


2022 MCIA-Level-1 Premium Files Test pdf - Free Dumps Collection [Q64-Q80]



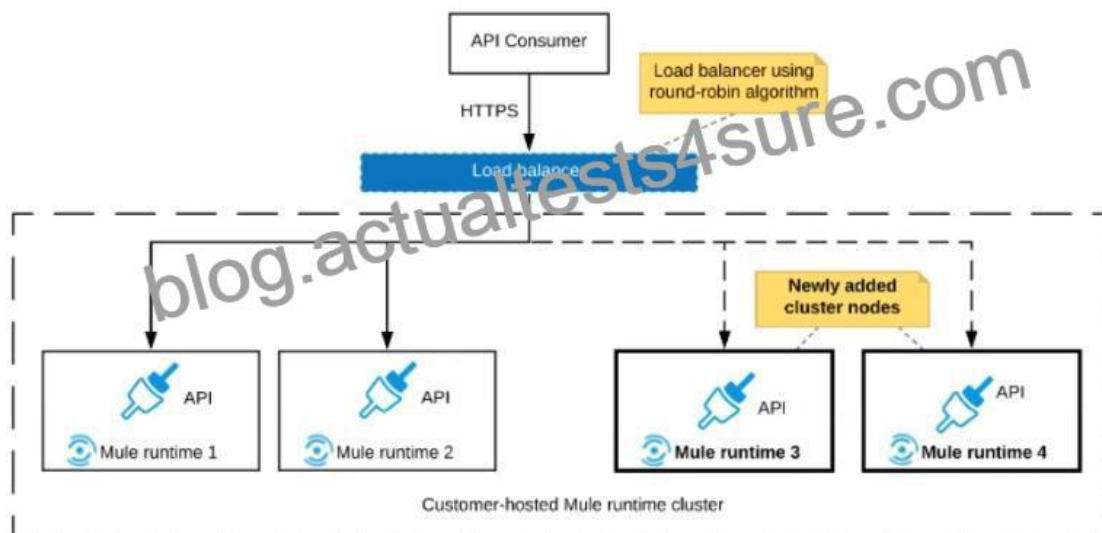
2022 MCIA-Level-1 Premium Files Test pdf - Free Dumps Collection

Get ready to pass the MCIA-Level-1 Exam right now using our MuleSoft Certified Architect Exam Package

MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam Certified Professional salary

The average salary of a MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam Certified Expert in:

- Europe - 70,500 EURO- United State - 100,200 USD- India. - 14,00,327 INR- England - 75,000 POUND **NO.64** 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?

- * 50% reduction in the response time of the API
- * 100% increase in the throughput of the API
- * 50% reduction in the JVM heap memory consumed by each node
- * 50% reduction in the number of requests being received by each node

NO.65 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?

- * 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
- * Application C and P both use the Object Store connector to access the Anypoint Object Store v2
- * 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
- * Application C and P both use the Object Store connector to access a persistent object store

NO.66 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?

- * Implement transformation logic in the new Mule application using DataWeave, replicating the transformation logic of existing transformation services
- * Implement transformation logic in the new Mule application using DataWeave, invoking existing transformation services when possible
- * Extend the existing transformation services with new transformation logic and invoke them from the new Mule application
- * Implement and expose all transformation logic as microservices using DataWeave, so it can be reused by any application component that needs it, including the new Mule application

NO.67 What is a recommended practice when designing an integration Mule 4 application that reads a large XML payload as a stream?

- * 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
- * The payload should be dealt with as an XML stream, without converting it to a single Java object (POJO)
- * The payload size should NOT exceed the maximum available heap memory of the Mule runtime on which the Mule application executes
- * The payload must be cached using a Cache scope If It Is to be sent to multiple backend systems

NO.68 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?

- * 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
- * Store the credentials in properties files in a shared folder within the organization's data center Have the Mule applications load properties files from this shared location at startup
- * Segregate the credentials for each backend system into environment-specific properties files Package these properties files in each Mule application, from where they are loaded at startup
- * 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

NO.69 A mule application is deployed to a Single Cloudhub worker and the public URL appears in Runtime Manager as the APP URL.

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

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

Immediately after the new property values are applied in runtime manager, how is the current Mule application deployment affected

and how will future web client requests to the Mule application be handled?

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

NO.70 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?

- * MuleSoft-hosted (CloudHub) dedicated load balancers
- * Client (application) management
- * Virtual private clouds
- * Permissions

:

To use a dedicated load balancer in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.

NO.71 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?

- * Keep the default logging configuration in RuntimeManager

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

- * Disable Cloudhub logging in Runtime Manager

Define the splunk appender in EACH Mule application's log4j2.xml file

- * 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.

