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

NEW QUESTION 9

A high-volume eCommerce retailer receives thousands of orders per hour and requires notification of its order management, warehouse, and billing system for subsequent processing within 15 minutes of order submission through its website.

Which integration technology, when used for its typical and intended purpose, meets the retailer??s requirements for this use case?

- A. Managed File Transfer (MFT)
- B. Publish/Subscriber Messaging Bus (Pub/Sub)
- C. Enterprise Data Warehouse (EDW)
- D. Extract Transform Load (ETL)

Answer: B

NEW QUESTION 10

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

NEW QUESTION 10

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: A

NEW QUESTION 11

A gaming company has implemented an API as a Mule application and deployed the API implementation to a CloudHub 2.0 private space. The API implementation must connect to a mainframe application running in the customer??s on-premises corporate data center and also to a Kafka cluster running in an Amazon AWS VPC.


What is the most efficient way to enable the API to securely connect from its private space to the mainframe application and Kafka cluster?

- A. In Runtime Manager, set up VPC peering between the CloudHub 2.0 private network and the on-premises data center.In the AWS account, set up VPC peering between the AWS VPC and the CloudHub 2.0 private network.




B) In the AWS account, attach the CloudHub 2.0 private space to an AWS transit gateway that routes from the CloudHub 2.0 private space to the on-premises data centre

B. InRuntime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the AWS VPC.




C) In Runtime Manager, configure an Anypoint VPN to route from the CloudHub 2.0 private space to the on-premises data centre

C. In the AWS account, attach the private..



D) In the AWS account, attach the private space directly to the AWS VPC, In the AWS account, use an AWS transit gateway to route from the AWS VPC to the on-premises data center.



D. Option A

E. Option B

F. Option C

G. Option D

Answer: B

NEW QUESTION 12

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

NEW QUESTION 13

A Mule application currently writes to two separate SQL Server database instances across the internet using a single XA transaction. It is proposed to split this one transaction into two separate non-XA transactions with no other changes to the Mule application.

What non-functional requirement can be expected to be negatively affected when implementing this change?

- A. Throughput
- B. Consistency
- C. Response time
- D. Availability

Answer: B

NEW QUESTION 17


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: CE

NEW QUESTION 18

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 characteristic
- B. Then configure the API implementation to use the mocking service and conduct the performance tests
- C. Use MUnit to simulate standard responses from the backend system then conduct performance tests to identify other bottlenecks in the system
- D. Include logic within the API implementation that bypasses invocations of the backend system in a performance test situation
- E. Instead invoking local stubs that replicate typical backend system responses then conduct performance tests using this API Implementation
- F. 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

NEW QUESTION 19

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 flows
- 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 23

Refer to the exhibit.



A business process involves the receipt of a file from an external vendor over SFTP. The file needs to be parsed and its content processed, validated, and ultimately persisted to a database. The delivery mechanism is expected to change in the future as more vendors send similar files using other mechanisms such as file transfer or HTTP POST.

What is the most effective way to design for these requirements in order to minimize the impact of future change?

- A. Use a MuleSoft Scatter-Gather and a MuleSoft Batch Job to handle the different files coming from different sources
- B. Create a Process API to receive the file and process it using a MuleSoft Batch Job while delegating the data save process to a System API
- C. Create an API that receives the file and invokes a Process API with the data contained In the file, then have the Process API process the data using a MuleSoft Batch Job and other System APIs as needed
- D. Use a composite data source so files can be retrieved from various sources and delivered to a MuleSoft Batch Job for processing

Answer: C

NEW QUESTION 24

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

NEW QUESTION 26

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(<https://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 the mule 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: C

NEW QUESTION 28

Mule application muleA deployed in cloudhub uses Object Store v2 to share data across instances. As a part of new requirement , application muleB which is deployed in same region wants to access this Object Store.

Which of the following option you would suggest which will have minimum latency in this scenario?

- A. Object Store REST API
- B. Object Store connector
- C. Both of the above option will have same latency
- D. Object Store of one mule application cannot be accessed by other mule application.

Answer: A

NEW QUESTION 32

An external API frequently invokes an Employees System API to fetch employee data from a MySQL database. The architect must design a caching strategy to query the database only when there is an update to the Employees table or else return a cached response in order to minimize the number of redundant transactions being handled by the database.

- A. Use an On Table Row operation configured with the Employees table, call invalidate cache, and hardcode the new Employees data to cache
- B. Use an object-store-caching- strategy and set the expiration interval to 1 hour.
- C. Use an On Table Row operation configured with the Employees table and call invalidate cache
- D. Use an object-store-caching-strategy and the default expiration interval.
- E. Use a Scheduler with a fixed frequency set to every hour to trigger an invalidate cache flow
- F. Use an object-store-caching-strategy and the default expiration interval.
- G. Use a Scheduler with a fixed frequency set to every hour, triggering an invalidate cache flow
- H. Use an object-store-caching-strategy and set the expiration interval to 1 hour.

Answer: B

NEW QUESTION 34

An organization is designing multiple new applications to run on CloudHub in a single Anypoint VPC and that must share data using a common persistent Anypoint object store V2 (OSv2).

Which design gives these mule applications access to the same object store instance?

- A. AVM connector configured to directly access the persistence queue of the persistent object store
- B. An Anypoint MQ connector configured to directly access the persistent object store
- C. Object store V2 can be shared across cloudhub applications with the configured osv2 connector
- D. The object store V2 rest API configured to access the persistent object store

Answer: C

NEW QUESTION 36

In Anypoint Platform, a company wants to configure multiple identity providers (IdPs) for multiple lines of business (LOBs). Multiple business groups, teams, and environments have been defined for these LOBs.

What Anypoint Platform feature can use multiple IdPs across the company's business groups, teams, and environments?

- A. MuleSoft-hosted (CloudHub) dedicated load balancers
- B. Client (application) management
- C. Virtual private clouds
- D. Permissions

Answer: A

NEW QUESTION 39

An organization is designing a mule application to support an all or nothing transaction between several database operations and some other connectors so that they all roll back if there is a problem with any of the connectors

Besides the database connector , what other connector can be used in the transaction.

- A. VM
- B. Anypoint MQ
- C. SFTP
- D. ObjectStore

Answer: A

NEW QUESTION 43

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: A

NEW QUESTION 44

Following MuleSoft best practices, what MuleSoft runtime deployment option best meets the company's goals to begin its digital transformation journey?

