



Salesforce

Exam Questions MuleSoft-Integration-Architect-I

Salesforce Certified MuleSoft Integration Architect 1 (SP24) Exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

An organization plans to use the Anypoint Platform audit logging service to log Anypoint MQ actions. What consideration must be kept in mind when leveraging Anypoint MQ Audit Logs?

- A. Anypoint MQ Audit Logs include logs for sending, receiving, or browsing messages
- B. Anypoint MQ Audit Logs include logs for failed Anypoint MQ operations
- C. Anypoint MQ Audit Logs include logs for queue create, delete, modify, and purge operations

Answer: C

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

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

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

Answer: D

NEW QUESTION 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

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 7

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 8

Which Anypoint Platform component helps integration developers discover and share reusable APIs, connectors, and templates?

- A. Anypoint Exchange
- B. API Manager
- C. Anypoint Studio
- D. Design Center

Answer: A

NEW QUESTION 9

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: B

NEW QUESTION 10

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 10

An organization has decided on a cloudhub migration strategy that aims to minimize the organizations own IT resources. Currently, the organization 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 14

A Mule application is synchronizing customer data between two different database systems.

What is the main benefit of using XA transaction over local transactions to synchronize these two database system?

- A. Reduce latency
- B. Increase throughput
- C. Simplifies communication
- D. Ensure consistency

Answer: D

NEW QUESTION 15

A Mule application is running on a customer-hosted Mule runtime in an organization's network. The Mule application acts as a producer of asynchronous Mule events. Each Mule event must be broadcast to all interested external consumers outside the Mule application. The Mule events should be published in a way that is guaranteed in normal situations and also minimizes duplicate delivery in less frequent failure scenarios.

The organizational firewall is configured to only allow outbound traffic on ports 80 and 443. Some external event consumers are within the organizational network, while others are located outside the firewall.

What Anypoint Platform service is most idiomatic (used for its intended purpose) for publishing these Mule events to all external consumers while addressing the desired reliability goals?

- A. CloudHub VM queues
- B. Anypoint MQ
- C. Anypoint Exchange
- D. CloudHub Shared Load Balancer

Answer: B

NEW QUESTION 16

A rate limiting policy has been applied to a soap V1.2 API published in CloudHub. The API implementation catches errors in a global error handler on error propagate in the main flow for HTTP: RETRY_EXHAUSTED with HTTP status set to 429 and any with the HTTP status set to 500.

What is the expected HTTP status when the client exceeds the quota of the API calls?

- A. HTTP status 429 as defined in the HTTP:RETRY EXHAUSTED error handler in the API
- B. HTTP status 500 as defined in the ANY error handler in the API since an API:RETRY_EXHAUSTED will be generated
- C. HTTP status 401 unauthorized for policy violation
- D. HTTP status 400 from the rate-limiting policy violation since the call does not reach the back-end


Answer: A

NEW QUESTION 21

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 center


B. In Runtime 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 24

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 26

According to MuleSoft's API development best practices, which type of API development approach starts with writing and approving an API contract?

- A. Implement-first
 B. Catalyst
 C. Agile
 D. Design-first

Answer: D

NEW QUESTION 29

A global organization operates datacenters in many countries. There are private network links between these datacenters because all business data (but NOT metadata) must be exchanged over these private network connections.

The organization does not currently use AWS in any way.

The strategic decision has Just been made to rigorously minimize IT operations effort and investment going forward.

What combination of deployment options of the Anypoint Platform control plane and runtime plane(s) best serves this organization at the start of this strategic journey?

- A. MuleSoft-hosted Anypoint Platform control plane CloudHub Shared Worker Cloud in multiple AWS regions
- B. Anypoint Platform - Private Cloud Edition Customer-hosted runtime plane in each datacenter
- C. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in multiple AWS regions
- D. MuleSoft-hosted Anypoint Platform control plane Customer-hosted runtime plane in each datacenter

Answer: D

NEW QUESTION 30

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 35

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 39

In a Mule Application, a flow contains two (2) JMS consume operations that are used to connect to a JMS broker and consume messages from two(2) JMS destination. The Mule application then joins the two JMS messages together.

The JMS broker does not implement high availability (HA) and periodically experiences scheduled outages of upto 10 mins for routine maintenance.

What is the most idiomatic (used for its intended purpose) way to build the mule flow so it can best recover from the expected outages?