- * Keep the default logging configuration in Runtime Manager

Define the Splunk appender in EACH Mule application log4j2.xml file

NO.72 What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

- * Only a DLB allows the configuration of a custom TLS server certificate
- * Only the SLB can forward HTTP traffic to the VPC-internal ports of the CloudHub workers
- * Both a DLB and the SLB allow the configuration of access control via IP whitelists
- * Both a DLB and the SLB implement load balancing by sending HTTP requests to workers with the lowest workloads
- * Shared load balancers don't allow you to configure custom SSL certificates or proxy rules

- * Dedicated Load Balancer are optional but you need to purchase them additionally if needed.
- * TLS is a cryptographic protocol that provides communications security for your Mule app. TLS offers many different ways of exchanging keys for authentication, encrypting data, and guaranteeing message integrity.
- * The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.
- * Only a DLB allows the configuration of a custom TLS server certificate
- * DLB enables you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.
- * To use a DLB in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.
- * MuleSoft Reference: <https://docs.mulesoft.com/runtime-manager/dedicated-load-balancer-tutorial> Additional Info on SLB Vs DLB:

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

NO.73 An organization's governance process requires project teams to get formal approval from all key stakeholders for all new Integration design specifications. An integration Mule application is being designed that interacts with various backend systems. The Mule application will be created using Anypoint Design Center or Anypoint Studio and will then be deployed to a customer-hosted runtime.

What key elements should be included in the integration design specification when requesting approval for this Mule application?

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

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

NO.74 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 for the service that can be exposed by the CloudHub Shared Load Balancer (SLB) to external web clients over the public internet?

- * Underlying Mule applications need to implement own certificates
- * Only MuleSoft provided certificates can be used for server side certificate
- * Only self signed certificates can be used
- * All certificates which can be used in shared load balancer need to get approved by raising support ticket

Correct answer is Only MuleSoft provided certificates can be used for server side certificate

* The CloudHub Shared Load Balancer terminates TLS connections and uses its own server-side certificate.


* You would need to use dedicated load balancer which can enable you to define SSL configurations to provide custom certificates and optionally enforce two-way SSL client authentication.

* To use a dedicated load balancer in your environment, you must first create an Anypoint VPC. Because you can associate multiple environments with the same Anypoint VPC, you can use the same dedicated load balancer for your different environments.

Additional Info on SLB Vs DLB:

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

NO.75 An organization needs to enable access to their customer data from both a mobile app and a web application, which each need access to common fields as well as certain unique fields. The data is available partially in a database and partially in a 3rd-party CRM system. What APIs should be created to best fit these design requirements?



- * A Process API that contains the data required by both the web and mobile apps, allowing these applications to invoke it directly and access the data they need thereby providing the flexibility to add more fields in the future without needing API changes.
- * One set of APIs (Experience API, Process API, and System API) for the web app, and another set for the mobile app.
- * Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system
- * A common Experience API used by both the web and mobile apps, but separate Process APIs for the web and mobile apps that interact with the database and the CRM System.

Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs

Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web

applications will have different set of requirements which cannot be fulfilled by single experience layer API Option : Common

Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of


flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.

Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led

connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates

duplicity of code.

Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system



Lets analyze the situation in regards to the different options available Option : A common Experience API but separate Process APIs

Analysis : This solution will not work because having common experience layer will not help the purpose as mobile and web

applications will have different set of requirements which cannot be fulfilled by single experience layer API Option : Common

Process API Analysis : This solution will not work because creating a common process API will impose limitations in terms of

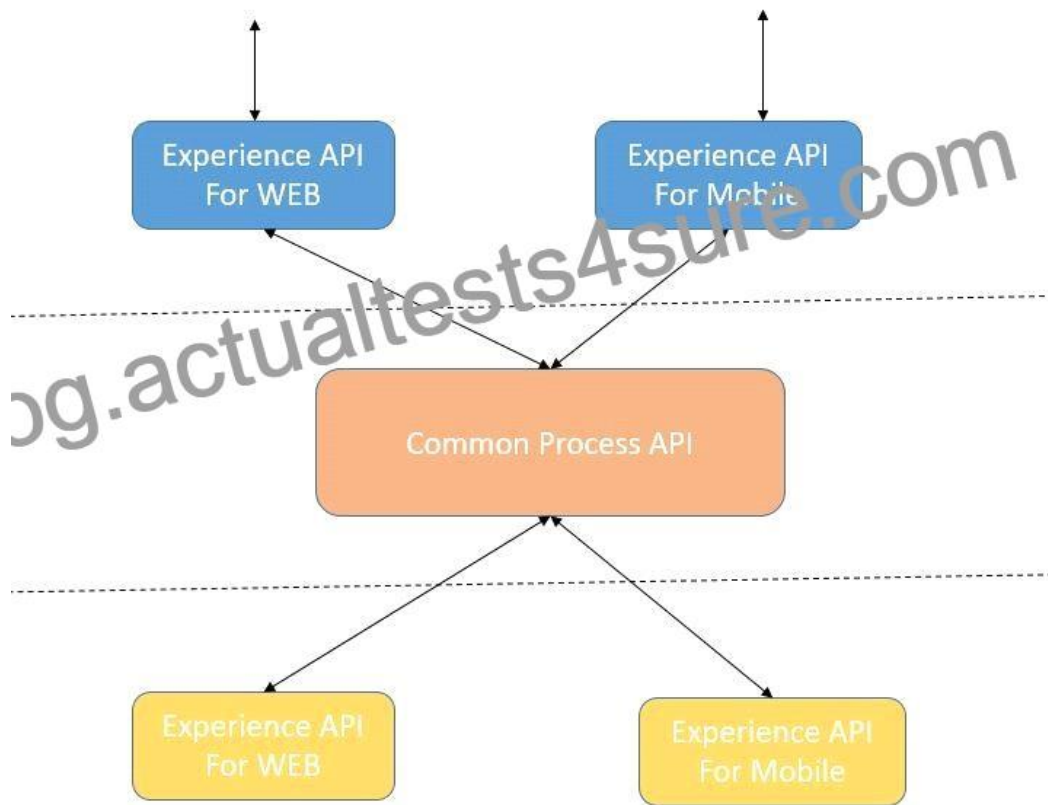
flexibility to customize API;s as per the requirements of different applications. It is not a recommended approach.