- A. Runtime Fabric on VMs/bare metal
- B. CloudHub runtimes
- C. Customer-hosted runtimes provisioned by a MuleSoft services partner
- D. Customer-hosted self-provisioned runtimes

Answer: B

NEW QUESTION 47

When the mule application using VM is deployed to a customer-hosted cluster or multiple cloudhub workers, how are messages consumed by the Mule engine?

- A. in non-deterministic way
- B. by starting an XA transaction for each new message
- C. in a deterministic way
- D. the primary only in order to avoid duplicate processing

Answer: A

NEW QUESTION 51

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 exchangeWrite other Mule applications to subscribe to the Anypoint MQ exchange, transform each received message, 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 53

What is the MuleSoft-recommended best practice to share the connector and configuration information among the APIs?

- A. Build a Mule domain project, add the Database connector and configuration to it, and reference this one domain project from each System API
- B. Build a separate Mule domain project for each API, and configure each of them to use a file on a shared file store to load the configuration information dynamically
- C. Build another System API that connects to the database, and refactor all the other APIs to make requests through the new System API to access the database
- D. Create an API proxy for each System API and share the Database connector configuration with all the API proxies via an automated policy

Answer: A

NEW QUESTION 57

Which type of communication is managed by a service mesh in a microservices architecture?

- A. Communication between microservices runtime administrators
- B. Communication between microservices developers
- C. Communication between microservices
- D. Communication between trading partner services

Answer: C

NEW QUESTION 62

An organization plans to migrate all its Mule applications to Runtime Fabric (RTF). Currently, all Mule applications have been deployed to CloudHub using automated CI/CD scripts.

What steps should be taken to properly migrate the applications from CloudHub to RTF, while keeping the same automated CI/CD deployment strategy?

- A. A runtimefabric dependency should be added as a mule-plugin to the pom.xml file in all the Mule applications.
- B. runtimeFabric command-line parameter should be added to the CI/CD deployment scripts.
- C. A runtimeFabricDeployment profile should be added to Mule configuration properties YAML files in all the Mule applications.CI/CD scripts must be modified to use the new configuration properties.
- D. runtimefabricDeployment profile should be added to the pom.xml file in all the Mule application
- E. CI/CD scripts must be modified to use the new RTF profile.
- F. - The pom.xml and Mule configuration YAML files can remain unchanged in each Mule application.A --runtimeFabric command-line parameter should be added to the CI/CD deployment scripts

Answer: D

NEW QUESTION 67

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 72

An organization has implemented the cluster with two customer hosted Mule runtimes is hosting an application.

This application has a flow with a JMS listener configured to consume messages from a queue destination. As an integration architect can you advise which JMS listener configuration must be used to receive messages in all the nodes of the cluster?

- A. Use the parameter primaryNodeOnly= "false" on the JMS listener
- B. Use the parameter primaryNodeOnly= "false" on the JMS listener with a shared subscription
- C. Use the parameter primaryNodeOnly= "true" on the JMS listener with a non•shared subscription
- D. Use the parameter primaryNodeOnly= "true" on the JMS listener

Answer: B

NEW QUESTION 74

What is an example of data confidentiality?

- A. Signing a file digitally and sending it using a file transfer mechanism
- B. Encrypting a file containing personally identifiable information (PI)
- C. Providing a server's private key to a client for secure decryption of data during a two- way SSL handshake
- D. De-masking a person's Social Security number while inserting it into a database

Answer: B

NEW QUESTION 77

According to MuleSoft, which deployment characteristic applies to a microservices application architecture?

- A. Services exist as independent deployment artifacts and can be scaled -independently of other services
- B. All services of an application can be deployed together as single Java WAR file
- C. A deployment to enhance one capability requires a redeployment of all capabilities
- D. Core business capabilities are encapsulated in a single, deployable application

Answer: A

NEW QUESTION 78

An IT integration delivery team begins a project by gathering all of the requirements, and proceeds to execute the remaining project activities as sequential, non-repeating phases.

Which IT project delivery methodology is this team following?

- A. Kanban
- B. Scrum
- C. Waterfall
- D. Agile

Answer: C

NEW QUESTION 83

A new Mule application has been deployed through Runtime Manager to CloudHub 1.0 using a CI/CD pipeline with sensitive properties set as cleartext. The Runtime Manager Administrator opened a high priority incident ticket about this violation of their security requirements indicating these sensitive properties values must not be stored or visible in Runtime Manager but should be changeable in Runtime Manager by Administrators with proper permissions.


How can the Mule application be deployed while safely hiding the sensitive properties?

- A. Add an ArrayList of all the sensitive properties?? names in the mule-artifact.json file of the application
- B. Add encrypted versions of the sensitive properties as global configuration properties in the Mule application
- C. Add a new wrapper.java.additional.xx parameter for each sensitive property in the wrapper.conf file used by the CI/CD pipeline scripts
- D. Create a variable for each sensitive property and declare them as hidden in the CI/CD pipeline scripts

Answer: B

NEW QUESTION 84

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

NEW QUESTION 88

A system administrator needs to determine when permissions were last changed for an Anypoint Platform user. Which Anypoint Platform component should the administrator use to obtain this information?

- A. Audit Logging
- B. Anypoint Monitoring
- C. Anypoint Studio
- D. Mule Stack Traces

Answer: A

NEW QUESTION 92

In which order are the API Client, API Implementation, and API interface components called in a typical REST request?

- A. API Client > API implementation > API Interface
- B. API interface > API Client > API Implementation
- C. API Client > API Interface > API implementation
- D. API Implementation > API Interface > API Client

Answer: C

NEW QUESTION 94

An organization is creating a Mule application that will be deployed to CloudHub. The Mule application has a property named dbPassword that stores a database user's password.

The organization's security standards indicate that the dbPassword property must be hidden from every Anypoint Platform user after the value is set in the Runtime

Manager Properties tab.

What configuration in the Mule application helps hide the dbPassword property value in Runtime Manager?

- A. Use secure::dbPassword as the property placeholder name and store the cleartext (unencrypted) value in a secure properties placeholder file
- B. Use secure::dbPassword as the property placeholder name and store the property encrypted value in a secure properties placeholder file
- C. Add the dbPassword property to the secureProperties section of the pom.xml file
- D. Add the dbPassword property to the secureProperties section of the mule-artifact.json file

Answer: B

NEW QUESTION 98

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), for 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 rewritten DevOps 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 rewritten DevOps MUST STILL manage port conflicts
- D. NO, the Mule applications do NO need to be rewritten DevOps NO LONGER needs to manage port conflicts between the Mule applications.

Answer: C

NEW QUESTION 99

What requirement prevents using Anypoint MQ as the messaging broker for a Mule application?