- A. Configure a reconnection strategy for the JMS connector
- B. Enclose the two(2) JMS operation in an Until Successful scope
- C. Consider a transaction for the JMS connector
- D. Enclose the two(2) JMS operations in a Try scope with an Error Continue error handler

Answer: A

NEW QUESTION 40

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

NEW QUESTION 44

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 47

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 52

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 54


According to MuleSoft, which principle is common to both Service Oriented Architecture (SOA) and API-led connectivity approaches?

- A. Service centralization
- B. Service statefulness
- C. Service reusability
- D. Service interdependence

Answer: C

NEW QUESTION 57

Refer to the exhibit.



A business process involves two APIs that interact with each other asynchronously over HTTP. Each API is implemented as a Mule application. API 1 receives the initial HTTP request and invokes API 2 (in a fire and forget fashion) while API 2, upon completion of the processing, calls back into API 1 to notify about completion of the asynchronous process.

Each API is deployed to multiple redundant Mule runtimes and a separate load balancer, and is deployed to a separate network zone.

In the network architecture, how must the firewall rules be configured to enable the above interaction between API 1 and API 2?

- A. To authorize the certificate to be used both APIs
- B. To enable communication from each API's Mule Runtimes and Network zone to the load balancer of the other API
- C. To open direct two-way communication between the Mule Runtimes of both APIs
- D. To allow communication between load balancers used by each API

Answer: B

NEW QUESTION 61

A bank is implementing a REST API in a Mule application to receive an array of accounts from an online banking platform user interface (UI), retrieve account balances for those accounts from a backend Finance system, and then return the account balances so they can be displayed in the online banking platform UI. As part of the processing, the MuleSoft API also needs to insert the retrieved account data into an AuditDatabase for auditing purposes. The auditing process should not add latency to the account balance retrieval response back to the online banking platform UI.

The retrieveBalances flow in the Mule application is designed to use an operation in a connector to the Finance system (the Finance operation) that can only look up one account record at a time, and an operation from a different connector to the Audit system (the Audit operation) that can only insert one account record at a time.

To best meet the performance-related requirements, what scope or scopes should be used and how should they be used to incorporate the Finance operation and Audit operation into the retrieveBalances flow?

- A. Wrap the Finance operation in a Parallel For-Each scope
- B. Wrap the Audit operation in an Async scope.
- C. Wrap the Finance operation in an Until-Successful scope
- D. Wrap the Audit operation in a Try-Catch scope.
- E. Wrap both connector operations in an Async scope.
- F. Wrap both connector operations in a For-Each scope.

Answer: A

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 65

An organization needs to procure an enterprise software system to increase cross-selling opportunities and better track prospect data.

Which category of enterprise software has these core capabilities, when used for its typical and intended purpose?

- A. Supply Chain Management (SCM)
- B. IT Service Management (ITSM)
- C. Business-to-Business (B2B)
- D. Customer Relationship Management (CRM)

Answer: D

NEW QUESTION 70

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 71

Which productivity advantage does Anypoint Platform have to both implement and manage an API?

- A. Automatic API proxy generation
- B. Automatic API specification generation
- C. Automatic API semantic versioning
- D. Automatic API governance

Answer: A

NEW QUESTION 75

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 78

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 80

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 82

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 83

What are two reasons why a typical MuleSoft customer favors a MuleSoft-hosted Anypoint Platform runtime plane over a customer-hosted runtime for its Mule application deployments? (Choose two.)

- A. Reduced application latency
- B. Increased application isolation
- C. Reduced time-to-market for the first application
- D. Increased application throughput
- E. Reduced IT operations effort

Answer: CE

NEW QUESTION 85

As part of a growth strategy, a supplier signs a trading agreement with a large customer. The customer sends purchase orders to the supplier according to the ANSI X12 EDI standard, and the supplier creates the orders in its ERP system using the information in the EDI document.

The agreement also requires that the supplier provide a new RESTful API to process request from the customer for current product inventory level from the supplier's ERP system.

Which two fundamental integration use cases does the supplier need to deliver to provide an end-to-end solution for this business scenario? (Choose two.)

- A. Synchronized data transfer
- B. Sharing data with external partners
- C. User interface integration
- D. Streaming data ingestion
- E. Data mashups

Answer: AB

NEW QUESTION 88

An organization will deploy Mule applications to Cloudhub, Business requirements mandate that all application logs be stored ONLY in an external splunk consolidated logging service and NOT in Cloudhub.

