

MCIA-Level-1 Dumps

MuleSoft Certified Integration Architect - Level 1

<https://www.certleader.com/MCIA-Level-1-dumps.html>



NEW QUESTION 1

Mule applications need to be deployed to CloudHub so they can access on-premises database systems. These systems store sensitive and hence tightly protected data, so are not accessible over the internet.

What network architecture supports this requirement?

- A. An Anypoint VPC connected to the on-premises network using an IPsec tunnel or AWSDirectConnect, plus matching firewall rules in the VPC and on-premises network
- B. Static IP addresses for the Mule applications deployed to the CloudHub Shared Worker Cloud, plus matching firewall rules and IPwhitelisting in the on-premises network
- C. An Anypoint VPC with one Dedicated Load Balancer fronting each on-premises database system, plus matching IP whitelisting in the load balancer and firewall rules in the VPC and on-premises network
- D. Relocation of the database systems to a DMZ in the on-premises network, with Mule applications deployed to the CloudHub Shared Worker Cloud connecting only to the DMZ

Answer: A

NEW QUESTION 2

An API client is implemented as a Mule application that includes an HTTP Request operation using a default configuration. The HTTP Request operation invokes an external API that follows standard HTTP status code conventions, which causes the HTTP Request operation to return a 4xx status code.

What is a possible cause of this status code response?

- A. An error occurred inside the external API implementation when processing the HTTP request that was received from the outbound HTTP Request operation of the Mule application
- B. The external API reported that the API implementation has moved to a different external endpoint
- C. The HTTP response cannot be interpreted by the HTTP Request operation of the Mule application after it was received from the external API
- D. The external API reported an error with the HTTP request that was received from the outbound HTTP Request operation of the Mule application

Answer: D

NEW QUESTION 3


An API has been unit tested and is ready for integration testing. The API is governed by a Client ID Enforcement policy in all environments. What must the testing team do before they can start integration testing the API in the Staging environment?

- A. They must access the API portal and create an API notebook using the Client ID and Client Secret supplied by the API portal in the Staging environment
- B. They must request access to the API instance in the Staging environment and obtain a Client ID and Client Secret to be used for testing the API
- C. They must be assigned as an API version owner of the API in the Staging environment
- D. They must request access to the Staging environment and obtain the Client ID and Client Secret for that environment to be used for testing the API

Answer: B

NEW QUESTION 4

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 5

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 6

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 Mute 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: B

NEW QUESTION 7

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


What is an expected outcome of this transition to container-based Mule application deployments?

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

Answer: A

NEW QUESTION 8

Refer to the exhibit.



A Mule application is deployed to a multi-node Mule runtime cluster. The Mule application uses the competing consumer pattern among its cluster replicas to

receive JMS messages from a JMS queue. To process each received JMS message, the following steps are performed in a flow:

Step 1: The JMS Correlation ID header is read from the received JMS message.

Step 2: The Mule application invokes an idempotent SOAP webservice over HTTPS, passing the JMS Correlation ID as one parameter in the SOAP request.

Step 3: The response from the SOAP webservice also returns the same JMS Correlation ID.

Step 4: The JMS Correlation ID received from the SOAP webservice is validated to be identical to the JMS Correlation ID received in Step 1.

Step 5: The Mule application creates a response JMS message, setting the JMS Correlation ID message header to the validated JMS Correlation ID and publishes that message to a response JMS queue.

Where should the Mule application store the JMS Correlation ID values received in Step 1 and Step 3 so that the validation in Step 4 can be performed, while also making the overall Mule application highly available, fault-tolerant, performant, and maintainable?

- A. Both Correlation ID values should be stored in a persistent object store
- B. Both Correlation ID values should be stored in a non-persistent object store
- C. The Correlation ID value in Step 1 should be stored in a persistent object store. The Correlation ID value in step 3 should be stored as a Mule event variable/attribute
- D. Both Correlation ID values should be stored as Mule event variables/attributes

Answer: C

NEW QUESTION 9

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 RAML archetype and reusable trait definitions to be reused across API implementations
- B. A custom asset with the default API implementation
- 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 10

What Anypoint Connectors support transactions?

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

Answer: A

NEW QUESTION 10

An organization has various integrations implemented as Mule applications. Some of these Mule applications are deployed to customhosted 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 11

Additional nodes are being added to an existing customer-hosted Mule runtime cluster to improve performance. Mule applications deployed to this cluster are invoked by API clients through a load balancer.

What is also required to carry out this change?

- A. A new load balancer must be provisioned to allow traffic to the new nodes in a round-robin fashion
- B. External monitoring tools or log aggregators must be configured to recognize the new nodes
- C. API implementations using an object store must be adjusted to recognize the new nodes and persist to them
- D. New firewall rules must be configured to accommodate communication between API clients and the new nodes

Answer: C

NEW QUESTION 14

What metrics about API invocations are available for visualization in custom charts using Anypoint Analytics?

- A. Request size, request HTTP verbs, response time
- B. Request size, number of requests, JDBC Select operation result set size
- C. Request size, number of requests, JDBC Select operation response time
- D. Request size, number of requests, response size, response time

Answer: D

NEW QUESTION 16

An organization is designing an integration solution to replicate financial transaction data from a legacy system into a data warehouse (DWH). The DWH must contain a daily snapshot of financial transactions, to be delivered as a CSV file. Daily transaction volume exceeds tens of millions of records, with significant spikes in volume during popular shopping periods. What is the most appropriate integration style for an integration solution that meets the organization's current requirements?

- A. API-led connectivity
- B. Batch-triggered ETL
- C. Event-driven architecture
- D. Microservice architecture

Answer: D

NEW QUESTION 17

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

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

Answer: B

NEW QUESTION 19

An Order microservice and a Fulfillment microservice are being designed to communicate with their clients through message-based integration (and NOT through API invocations). The Order microservice publishes an Order message (a kind of command message) containing the details of an order to be fulfilled. The intention is that Order messages are only consumed by one Mule application, the Fulfillment microservice. The Fulfillment microservice consumes Order messages, fulfills the order described therein, and then publishes an OrderFulfilled message (a kind of event message). Each OrderFulfilled message can be consumed by any interested Mule application, and the Order microservice is one such Mule application. What is the most appropriate choice of message broker(s) and message destination(s) in this scenario?

- A. Order messages are sent to an Anypoint MQ exchangeOrderFulfilled messages are sent to an Anypoint MQ queueBoth microservices interact with Anypoint MQ as the message broker, which must therefore scale to support the load of both microservices
- B. Order messages are sent to a JMS queueOrderFulfilled messages are sent to a JMS topicBoth microservices interact with the same JMS provider (message broker) instance, which must therefore scale to support the load of both microservices
- C. Order messages are sent directly to the Fulfillment microserviceOrderFulfilled messages are sent directly to the Order microserviceThe Order microservice interacts with one AMQP-compatible message broker and the Fulfillment microservice interacts with a different AMQP-compatible message broker, so that both message brokers can be chosen and scaled to best support the load of each microservice
- D. Order messages are sent to a JMS queueOrderFulfilled messages are sent to a JMS topicThe Order microservice interacts with one JMS provider (message broker) and the Fulfillment microservice interacts with a different JMS provider, so that both message brokers can be chosen and scaled to best support the load of each microservice

Answer: D

NEW QUESTION 21

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

Answer: B

NEW QUESTION 26

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

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

Answer: A

NEW QUESTION 27

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your MCIA-Level-1 Exam with Our Prep Materials Via below:

<https://www.certleader.com/MCIA-Level-1-dumps.html>