- A. When the payload sent through the message broker must use XML format
- B. When the payload sent through the message broker must be encrypted
- C. When the messaging broker must support point-to-point messaging
- D. When the messaging broker must be deployed on-premises

Answer: D

NEW QUESTION 100

Which role is primarily responsible for building API implementation as part of a typical MuleSoft integration project?

- A. API Developer
- B. API Designer
- C. Integration Architect
- D. Operations

Answer: A

NEW QUESTION 104

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

NEW QUESTION 106


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

NEW QUESTION 107

Refer to the exhibit.



A Mule application is deployed to a cluster of two customer-hosted Mule runtimes. The Mule application has a flow that polls a database and another flow with an HTTP Listener.

HTTP clients send HTTP requests directly to individual cluster nodes.

What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has failed, but before that node is restarted?

- A. Database polling continues Only HTTP requests sent to the remaining node continue to be accepted
- B. Database polling stops All HTTP requests continue to be accepted
- C. Database polling continues All HTTP requests continue to be accepted, but requests to the failed node Incur increased latency
- D. Database polling stops All HTTP requests are rejected

Answer: A

NEW QUESTION 112

A retailer is designing a data exchange interface to be used by its suppliers. The interface must support secure communication over the public internet. The interface must also work with a wide variety of programming languages and IT systems used by suppliers.

What are suitable interface technologies for this data exchange that are secure, cross- platform, and internet friendly, assuming that Anypoint Connectors exist for these interface technologies?

- A. EDJFACT XML over SFTP JSON/REST over HTTPS
- B. SOAP over HTTPS HOP over TLS gRPC over HTTPS
- C. XML over ActiveMQ XML over SFTP XML/REST over HTTPS
- D. CSV over FTP YAML over TLS JSON over HTTPS

Answer: C

NEW QUESTION 116

A set of integration Mule applications, some of which expose APIs, are being created to enable a new business process. Various stakeholders may be impacted by this. These stakeholders are a combination of semi-technical users (who understand basic integration terminology and concepts such as JSON and XML) and technically skilled potential consumers of the Mule applications and APIs.

What Is an effective way for the project team responsible for the Mule applications and APIs being built to communicate with these stakeholders using Anypoint Platform and its supplied toolset?

- A. Use Anypoint Design Center to implement the Mule applications and APIs and give the various stakeholders access to these Design Center projects, so they can collaborate and provide feedback
- B. Create Anypoint Exchange entries with pages elaborating the integration design, including API notebooks (where applicable) to help the stakeholders understand and interact with the Mule applications and APIs at various levels of technical depth
- C. Use Anypoint Exchange to register the various Mule applications and APIs and share the RAML definitions with the stakeholders, so they can be discovered
- D. Capture documentation about the Mule applications and APIs inline within the Mule integration flows and use Anypoint Studio's Export Documentation feature to provide an HTML version of this documentation to the stakeholders

Answer: B

NEW QUESTION 121

An integration team follows MuleSoft's recommended approach to full lifecycle API development.

Which activity should this team perform during the API implementation phase?

- A. Validate the API specification
- B. Use the API specification to build the MuleSoft application
- C. Design the API specification
- D. Use the API specification to monitor the MuleSoft application

Answer: B

NEW QUESTION 125

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

NEW QUESTION 130

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 queues One or more ActiveMQ dead-letter queues for manual processing
- B. One or more On Error scopes to assist calling the back-end system An 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 system An Until Successful scope containing Object Store components for long retries A dead-letter object store configured in the Mule application

Answer: A

NEW QUESTION 132

An API client makes an HTTP request to an API gateway with an Accept header containing the value application/json.

What is a valid HTTP response payload for this request in the client requested data format?

- A. <status>healthy</status>
- B. {"status": "healthy"}
- C. status(??healthy")
- D. status: healthy

Answer: B

NEW QUESTION 137

What requires configuration of both a key store and a trust store for an HTTP Listener?

- A. Support for TLS mutual (two-way) authentication with HTTP clients
- B. Encryption of requests to both subdomains and API resource endpoints fhttpS://aDi.customer.com/ and https://customer.com/api)
- C. Encryption of both HTTP request and HTTP response bodies for all HTTP clients
- D. Encryption of both HTTP request header and HTTP request body for all HTTP clients

Answer: A

NEW QUESTION 141

An organization's IT team must secure all of the internal APIs within an integration solution by using an API proxy to apply required authentication and authorization policies.

Which integration technology, when used for its intended purpose, should the team choose to meet these requirements if all other relevant factors are equal?

- A. API Management (APIM)
- B. Robotic Process Automation (RPA)
- C. Electronic Data Interchange (EDI)
- D. Integration Platform-as-a-service (PaaS)

Answer: A

NEW QUESTION 142

A Mule application is being designed to do the following:

Step 1: Read a SalesOrder message from a JMS queue, where each SalesOrder consists of a header and a list of SalesOrderLineItems.

Step 2: Insert the SalesOrder header and each SalesOrderLineItem into different tables in an RDBMS.

Step 3: Insert the SalesOrder header and the sum of the prices of all its SalesOrderLineItems into a table in a different RDBMS.

No SalesOrder message can be lost and the consistency of all SalesOrder-related information in both RDBMSs must be ensured at all times.

What design choice (including choice of transactions) and order of steps addresses these requirements?

- A. 1) Read the JMS message (NOT in an XA transaction) 2) Perform BOTH DB inserts in ONE DB transaction 3) Acknowledge the JMS message
- B. 1) Read the JMS message (NOT in an XA transaction) 2) Perform EACH DB insert in a SEPARATE DB transaction 3) Acknowledge the JMS message

- C. 1) Read the JMS message in an XA transaction2) In the SAME XA transaction, perform BOTH DB inserts but do NOT acknowledge the JMS message
D. 1) Read and acknowledge the JMS message (NOT in an XA transaction)2) In a NEW XA transaction, perform BOTH DB inserts

Answer: A

NEW QUESTION 143

An organization has chosen Mulesoft for their integration and API platform.

According to the Mulesoft catalyst framework, what would an integration architect do to create achievement goals as part of their business outcomes?

- A. Measure the impact of the centre for enablement
B. build and publish foundational assets
C. agree upon KPI's and help develop and overall success plan
D. evangelize API's

Answer: C

NEW QUESTION 145

According to MuleSoft, a synchronous invocation of a RESTful API using HTTP to get an individual customer record from a single system is an example of which system integration interaction pattern?

- A. Request-Reply
B. Multicast
C. Batch
D. One-way

Answer: A

NEW QUESTION 147

An integration team uses Anypoint Platform and follows MuleSoft's recommended approach to full lifecycle API development.