In order to most easily store Mule application logs ONLY in Splunk, how must Mule application logging be configured in Runtime Manager, and where should the log4j2 splunk appender be defined?

- A. Keep the default logging configuration in RuntimeManagerDefine the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manager to

support at Mule application deployments.

B. Disable Cloudhub logging in Runtime Manager Define the splunk appender in EACH Mule application??s log4j2.xml file

C. Disable Cloudhub logging in Runtime Manager Define the splunk appender in ONE global log4j.xml file that is uploaded once to Runtime Manger to support at Mule application deployments.

D. Keep the default logging configuration in Runtime Manager Define the Splunk appender in EACH Mule application log4j2.xml file

Answer: B

NEW QUESTION 93

An organization has deployed runtime fabric on an eight node cluster with performance profile. An API uses and non persistent object store for maintaining some of its state data. What will be the impact to the stale data if server crashes?

A. State data is preserved

B. State data is rolled back to a previously saved version

C. State data is lost

D. State data is preserved as long as more than one more is unaffected by the crash

Answer: C

NEW QUESTION 94

A manufacturing company plans 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 choice of runtime plane meets these organizational requirements?

A. CloudHub 2.0

B. Customer-hosted self-provisioned runtime plane


C. Anypoint Service Mesh

D. Anypoint Runtime Fabric

Answer: D

NEW QUESTION 98

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 100

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 102

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 106

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 108

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 113

As a part of business requirement, old CRM system needs to be integrated using Mule application. CRM system is capable of exchanging data only via SOAP/HTTP protocol. As an integration architect who follows API led approach, what is the below step you will perform so that you can share document with CRM team?

- A. Create RAML specification using Design Center
- B. Create SOAP API specification using Design Center
- C. Create WSDL specification using text editor
- D. Create WSDL specification using Design Center

Answer: C

NEW QUESTION 114

An application deployed to a runtime fabric environment with two cluster replicas is designed to periodically trigger a 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 process all 1,00,000 records from scratch leading to duplicate record processing
- D. A new placement replica will be available and will take over processing the remaining 99,000 records

Answer: D

NEW QUESTION 118

The company's FTPS server login username and password

- A. TLS context trust store containing a public certificate for the compan
- B. The company's PGP public key that was used to sign the files
- C. The partner's PGP public key used by the company to login to the FTPS serve
- D. A TLS context key store containing the private key for the companyThe partner's PGP private key that was used to sign the files
- E. The company's FTPS server login username and passwor
- F. A TLS context trust store containing a public certificate for ftps.partner.comThe partner's PGP public key that was used to sign the files
- G. The partner's PGP public key used by the company to login to the FTPS serve
- H. A TLS context key store containing the private key for ftps.partner.comThe company's PGP private key that was used to sign the files

Answer: C

NEW QUESTION 121

According to MuleSoft, which major benefit does a Center for Enablement (C4E) provide for an enterprise and its lines of business?

- A. Enabling Edge security between the lines of business and public devices
- B. Centralizing project management across the lines of business
- C. Centrally managing return on investment (ROI) reporting from lines of business to leadership
- D. Accelerating self-service by the lines of business

Answer: D

NEW QUESTION 122

An organization has an HTTPS-enabled Mule application named Orders API that receives requests from another Mule application named Process Orders. The communication between these two Mule applications must be secured by TLS mutual authentication (two-way TLS).

At a minimum, what must be stored in each truststore and keystore of these two Mule applications to properly support two-way TLS between the two Mule applications while properly protecting each Mule application's keys?

- A. Orders API truststore: The Orders API public keyProcess Orders keystore: The Process Orders private key and public key
- B. Orders API truststore: The Orders API private key and public key Process Orders keystore: The Process Orders private key public key
- C. Orders API truststore: The Process Orders public keyOrders API keystore: The Orders API private key and public key Process Orders truststore: The Orders API public keyProcess Orders keystore: The Process Orders private key and public key
- D. Orders API truststore: The Process Orders public key Orders API keystore: The Orders API private key Process Orders truststore: The Orders API public keyProcess Orders keystore: The Process Orders private key

Answer: C

NEW QUESTION 123

A manufacturing company is developing a new set of APIs for its retail business. One of the APIs is a Master Look Up API, which is a System API, The API uses a persistent object-store. This API will be used by almost all other APIs to provide master lookup data.

