MuleSoft

Exam Questions MCIA-Level-1

MuleSoft Certified Integration Architect - Level 1

NEW QUESTION 1
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 2
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?

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

Answer: A

NEW QUESTION 3
Refer to the exhibit.

A Mule application is deployed to a cluster of two customer-hosted Mute runtimes. The Mute application has a flow that polls a database and another flow with an HTTP Listener.
HTTP clients send HTTP requests directly to individual cluster nodes.
What happens to database polling and HTTP request handling in the time after the primary (master) node of the cluster has railed, but before that node is restarted?

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

Answer: A

NEW QUESTION 4
Refer to the exhibit.

An organization uses a 2-node Mute 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
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 6
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 7
What comparison is true about a CloudHub Dedicated Load Balancer (DLB) vs. the CloudHub Shared Load Balancer (SLB)?

A. Both a DLB and the SLB implement load balancing by sending HTTP requests to workers with the lowest workloads
B. Both a DLB and the SLB allow the configuration of access control via IP whitelists
C. Only a DLB allows the configuration of a custom TLS server certificate
D. Only the SLB can forward HTTP traffic to the VPC-internal ports of the CloudHub workers

Answer: B

NEW QUESTION 8
What aspect of logging is only possible for Mule applications deployed to customer-hosted Mule runtimes, but NOT for Mule applications deployed to CloudHub?

A. To send Mule application log entries to Splunk
B. To change tog4j2 tog levels in Anypoint Runtime Manager without having to restart the Mule application
C. To log certain messages to a custom log category
D. To directly referenceone shared and customized log4j2.xml file from multiple Mule applications

Answer: D

NEW QUESTION 9
An integration Mule application is deployed to a customer-hosted multi-node Mule 4 runtime duster. The Mule application uses a Listener operation of a JMS connector to receive incoming messages from a JMS queue.
How are the messages consumed by the Mule application?

A. Depending on the JMS provider's configuration, either all messages are consumed by ONLY the primary cluster node or else ALL messages are consumed by ALL cluster nodes
B. Regardless of the Listener operation configuration, all messages are consumed by ALL cluster nodes
C. Depending on the Listener operation configuration, either all messages are consumed by ONLY the primary cluster node or else EACH message is consumed by ANY ONE cluster node
D. Regardless of the Listener operation configuration, all messages are consumed by ONLY the primary cluster node

Answer: C

NEW QUESTION 10
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 10
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 15
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 servic
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 17
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 18
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 C and P both use the Object Store connector to access the Anypoint Object Store v2
B. Application C and P both use the Object Store connector to access a persistent object store
C. Application C uses the Object Store connector to access a persistent objectApplication P accesses the persistent object store via the Object Store REST API
D. Application P uses the Object Store connector to access a persistent object storeApplication C accesses this persistent object store via the Object Store REST API through an IPsec tunnel

Answer: A

NEW QUESTION 20
Refer to the exhibit. A Mule application is being designed to expose a SOAP web service to its clients.
What language is typically used inside the web service's interface definition to define the data structures that the web service is expected to exchange with its clients?

A. JSON Schema
B. RAML
C. WSDL
D. XSD

Answer: C

NEW QUESTION 21
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry

* MCIA-Level-1 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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