Which step should the team's API designer take before the API developers implement the API Specification?

- A. Generate test cases using MUnit so the API developers can observe the results of running the API
B. Use the scaffolding capability of Anypoint Studio to create an API portal based on the API specification
C. Publish the API specification to Exchange and solicit feedback from the API's consumers
D. Use API Manager to version the API specification

Answer: C

NEW QUESTION 148

A Mule application contains a Batch Job with two Batch Steps (Batch_Step_1 and Batch_Step_2). A payload with 1000 records is received by the Batch Job. How many threads are used by the Batch Job to process records, and how does each Batch Step process records within the Batch Job?

- A. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and RECORDS are processed IN PARALLEL within and between the two Batch Steps
B. Each Batch Job uses a SINGLE THREAD for all Batch steps Each Batch step instance receives ONE record at a time as the payload, and RECORDS are processed IN ORDER, first through Batch_Step_1 and then through Batch_Step_2
C. Each Batch Job uses a SINGLE THREAD to process a configured block size of record Each Batch Step instance receives A BLOCK OF records as the payload, and BLOCKS of records are processed IN ORDER
D. Each Batch Job uses SEVERAL THREADS for the Batch Steps Each Batch Step instance receives ONE record at a time as the payload, and BATCH STEP INSTANCES execute IN PARALLEL to process records and Batch Steps in ANY order as fast as possible

Answer: A

NEW QUESTION 149

An organization is in the process of building automated deployments using a CI/CD process. As a part of automated deployments, it wants to apply policies to API Instances.


What tool can the organization use to promote and deploy API Manager policies?

- A. Anypoint CLI
B. MUnit Maven plugin
C. Mule Maven plugin
D. Runtime Manager agent

Answer: A

NEW QUESTION 150

Refer to the exhibit.




An organization is designing a Mule application to receive data from one external business partner. The two companies currently have no shared IT infrastructure and do not want to establish one. Instead, all communication should be over the public internet (with no VPN). What Anypoint Connector can be used in the organization's Mule application to securely receive data from this external business partner?

- A. File connector
- B. VM connector
- C. SFTP connector
- D. Object Store connector

Answer: C

NEW QUESTION 154

Refer to the exhibit.




A shopping cart checkout process consists of a web store backend sending a sequence of API invocations to an Experience API, which in turn invokes a Process API. All API invocations are over HTTPS POST. The Java web store backend executes in a Java EE application server, while all API implementations are Mule applications executing in a customer -hosted Mule runtime.

End-to-end correlation of all HTTP requests and responses belonging to each individual checkout Instance is required. This is to be done through a common correlation ID, so that all log entries written by the web store backend, Experience API implementation, and Process API implementation include the same correlation ID for all requests and responses belonging to the same checkout instance.

What is the most efficient way (using the least amount of custom coding or configuration) for the web store backend and the implementations of the Experience API and Process API to participate in end-to-end correlation of the API invocations for each checkout instance?


A) The web store backend, being a Java EE application, automatically makes use of the thread-local correlation ID generated by the Java EE application server and automatically transmits that to the Experience API using HTTP-standard headers

No special code or configuration is included in the web store backend, Experience API, and Process API implementations to generate and manage the correlation ID




B) The web store backend generates a new correlation ID value at the start of checkout and sets it on the X-CORRELATION-Id HTTP request header In each API invocation belonging to that checkout


No special code or configuration is included in the Experience API and Process API implementations to generate and manage the correlation ID



C) The Experience API implementation generates a correlation ID for each incoming HTTP request and passes it to the web store backend in the HTTP response, which includes it in all subsequent API invocations to the Experience API. The Experience API implementation must be coded to also propagate the correlation ID to the Process API in a suitable HTTP request header




D) The web store backend sends a correlation ID value in the HTTP request body In the way required by the Experience API The Experience API and Process API implementations must be coded to receive the custom correlation ID In the HTTP requests and propagate It in suitable HTTP request headers



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

Answer: B

NEW QUESTION 159
Refer to the exhibit.



Anypoint Platform supports role-based access control (RBAC) to features of the platform. An organization has configured an external Identity Provider for identity management with Anypoint Platform.

What aspects of RBAC must ALWAYS be controlled from the Anypoint Platform control plane and CANNOT be controlled via the external Identity Provider?

- A. Controlling the business group within Anypoint Platform to which the user belongs
- B. Assigning Anypoint Platform permissions to a role
- C. Assigning Anypoint Platform role(s) to a user
- D. Removing a user's access to Anypoint Platform when they no longer work for the organization

Answer: B

NEW QUESTION 160

A corporation has deployed multiple mule applications implementing various public and private API's to different cloudhub workers. These API's are critical applications that must be highly available and in line with the reliability SLA as defined by stakeholders.

How can API availability (liveliness or readiness) be monitored so that Ops team receives outage notifications?

- A. Enable monitoring of individual applications from Anypoint monitoring
- B. Configure alerts with failure conditions in runtime manager
- C. Configure alerts failure conditions in API manager
- D. Use any point functional monitoring test API's functional behavior

Answer: A

NEW QUESTION 163

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

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

Answer: C

NEW QUESTION 165

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 Mule 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

NEW QUESTION 167

A company is implementing a new Mule application that supports a set of critical functions driven by a REST API enabled, claims payment rules engine hosted on Oracle ERP. As designed the mule application requires many data transformation operations as it performs its batch processing logic.

The company wants to leverage and reuse as many of its existing Java-based capabilities (classes, objects, data model etc.) as possible

What approach should be considered when implementing required data mappings and transformations between Mule application and Oracle ERP in the new Mule application?

- A. Create a new metadata RAML classes in Mule from the appropriate Java objects and then perform transformations via Dataweave
- B. From the mule application, transform via the XSLT model
- C. Transform by calling any suitable Java class from Dataweave
- D. Invoke any of the appropriate Java methods directly, create metadata RAML classes and then perform required transformations via Dataweave

Answer: C

NEW QUESTION 172

A company is building an application network and has deployed four Mule APIs: one experience API, one process API, and two system APIs. The logs from all the APIs are aggregated in an external log aggregation tool. The company wants to trace messages that are exchanged between multiple API implementations. What is the most idiomatic (based on its intended use) identifier that should be used to implement Mule event tracing across the multiple API implementations?

- A. Mule event ID
- B. Mule correlation ID
- C. Client's IP address
- D. DataWeave UUID

Answer: B

NEW QUESTION 173

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: CD

NEW QUESTION 178

Which Exchange asset type represents configuration modules that extend the functionality of an API and enforce capabilities such as security?