The Master Look Up API is deployed on two CloudHub workers of 0.1 vCore each because there is a lot of master data to be cached. Most of the master lookup data is stored as a key-value pair. The cache gets refreshed if the key is not found in the cache.

During performance testing, it was determined that the Master Look Up API has a high response time due to the latency of database queries executed to fetch the master lookup data.

What two methods can be used to resolve these performance issues?

Choose 2 answers

- A. Implement the HTTP caching policy for all GET endpoints for the Master Look Up API
- B. Implement an HTTP caching policy for all GET endpoints in the Master Look Up API
- C. Implement locking to synchronize access to the Object Store
- D. Upgrade the vCore size from 0.1 vCore to 0.2 vCore

Answer: AD

NEW QUESTION 127

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 availabilty 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 129

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

NEW QUESTION 134

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 135

A large life sciences customer plans to use the Mule Tracing module with the Mapped Diagnostic Context (MDC) logging operations to enrich logging in its Mule application and to improve tracking by providing more context in the Mule application logs. The customer also wants to improve throughput and lower the message processing latency in its Mule application flows.

After installing the Mule Tracing module in the Mule application, how should logging be performed in flows in Mule applications, and what should be changed in the log4j2.xml files?

- A. In the flows, add Mule Tracing module Set logging variable operations before any Core Logger components. In log4j2.xml files, change the appender's pattern layout to use %MDC and then assign the appender to a Logger or Root element.
- B. In the flows, add Mule Tracing module Set logging variable operations before any Core Logger components. In log4j2.xml files, change the appender's pattern layout to use the %MDC placeholder and then assign the appender to an AsyncLogger element.
- C. In the flows, add Mule Tracing module Set logging variable operations before any Core Logger components. In log4j2.xml files, change the appender's pattern layout to use %asyncLogger placeholder and then assign the appender to an AsyncLogger element.
- D. In the flows, wrap Logger components in Async scope
- E. In log4j2.xml files, change the appender's pattern layout to use the %asyncLoggerplaceholder and then assign the appender to a Logger or Root element.

Answer: A

NEW QUESTION 137

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 139

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 scope that participates in the local transaction.

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

- A. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
- B. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back
- C. Execution of all routes within Scatter-Gather occurs sequentially Any error that occurs inside Scatter-Gather will result in a roll back of all the database operations
- D. Execution of all routes within Scatter-Gather occurs in parallel Any error that occurs inside Scatter-Gather will be handled by error handler and will not result in roll back

Answer: A

NEW QUESTION 144

An organization is designing the following two Mule applications that must share data via a common persistent object store instance:

- Mule application P will be deployed within their on-premises datacenter.
- Mule application C will run on CloudHub in an Anypoint VPC.

The object store implementation used by CloudHub is the Anypoint Object Store v2 (OSv2).

what type of object store(s) should be used, and what design gives both Mule applications access to the same object store instance?

- A. Application P uses the Object Store connector to access a persistent object store Application C accesses this persistent object store via the Object Store REST API through an IPsec tunnel
- B. Application C and P both use the Object Store connector to access the Anypoint Object Store v2
- C. Application C uses the Object Store connector to access a persistent object Application P accesses the persistent object store via the Object Store REST API
- D. Application C and P both use the Object Store connector to access a persistent object store

Answer: C

NEW QUESTION 148

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 transaction 2) 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 152

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. Use async 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. Use async logger at the level equal to DEBUG or TRACE and use pattern layout 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 154

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 157

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 158

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 162

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 163

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 processing There 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 processing There 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: A

NEW QUESTION 165

A stock trading company handles millions of trades a day and requires excellent performance and reliability within its stock trading system. The company operates a number of event-driven APIs Implemented as Mule applications that are hosted on various customer-hosted Mule clusters and needs to enable message exchanges between the APIs within their internal network using shared message queues.

What is an effective way to meet the cross-cluster messaging requirements of its event- driven APIs?

- A. Non-transactional JMS operations with a reliability pattern and manual acknowledgements
- B. Persistent VM queues with automatic acknowledgements
- C. JMS transactions with automatic acknowledgements
- D. extended Architecture (XA) transactions and XA connected components with manual acknowledgements

Answer: C

NEW QUESTION 170

What is a core pillar of the MuleSoft Catalyst delivery approach?