Option : Separate set of API's for both the applications Analysis : This goes against the principle of Anypoint API-led

connectivity approach which promotes creating reusable assets. This solution may work but this is not efficient solution and creates

duplicity of code.

Hence the correct answer is: Separate Experience APIs for the mobile and web app, but a common Process API that invokes separate System APIs created for the database and CRM system



NO.76 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?

* 1) Read the JMS message (NOT in an XA transaction)

2) Perform BOTH DB inserts in ONE DB transaction

3) Acknowledge the JMS message

* 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

* 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

* 1) Read and acknowledge the JMS message (NOT in an XA transaction)

2) In a NEW XA transaction, perform BOTH DB inserts

* Option A says "Perform EACH DB insert in a SEPARATE DB transaction". In this case if first DB insert is successful and second one fails then first insert won't be rolled back causing inconsistency. This option is ruled out.

* Option D says Perform BOTH DB inserts in ONE DB transaction.

Rule of thumb is when one or more DB connections are required we must use XA transaction as local transactions support only one resource. So this option is also ruled out.

* Option B acknowledges the before DB processing, so message is removed from the queue. In case of system failure at later point, message can't be retrieved.

* Option C is Valid: Though it says "do not ack JMS message", message will be auto acknowledged at the end of transaction. Here is how we can ensure all components are part of XA transaction:

<https://docs.mulesoft.com/jms-connector/1.7/jms-transactions> Additional Information about transactions:

* XA Transactions; You can use an XA transaction to group together a series of operations from multiple transactional resources, such as JMS, VM or JDBC resources, into a single, very reliable, global transaction.

* The XA (eXtended Architecture) standard is an X/Open group standard which specifies the interface between a global transaction manager and local transactional resource managers.

The XA protocol defines a 2-phase commit protocol which can be used to more reliably coordinate and sequence a series of "all or nothing" operations across multiple servers, even servers of different types

* Use JMS ack if

"Acknowledgment should occur eventually, perhaps asynchronously

"The performance of the message receipt is paramount

"The message processing is idempotent

"For the choreography portion of the SAGA pattern

* Use JMS transactions

"For all other times in the integration you want to perform an atomic unit of work

"When the unit of work comprises more than the receipt of a single message

"To simply and unify the programming model (begin/commit/rollback)

NO.77 To implement predictive maintenance on its machinery equipment, ACME Tractors has installed thousands of IoT sensors that will send data for each machinery asset as sequences of JMS messages, in near real-time, to a JMS queue named SENSOR_DATA on a JMS server. The Mule application contains a JMS Listener operation configured to receive incoming

messages from the JMS servers SENSOR_DATA JMS queue. The Mule application persists each received JMS message, then sends a transformed version of the corresponding Mule event to the machinery equipment back-end systems.

The Mule application will be deployed to a multi-node, customer-hosted Mule runtime cluster. Under normal conditions, each JMS message should be processed exactly once.

How should the JMS Listener be configured to maximize performance and concurrent message processing of the JMS queue?

* Set numberOfConsumers = 1

Set primaryNodeOnly = false

* Set numberOfConsumers = 1

Set primaryNodeOnly = true


* Set numberOfConsumers to a value greater than one

Set primaryNodeOnly = true

* Set numberOfConsumers to a value greater than one

Set primaryNodeOnly = false

NO.78 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?

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

NO.79 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?

- * 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
- * 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
- * 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
- * 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

NO.80 A set of integration Mule applications, some of which expose APIs, are being created to enable a new business process. Various stakeholders may be impacted by this. These stakeholders are a combination of semi- technical users (who understand basic integration terminology and concepts such as JSON and XML) and technically skilled potential consumers of the Mule applications and APIs.

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

- * 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
- * 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
- * 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
- * Use Anypoint Exchange to register the various Mule applications and APIs and share the RAML definitions with the stakeholders, so they can be discovered

What is the duration, language, and format of MuleSoft mcia - Level 1: MuleSoft Certified Integration Architect - Level 1 Exam - Language: English- Passing score: 70%- Type of Questions: Single and Multiple Choice.- Number of Questions: 58- Length of Examination: 120 minutes **Master 2022 Latest The Questions MuleSoft Certified Architect and Pass MCIA-Level-1 Real Exam!:** <https://www.actualtests4sure.com/MCIA-Level-1-test-questions.html>