- A. Rulesets
- B. Policies
- C. REST APIs
- D. Connectors

Answer: B

NEW QUESTION 183

As a part of design, Mule application is required call the Google Maps API to perform a distance computation. The application is deployed to cloudhub. At the minimum what should be configured in the TLS context of the HTTP request configuration to meet these requirements?

- A. The configuration is built-in and nothing extra is required for the TLS context
- B. Request a private key from Google and create a PKCS12 file with it and add it in keyStore as a part of TLS context
- C. Download the Google public certificate from a browser, generate JKS file from it and add it in key store as a part of TLS context
- D. Download the Google public certificate from a browser, generate a JKS file from it and add it in Truststore as part of the TLS context

Answer: D

NEW QUESTION 187

A new upstream API is being designed to offer an SLA of 500 ms median and 800 ms maximum (99th percentile) response time. The corresponding API implementation needs to sequentially invoke 3 downstream APIs of very similar complexity. The first of these downstream APIs offers the following SLA for its response time: median: 100 ms, 80th percentile: 500 ms, 95th percentile: 1000 ms. If possible, how can a timeout be set in the upstream API for the invocation of the first downstream API to meet the new upstream API's desired SLA?

- A. Set a timeout of 100 ms; that leaves 400 ms for the other two downstream APIs to complete
- B. Do not set a timeout; the invocation of this API is mandatory and so we must wait until it responds
- C. Set a timeout of 50 ms; this times out more invocations of that API but gives additional room for retries
- D. No timeout is possible to meet the upstream API's desired SLA; a different SLA must be negotiated with the first downstream API or invoke an alternative API

Answer: D

NEW QUESTION 188

A project team uses RAML specifications to document API functional requirements and deliver API definitions. As per the current legal requirement, all designed API definitions to be augmented with an additional non-functional requirement to protect the services from a high rate of requests according to define service level agreements.

Assuming that the project is following Mulesoft API governance and policies, how should the project team convey the necessary non-functional requirement to stakeholders?

- A. Create proxies in API manager for the non functional requirement and publish to exchange
- B. Add all non functional requirements as comments to RAML specification and publish to exchange
- C. Create various SLA's in API manager for the non functional requirement and publish to exchange
- D. Update API definitions with the fragment for the appropriate policy and publish to exchange

Answer: D

NEW QUESTION 193

A Kubernetes controller automatically adds another pod replica to the resource pool in response to increased application load. Which scalability option is the controller implementing?

- A. Down
- B. Diagonal
- C. Vertical
- D. Horizontal

Answer: D

NEW QUESTION 198

As a part of project, existing java implementation is being migrated to Mulesoft. Business is very tight on the budget and wish to complete the project in most economical way possible.

Canonical object model using java is already a part of existing implementation. Same object model is required by mule application for a business use case. What is the best way to achieve this?

- A. Make use of Java module
- B. Create similar model for Mule applications
- C. Create a custom application to read Java code and make it available for Mule application
- D. Use Anypoint exchange

Answer: A

NEW QUESTION 201

A payment processing company has implemented a Payment Processing API Mule application to process credit card and debit card transactions. Because the Payment Processing API handles highly sensitive information, the payment processing company requires that data must be encrypted both In-transit and at-rest. To meet these security requirements, consumers of the Payment Processing API must create request message payloads in a JSON format specified by the API, and the message payload values must be encrypted.

How can the Payment Processing API validate requests received from API consumers?

- A. A Transport Layer Security (TLS) - Inbound policy can be applied in API Manager to decrypt the message payload and the Mule application implementation can then use the JSON Validation module to validate the JSON data
- B. The Mule application implementation can use the APIkit module to decrypt and then validate the JSON data
- C. The Mule application implementation can use the Validation module to decrypt and then validate the JSON data
- D. The Mule application implementation can use DataWeave to decrypt the message payload and then use the JSON Scheme Validation module to validate the JSON data

Answer: A

NEW QUESTION 203

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 load the properties at startup by invoking that credentials service

Answer: D

NEW QUESTION 204

An organization plans to migrate its deployment environment from an onpremises cluster to a Runtime Fabric (RTF) cluster. The on-premises Mule applications are currently configured with persistent object stores.

There is a requirement to enable Mule applications deployed to the RTF cluster to store and share data across application replicas and through restarts of the entire RTF cluster,

How can these reliability requirements be met?

- A. Replace persistent object stores with persistent VM queues in each Mule application deployment
- B. Install the Object Store pod on one of the cluster nodes
- C. Configure Anypoint Object Store v2 to share data between replicas in the RTF cluster
- D. Configure the Persistence Gateway in the RTF installation

Answer: C

NEW QUESTION 209

Customer has deployed mule applications to different customer hosted mule run times. Mule applications are managed from Anypoint platform.

What needs to be configured to monitor these Mule applications from Anypoint monitoring and what sends monitoring data to Anypoint monitoring?

- A. Enable monitoring of individual applications from runtime manager application settings Runtime manager agent sends monitoring data from the mule applications to Anypoint monitoring
- B. Install runtime manager agent on each mule runtime Runtime manager agent sends monitoring data from the mule applications to Anypoint monitoring
- C. Anypoint monitoring agent on each mule runtime Anypoint monitoring agent sends monitoring data from the mule applications to Anypoint monitoring
- D. By default, Anypoint monitoring agent will be installed on each Mule run time Anypoint Monitoring agent automatically sends monitoring data from the Mule applications to Anypoint monitoring

Answer: A

NEW QUESTION 214

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 cacheUse an object store caching strategy and set expiration interval to 1-hour

Answer: A

NEW QUESTION 219

What limits if a particular Anypoint Platform user can discover an asset in Anypoint Exchange?

- A. Design Center and RAML were both used to create the asset
- B. The existence of a public Anypoint Exchange portal to which the asset has been published
- C. The type of the asset in Anypoint Exchange
- D. The business groups to which the user belongs

Answer: D

NEW QUESTION 223

A stock broking company makes use of CloudHub VPC to deploy Mule applications. Mule application needs to connect to a database application in the customers on-premises corporate data center and also to a Kafka cluster running in AWS VPC.

How is access enabled for the API to connect to the database application and Kafka cluster securely?

- A. Set up a transit gateway to the customers on-premises corporate datacenter to AWS VPC
- B. Setup AnyPoint VPN to the customer's on-premise corporate data center and VPC peering with AWS VPC
- C. Setup VPC peering with AWS VPC and the customers devices corporate data center
- D. Setup VPC peering with the customers onto my service corporate data center and Anypoint VPN to AWS VPC

Answer: B

NEW QUESTION 227

A Mule application is built to support a local transaction for a series of operations on a single database. The Mule application has a Scatter-Gather that participates in the local transaction.