- A. Business outcomes
- B. Technology centralization
- C. Process thinking
- D. Scope reduction

Answer: A

NEW QUESTION 171

Which Exchange asset type represents a complete API specification in RAML or OAS format?

- A. Connectors
- B. REST APIs
- C. API Spec Fragments
- D. SOAP APIs

Answer: B

NEW QUESTION 176

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 177

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

NEW QUESTION 182

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 187

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

NEW QUESTION 191

49 of A popular retailer is designing a public API for its numerous business partners. Each business partner will invoke the API at the URL 58.

<https://api.acme.com/partnefs/v1>. The API implementation is estimated to require deployment to 5 CloudHub workers.

The retailer has obtained a public X.509 certificate for the name apl.acme.com, signed by a reputable CA, to be used as the server certificate.

Where and how should the X.509 certificate and Mule applications be used to configure load balancing among the 5 CloudHub workers, and what DNS entries should be configured in order for the retailer to support its numerous business partners?

- A. Add the X.509 certificate to the Mule application's deployable archive, then configure a CloudHub Dedicated Load Balancer (DLB) for each of the Mule application's CloudHub workersCreate a CNAME for api.acme.com pointing to the DLB's A record
- B. Add the X.509 certificate to the CloudHub Shared Load Balancer (SLB), not to the Mule applicationCreate a CNAME for api.acme.com pointing to the SLB's A record
- C. Add the X.509 certificate to a CloudHub Dedicated Load Balancer (DLB), not to the Mule applicationCreate a CNAME for api.acme.com pointing to the DLB's A record
- D. Add the x.509 certificate to the Mule application's deployable archive, then configure the CloudHub Shared Load Balancer (SLB) for each of the Mule application's CloudHub workersCreate a CNAME for api.acme.com pointing to the SLB's A record

Answer: C

NEW QUESTION 195

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 198

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

NEW QUESTION 200

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 202

Which key DevOps practice and associated Anypoint Platform component should a MuteSoft integration team adopt to improve delivery quality?

- A. A Continuous design with API Designer
- B. Automated testing with MUnit
- C. Passive monitoring with Anypoint Monitoring
- D. Manual testing with Anypoint Studio

Answer: B

NEW QUESTION 206

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

NEW QUESTION 210

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 application
- D. Use VM queues in the Mule application to allow any non-mule assets within the customer provisioned AWS VPC to subscribe to and receive messages

Answer: A

NEW QUESTION 214

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 216

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 220

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

NEW QUESTION 223

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 226

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 228

According to the National Institute of Standards and Technology (NIST), which cloud computing deployment model describes a composition of two or more distinct clouds that support data and application portability?

- A. Private cloud
- B. Hybrid cloud
- C. Public cloud
- D. Community cloud

Answer: B

NEW QUESTION 232

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

- A. sObject Platform Action API
- B. User Interface API
- C. Metadata API
- D. Process Rules API

Answer: C

NEW QUESTION 237

What is a recommended practice when designing an integration Mule 4 application that reads a large XML payload as a stream?

- A. The payload should be dealt with as a repeatable XML stream, which must only be traversed (iterated-over) once and CANNOT be accessed randomly from DataWeave expressions and scripts
- B. The payload should be dealt with as an XML stream, without converting it to a single Java object (POJO)
- C. The payload size should NOT exceed the maximum available heap memory of the Mule runtime on which the Mule application executes
- D. The payload must be cached using a Cache scope if it is to be sent to multiple backend systems

Answer: C

NEW QUESTION 239


One of the backend systems involved by the API implementation enforces rate limits on the number of requests a particular 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 responses Then conduct performance test using this API implementation
- B. Use MUnit to simulate standard responses from the back-end system
- C. Then conduct performance test to identify other bottlenecks in the system
- D. Create a Mocking service that replicates the back-end system's production performance characteristics Then 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 system
- F. Then upscale performance results to full production scale

Answer: C

NEW QUESTION 244

Refer to the exhibit.



An organization uses a 2-node Mule runtime cluster to host one stateless API implementation. The API is accessed over HTTPS through a load balancer that uses round-robin for load distribution.

Two additional nodes have been added to the cluster and the load balancer has been configured to recognize the new nodes with no other change to the load balancer.

What average performance change is guaranteed to happen, assuming all cluster nodes are fully operational?

