

MuleSoft

Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1



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.

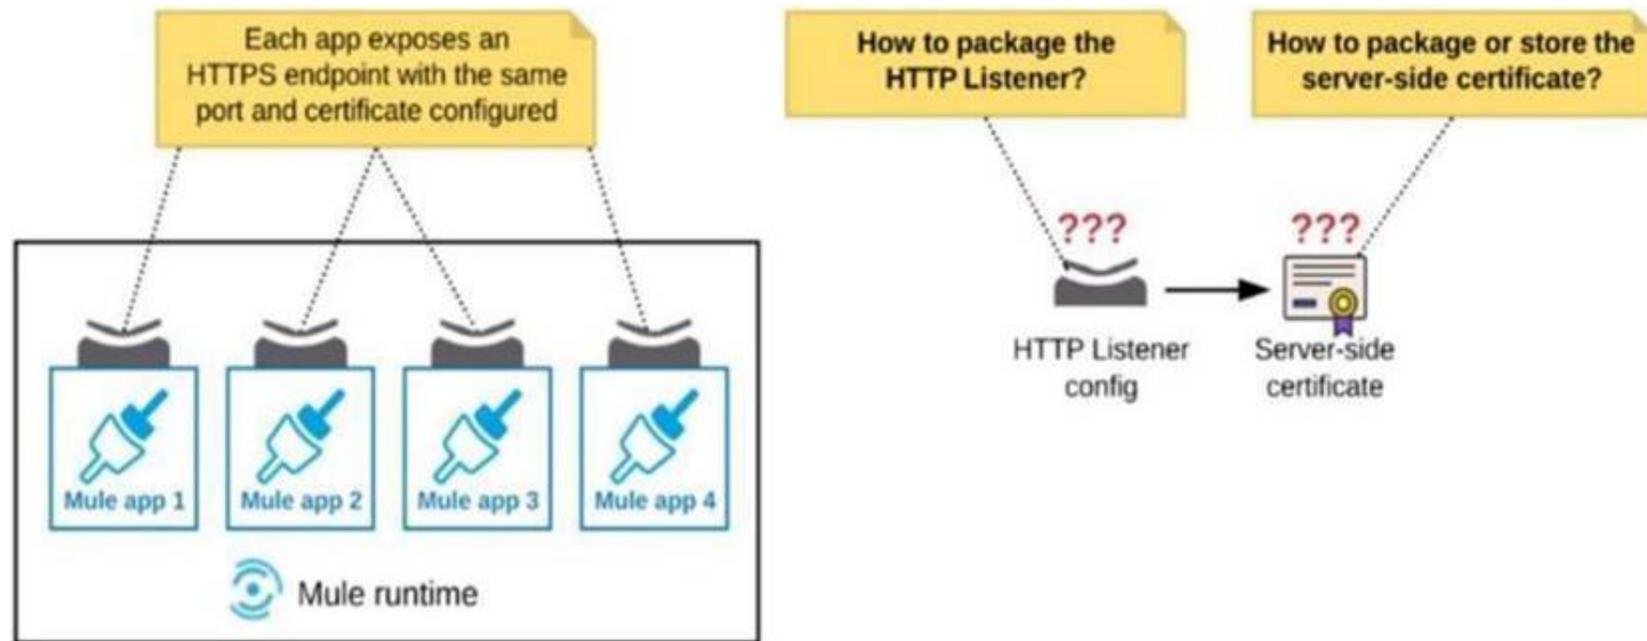
What network architecture supports this requirement?

- 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

Refer to the exhibit.



An organization deploys multiple Mule applications to the same customer -hosted Mule runtime. Many of these Mule applications must expose an HTTPS endpoint on the same port using a serverside certificate that rotates often.

What is the most effective way to package the HTTP Listener and package or store the server-side certificate when deploying these Mule applications, so the disruption caused by certificate rotation is minimized?

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

Answer: B

NEW QUESTION 3

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 4

An XA transaction is being configured that involves a JMS connector listening for Incoming JMS messages. What is the meaning of the timeout attribute of the XA transaction, and what happens after the timeout expires?

- A. The time that is allowed to pass between committing the transaction and the completion of the Mule flow After the timeout, flow processing triggers an error
- B. The time that is allowed to pass between receiving JMS messages on the same JMS connection After the timeout, a new JMS connection is established

- C. The time that is allowed to pass without the transaction being ended explicitly After the timeout, the transaction is forcefully rolled-back
- D. The time that is allowed to pass for state JMS consumer threads to be destroyed After the timeout, a new JMS consumer thread is created

Answer: C

NEW QUESTION 5


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 6

Refer to the exhibit.



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


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

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

Answer: C

NEW QUESTION 7

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 8

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: A

NEW QUESTION 9

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: B

NEW QUESTION 10

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 Mule 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 Mule application using DataWeave, invoking existing transformation services when possible

Answer: D

NEW QUESTION 10

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 13

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 storeThe 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 14


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/partners/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 api.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 application Create 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 application Create 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 17

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

Answer: C

NEW QUESTION 20

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 EACH DB insert in a SEPARATE DB transaction* 3. Acknowledge the JMS message
- B. * 1. Read and acknowledge the JMS message (NOT in an XA transaction)* 2. In a NEW XA transaction, perform BOTH DB inserts
- 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 the JMS message (NOT in an XA transaction)* 2. Perform BOTH DB inserts in ONE DB transaction* 3. Acknowledge the JMS message

Answer: C

NEW QUESTION 24

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 has been updated and hence the database driver library/JAR needs a minor version upgrade
- C. The database server was unavailable for four hours due to a major outage but is now fully operational again
- D. The credentials for accessing the database have been updated and the previous credentials are no longer valid

Answer: D

NEW QUESTION 28

What operation can be performed through a JMX agent enabled in a Mule application?

- A. View object store entries
- B. Replay an unsuccessful message
- C. Deploy a Mule application
- D. Set a particular log4j2 log level to TRACE

Answer: D

NEW QUESTION 32

An integration Mule application is being designed to process orders by submitting them to a backend 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 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: A

NEW QUESTION 33

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

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

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

Answer: B

NEW QUESTION 35

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 theMuleSoft-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 thebackend 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 37

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 40

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

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

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

Answer: B

NEW QUESTION 43

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 45

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 49

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 StepsEach 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
- B. Each Batch Job uses SEVERAL THREADS for the Batch StepsEach 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
- C. Each Batch Job uses a SINGLE THREAD for all Batch StepsEach 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
- D. Each Batch Job uses a SINGLE THREAD to process a configured block size of recordEach Batch Step instance receives A BLOCK OF records as the payload, and BLOCKS of records are processed IN ORDER

Answer: A

NEW QUESTION 51

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 52

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

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

Answer: A

NEW QUESTION 53

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

MCIA-Level-1 Practice Exam Features:

- * MCIA-Level-1 Questions and Answers Updated Frequently
- * MCIA-Level-1 Practice Questions Verified by Expert Senior Certified Staff
- * MCIA-Level-1 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * MCIA-Level-1 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The MCIA-Level-1 Practice Test Here](#)