What is the behavior of the Scatter-Gather when running within this local transaction?

- A. Execution of each route within the Scatter-Gather occurs sequentiallyAny error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
- B. Execution of all routes within the Scatter-Gather occurs in parallelAny error that occurs inside the Scatter-Gather will result in a rollback of all the database operations
- C. Execution of each route within the Scatter-Gather occurs sequentiallyAny error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations
- D. Execution of each route within the Scatter-Gather occurs in parallelAny error that occurs inside the Scatter-Gather will NOT result in a rollback of any of the database operations

Answer: A

NEW QUESTION 232

What Anypoint Connectors support transactions?

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

Answer: A

NEW QUESTION 233

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

NEW QUESTION 236

During a planning session with the executive leadership, the development team director presents plans for a new API to expose the data in the company's order database. An earlier effort to build an API on top of this data failed, so the director is recommending a design-first approach. Which characteristics of a design-first approach will help make this API successful?

- A. Building MUnit tests so administrators can confirm code coverage percentage during deployment
- B. Publishing the fully implemented API to Exchange so all developers can reuse the API
- C. Adding global policies to the API so all developers automatically secure the implementation before coding anything
- D. Developing a specification so consumers can test before the implementation is built

Answer: D

NEW QUESTION 239

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 243

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

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

Answer: D

NEW QUESTION 246

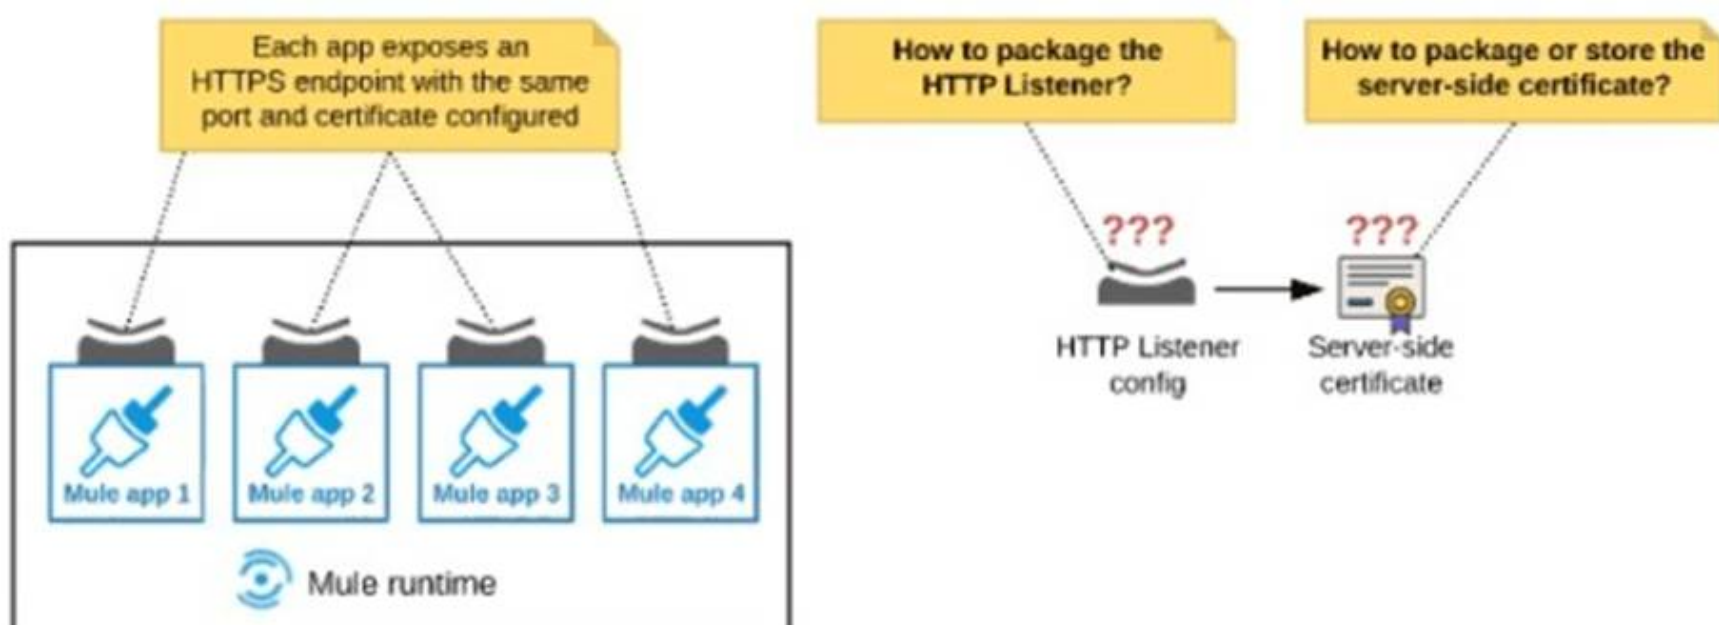
A company is modernizing its legal systems to accelerate access to applications and data while supporting the adoption of new technologies. The key to achieving this business goal is unlocking the companies' key systems and data including microservices running under Docker and Kubernetes containers using APIs. Considering the current aggressive backlog and project delivery requirements the company wants to take a strategic approach in the first phase of its transformation projects by quickly deploying APIs in mule runtime that are able to scale, connect to on-premises systems and migrate as needed. Which runtime deployment option supports company's goals?

- A. Customer hosted self provisioned runtimes
- B. Cloudhub runtimes
- C. Runtime fabric on self managed Kubernetes
- D. Runtime fabric on VMware metal

Answer: C

NEW QUESTION 249

Refer to the exhibit.



An organization deploys multiple Mule applications to the same customer-hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a server-side certificate that rotates often. What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

- A. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint Package the server- side certificate in ALL Mule APPLICATIONS that need to expose an HTTPS endpoint
- B. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing it from all Mule applications that need to expose an HTTPS endpoint
- C. Store the server-side certificate in a shared filesystem location in the Mule runtime's classpath, OUTSIDE the Mule DOMAIN or any Mule APPLICATION
- D. Package an HTTPS Listener configuration In all Mule APPLICATIONS that need to expose an HTTPS endpoint Package the server-side certificate in a NEW Mule DOMAIN project
- E. Package the HTTPS Listener configuration in a Mule DOMAIN project, referencing It from all Mule applications that need to expose an HTTPS endpoint
- F. Package the server- side certificate in the SAME Mule DOMAIN project Go to Set

Answer: B

NEW QUESTION 250

A mule application must periodically process a large dataset which varies from 6 GB to 8 GB from a back-end database and write transform data to an FTPS server using a properly configured batch 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 repeatable 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 255

A mule application is deployed to a Single Cloudhub worker and the public URL appears in Runtime Manager as the APP URL.