- A. 50% reduction in the response time of the API
- B. 100% increase in the throughput of the API
- C. 50% reduction in the JVM heap memory consumed by each node

D. 50% reduction In the number of requests being received by each node

Answer: D

NEW QUESTION 248

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 routed will be rolled back
- C. Scatter Gather router cannot be part of transaction

Answer: B

NEW QUESTION 251


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 255

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 259

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

NEW QUESTION 260

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 265


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: A

NEW QUESTION 270

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 275

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 278

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 application Change 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 application Change the naming convention to {API name}-{environment}-{region} and communicate this change to the consuming applications and users

Answer: D

NEW QUESTION 280

According to MuleSoft, what Action should an IT organization take regarding its technology assets in order to close the IT delivery.

- A. Make assets easily discoverable via a central repository
- B. Focus project delivery efforts on custom assets that meet the specific requirements of each individual line of business
- C. Create weekly meetings that all members of IT attend to present justification and request approval to use existing assets
- D. Hire additional staff to meet the demand for asset creation required for approved projects and timelines

Answer: A

NEW QUESTION 281

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 282


What is a defining characteristic of an integration-Platform-as-a-Service (iPaaS)?

- A. A Cloud-based
- B. No-code
- C. Code-first
- D. On-premises

Answer: A

NEW QUESTION 286

Refer to the exhibit.



A Mule application has an HTTP Listener that accepts HTTP DELETE requests. This Mule application is deployed to three CloudHub workers under the control of the CloudHub Shared Load Balancer.

A web client makes a sequence of requests to the Mule application's public URL.

How is this sequence of web client requests distributed among the HTTP Listeners running in the three CloudHub workers?

- A. Each request is routed to the PRIMARY CloudHub worker in the PRIMARY Availability Zone (AZ)
- B. Each request is routed to ONE ARBITRARY CloudHub worker in the PRIMARY Availability Zone (AZ)

- C. Each request is routed to ONE ARBITRARY CloudHub worker out of ALL three CloudHub workers
- D. Each request is routed (scattered) to ALL three CloudHub workers at the same time

Answer: C

NEW QUESTION 291

A new Mule application under development must implement extensive data transformation logic. Some of the data transformation functionality is already available as external transformation services that are mature and widely used across the organization; the rest is highly specific to the new Mule application.

The organization follows a rigorous testing approach, where every service and application must be extensively acceptance tested before it is allowed to go into production.

What is the best way to implement the data transformation logic for this new Mule application while minimizing the overall testing effort?

- A. Implement and expose all transformation logic as mlaoservices using DataWeave, so it can be reused by any application component that needs it, including the new Mule application
- B. Implement transformation logic in the new Mute application using DataWeave, replicating the transformation logic of existing transformation services
- C. Extend the existing transformation services with new transformation logic and Invoke them from the new Mule application
- D. Implement transformation logic in the new Mute application using DataWeave, invoking existing transformation services when possible

Answer: D

NEW QUESTION 292

What is required before an API implemented using the components of Anypoint Platform can be managed and governed (by applying API policies) on Anypoint Platform?

- A. The API must be published to Anypoint Exchange and a corresponding API instance ID must be obtained from API Manager to be used in the API implementation
- B. The API implementation source code must be committed to a source control management system (such as GitHub)
- C. A RAML definition of the API must be created in API designer so it can then be published to Anypoint Exchange
- D. The API must be shared with the potential developers through an API portal so API consumers can interact with the API

Answer: A

NEW QUESTION 296

A leading e-commerce giant will use Mulesoft API's on runtime fabric (RTF) to process customer orders. Some customer's sensitive information such as credit card information is also there as a part of a API payload.

What approach minimizes the risk of matching sensitive data to the original and can convert back to the original value whenever and wherever required?

- A. Apply masking to hide the sensitive information and then use API
- B. manager to detokenize the masking format to return the original value
- C. create a tokenization format and apply a tokenization policy to the API Gateway
- D. Used both masking and tokenization
- E. Apply a field level encryption policy in the API Gateway

Answer: C

NEW QUESTION 300

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 303

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

NEW QUESTION 307

.....

Relate Links

100% Pass Your MuleSoft-Integration-Architect-I Exam with ExamBible Prep Materials

<https://www.exambible.com/MuleSoft-Integration-Architect-I-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/>