Requests are sent by external web clients over the public internet to the mule application App url. Each of these requests routed to the HTTPS Listener event source of the running Mule application.

Later, the DevOps team edits some properties of this running Mule application in Runtime Manager.

Immediately after the new property values are applied in runtime manager, how is the current Mule application deployment affected and how will future web client requests to the Mule application be handled?

- A. Cloudhub will redeploy the Mule application to the OLD Cloudhub workerNew web client requests will RETURN AN ERROR until the Mule application is redeployed to the OLD Cloudhub worker
- B. CloudHub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests will RETURN AN ERROR until the NEW Cloudhub worker is available
- C. Cloudhub will redeploy the Mule application to a NEW Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker until the NEW Cloudhub worker is available.
- D. Cloudhub will redeploy the mule application to the OLD Cloudhub workerNew web client requests are ROUTED to the OLD Cloudhub worker BOTH before and after the Mule application is redeployed.


Answer: C

NEW QUESTION 256

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. RAML

Answer: B

NEW QUESTION 257

One of the backend systems involved by the API implementation enforces rate limits on the number of request a particle client can make. Both the back-end system and API implementation are deployed to several non-production environments including the staging environment and to a particular production environment. Rate limiting of the back-end system applies to all non-production environments. The production environment however does not have any rate limiting. What is the cost-effective approach to conduct performance test of the API implementation in the non-production staging environment?

- A. Including logic within the API implementation that bypasses in locations of the back-end system in the staging environment and invoke a Mocking service that replicates typical back-end system responsesThen conduct performance test using this API implementation
- B. Use MUnit to simulate standard responses from the back-end syste
- C. Then conduct performance test to identify other bottlenecks in the system
- D. Create a Mocking service that replicates the back-end system'sproduction performance characteristicsThen configure the API implementation to use the mocking service and conduct the performance test
- E. Conduct scaled-down performance tests in the staging environment against rate-limiting back-end syste
- F. Then upscale performance results to full production scale

Answer: C

NEW QUESTION 258

In one of the critical payment related mule application, transaction is being used . As an enhancement to implementation , scatter gather route is introduced which is also the part of transaction group. Scatter gather route has 4 routes. What will be the behavior of the Mule application in case of error occurs in 4th route of the scatter-gather router and transaction needs to be rolled back?

- A. Only errored route will be rolled back
- B. All routes will be rolled back
- C. Scatter Gather router cannot be part of transaction

Answer: B

NEW QUESTION 262

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 transactionsApply various compensating actions depending on where a failure occurs


Answer: C

NEW QUESTION 263

Refer to the exhibit.



A Mule application is being designed to be deployed to several CloudHub 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

NEW QUESTION 265

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

NEW QUESTION 270

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 deployed 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

NEW QUESTION 273


An application load balancer routes requests to a RESTful web API secured by Anypoint Flex Gateway. Which protocol is involved in the communication between the load balancer and the Gateway?

- A. SFTP
- B. HTTPS
- C. LDAP
- D. SMTP

Answer: B

NEW QUESTION 276

Refer to the exhibit.



A customer is running Mule applications on Runtime Fabric for Self-Managed Kubernetes (RTF-BYOKS) in a multi-cloud environment. Based on this configuration, how do Agents and Runtime Manager communicate, and what is exchanged between them?

- A. BLOCKING_IO, UBER
- B. UBER, Dedicated NIO Selector Pool
- C. CPU_LITE, CPU_INTENSIVE
- D. Shared NIO Selector Pool, CPU_LITE

Answer: D

NEW QUESTION 278

An organization designing a hybrid, load balanced, single cluster production environment. Due to performance service level agreement goals, it is looking into running the Mule applications in an active-active multi node cluster configuration.

What should be considered when running its Mule applications in this type of environment?

- A. All event sources, regardless of time, can be configured as the target source by the primary node in the cluster
- B. An external load balancer is required to distribute incoming requests throughout the cluster nodes
- C. A Mule application deployed to multiple nodes runs in an isolation from the other nodes in the cluster
- D. Although the cluster environment is fully installed configured and running, it will not process any requests until an outage condition is detected by the primary node in the cluster.

Answer: B

NEW QUESTION 280

The implementation of a Process API must change. What is a valid approach that minimizes the impact of this change on API clients?

- A. Implement required changes to the Process API implementation so that whenever possible, the Process API's RAML definition remains unchanged
- B. Update the RAML definition of the current Process API and notify API client developers by sending them links to the updated RAML definition
- C. Postpone changes until API consumers acknowledge they are ready to migrate to a new Process API or API version
- D. Implement the Process API changes in a new API implementation, and have the old API implementation return an HTTP status code 301 - Moved Permanently to inform API clients they should be calling the new API implementation

Answer: A

NEW QUESTION 282

A developer is examining the responses from a RESTful web service that is compliant with the Hypertext Transfer Protocol (HTTP/1.1) as defined by the Internet Engineering Task Force (IETF).

In this HTTP/1.1-compliant web service, which class of HTTP response status codes should be specified to indicate when client requests are successfully received, understood, and accepted by the web service?

- A. 3xx
- B. 2xx
- C. 4xx
- D. 5xx

Answer: B

NEW QUESTION 283

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 nodes from receiving further HTTP requests.

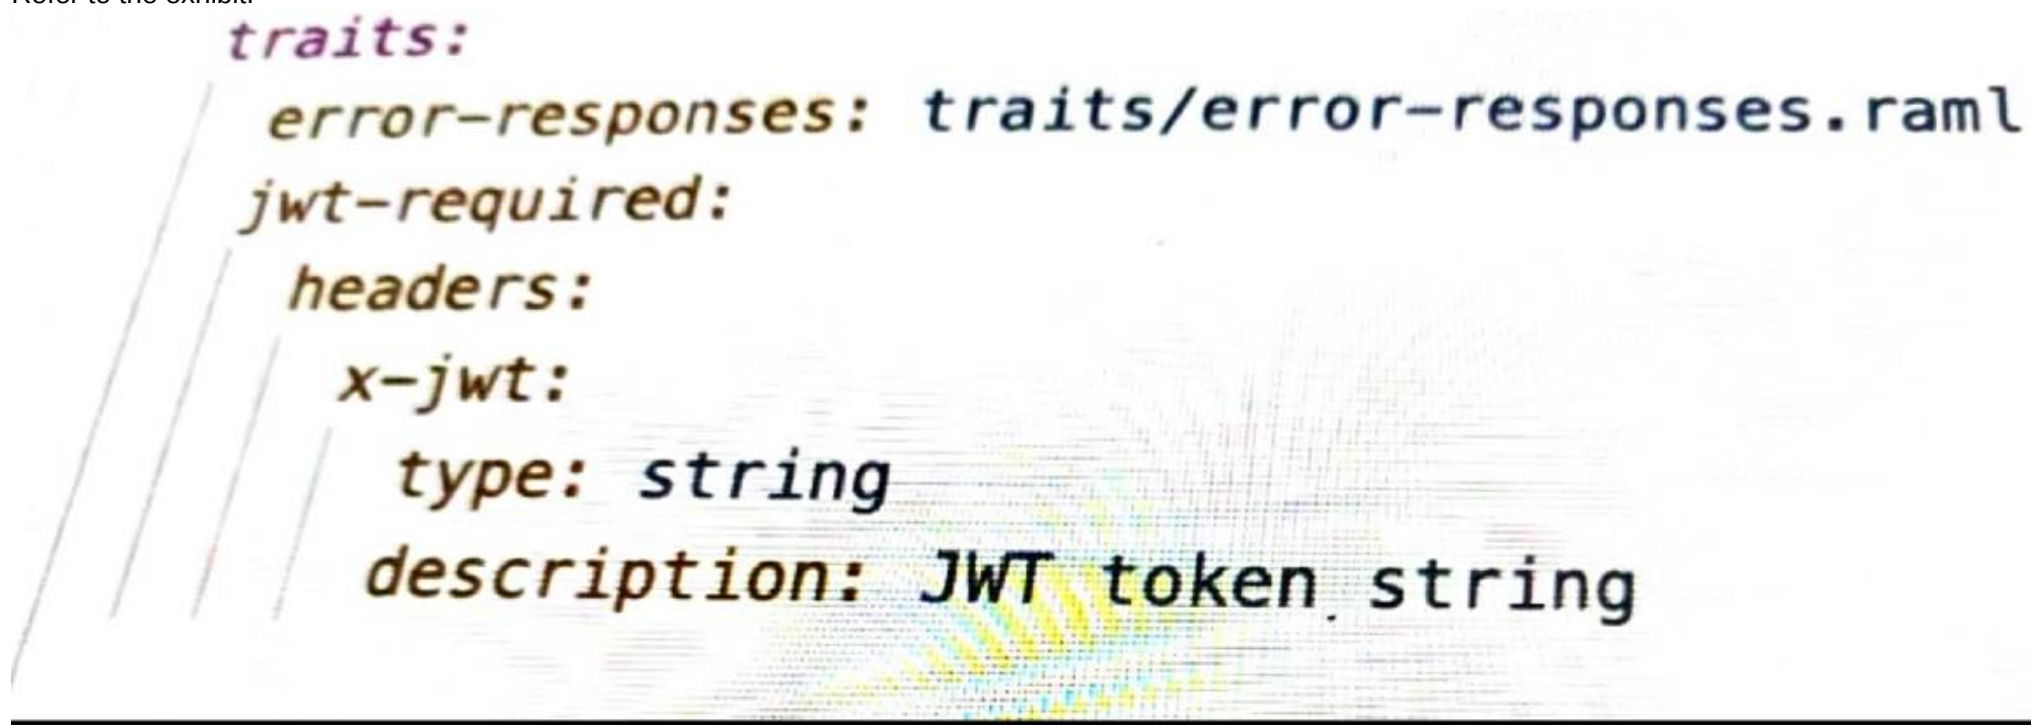
What performance-related consequence is guaranteed to happen to average, assuming the remaining cluster nodes are fully operational?

- 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: A

NEW QUESTION 286

Refer to the exhibit.



```
traits:
  error-responses: traits/error-responses.raml
  jwt-required:
    headers:
      x-jwt:
        type: string
        description: JWT token string
```

What is the type data format shown in the exhibit?

- A. JSON
- B. XML
- C. YAML
- D. CSV

Answer: C

NEW QUESTION 290

What is a key difference between synchronous and asynchronous logging from Mule applications?

- A. Synchronous logging writes log messages in a single logging thread but does not block the Mule event being processed by the next event processor
- B. Asynchronous logging can improve Mule event processing throughput while also reducing the processing time for each Mule event

- C. Asynchronous logging produces more reliable audit trails with more accurate timestamps
- D. Synchronous logging within an ongoing transaction writes log messages in the same thread that processes the current Mule event

Answer: B

NEW QUESTION 292

What is true about automating interactions with Anypoint Platform using tools such as Anypoint Platform REST API's, Anypoint CLI or the Mule Maven plugin?

- A. By default, the Anypoint CLI and Mule Maven plugin are not included in the Mule runtime
- B. Access to Anypoint Platform API's and Anypoint CLI can be controlled separately through the roles and permissions in Anypoint platform, so that specific users can get access to Anypoint CLI while others get access to the platform API's
- C. Anypoint Platform API's can only automate interactions with CloudHub while the Mule maven plugin is required for deployment to customer hosted Mule runtimes
- D. API policies can be applied to the Anypoint platform API's so that only certain LOS's has access to specific functions

Answer: A

NEW QUESTION 296

An organization is sizing an Anypoint VPC to extend their internal network to Cloudhub.

For this sizing calculation, the organization assumes 150 Mule applications will be deployed among three(3) production environments and will use Cloudhub's default zero- downtime feature. Each Mule application is expected to be configured with two(2) Cloudhub workers. This is expected to result in several Mule application deployments per hour.

- A. 10.0.0.0/21(2048 IPs)
- B. 10.0.0.0/22(1024IPs)
- C. 10.0.0.0/23(512 IPs)
- D. 10.0.0.0/24(256 IPs)

Answer: A

NEW QUESTION 300

What condition requires using a CloudHub Dedicated Load Balancer?

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

Answer: D

NEW QUESTION 301

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

MuleSoft-Integration-Architect-I Practice Exam Features:

- * MuleSoft-Integration-Architect-I Questions and Answers Updated Frequently
- * MuleSoft-Integration-Architect-I Practice Questions Verified by Expert Senior Certified Staff
- * MuleSoft-Integration-Architect-I Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * MuleSoft-Integration-Architect-I Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The MuleSoft-Integration-Architect-I Practice Test Here](#)