

Exam : Data Cloud Consultant

**Title : Salesforce Certified Data
Cloud Consultant**

Version : DEMO

1.A client wants to bring in loyalty data from a custom object in Salesforce CRM that contains a point balance for accrued hotel points and airline points within the same record. The client wants to split these point systems into two separate records for better tracking and processing.

What should a Data 360 Consultant recommend in this scenario?

- A. Create a junction object in Salesforce CRM and modify the ingestion strategy.
- B. Use batch transforms to create a second data lake object.
- C. Clone the data source object.
- D. Create a data kit from the data lake object and deploy it to the same Data 360 org.

Answer: B

Explanation:

The design point is to preserve source fidelity while shaping data only where Data 360 processing needs it. Here, Use batch transforms to create a second data lake object. fits because it changes the shape, keying, or refresh behavior at the Data 360 layer instead of forcing the source system to carry an analytics-specific design. In production, this keeps the upstream application simpler and gives the data team a repeatable way to prepare records for mapping, identity resolution, insights, or segmentation. The distractors fall short because they either move the problem into the wrong system, add needless duplication, ignore Data 360 object relationships, or rely on a feature built for a different lifecycle stage. In a real implementation, those choices usually create brittle pipelines, stale data, security exposure, or segments that look correct on paper but fail when activated. Thinking like an architect, the selected option places the logic where Data 360 can govern it and reuse it reliably.

2.A data architect has deployed an inbound change set containing a data kit into a new Salesforce Production org. Although the deployment status shows "Succeeded", the architect notices that the specific data stream configurations and calculated insights included in the kit are not yet active or visible in the Data 360 functional tabs.

What is the next required action the architect must perform to fully implement the configurations within the data kit?

- A. Install the data kit from the Data 360 Setup menu to map and activate the components.
- B. Re-upload the outbound change set from the source org to clear the metadata cache.
- C. Run a "Full Refresh" on all data streams within the newly deployed change set.
- D. Use the Salesforce Package Manager to verify the installation status of the components.

Answer: A

Explanation:

The lifecycle point is that Data 360 configuration is metadata, but it still has product-specific packaging and activation steps. Install the data kit from the Data 360 Setup menu to map and activate the components. is the right fit because Data 360 components are not fully usable just because generic deployment succeeded. The consultant must complete the Data 360-specific lifecycle step so streams, mappings, insights, and rules are correctly registered in the target environment. The distractors fall short because they either move the problem into the wrong system, add needless duplication, ignore Data 360 object relationships, or rely on a feature built for a different lifecycle stage. In a real implementation, those choices usually create brittle pipelines, stale data, security exposure, or segments that look correct on paper but fail when activated. Thinking like an architect, the selected option places the logic where Data 360 can govern it and reuse it reliably.

3.A Data 360 Consultant needs to provide real- time, multidimensional analytics for a derived insight, "Net Profit", to be used within Tableau Next. The requirements specify that the metrics must be as up to date as possible while allowing users to dynamically filter by region, period, and product category. Which approach should the consultant take?

- A. Define a real-time data graph.
- B. Create a calculated insight using SQL.
- C. Create a calculated field in the semantic layer.
- D. Configure a streaming transform.

Answer: C

Explanation:

The analytics requirement determines whether the logic belongs in an insight, a semantic metric, or a query-time layer. Create a calculated field in the semantic layer. matches the need because the business is asking for reusable metrics, dimensions, or aggregation behavior that consumers can filter and analyze. Data 360 separates raw harmonized data from analytical definitions so that dashboards, Tableau experiences, and segment logic remain consistent. The distractors fall short because they either move the problem into the wrong system, add needless duplication, ignore Data 360 object relationships, or rely on a feature built for a different lifecycle stage. In a real implementation, those choices usually create brittle pipelines, stale data, security exposure, or segments that look correct on paper but fail when activated. Thinking like an architect, the selected option places the logic where Data 360 can govern it and reuse it reliably. This is the nuance exam questions often test: the platform capability must match both the technical layer and the business timing requirement, not just sound related to data.

4.A Data 360 Consultant is preparing to implement Data 360.

Which ethic should the consultant adhere to regarding customer data?

- A. Carefully consider asking for sensitive data such as age, gender, and ethnicity.
- B. Give senior leaders in the firm access to customer data for audit purposes.
- C. Collect and use all of the data to create more personalized experiences.
- D. Map sensitive data to the same data model object for ease of deletion.

Answer: A

Explanation:

The governance lens is least privilege, purpose limitation, and honoring privacy operations against the individual profile. Carefully consider asking for sensitive data such as age, gender, and ethnicity. supports the governance requirement because Data 360 implementations should minimize unnecessary access, avoid over-collection, and process deletion or consent requests at the profile level where Salesforce expects them. The safest design is explicit, auditable, and limited to the business purpose. The distractors fall short because they either move the problem into the wrong system, add needless duplication, ignore Data 360 object relationships, or rely on a feature built for a different lifecycle stage. In a real implementation, those choices usually create brittle pipelines, stale data, security exposure, or segments that look correct on paper but fail when activated. Thinking like an architect, the selected option places the logic where Data 360 can govern it and reuse it reliably. This is the nuance exam questions often test: the platform capability must match both the technical layer and the business timing requirement, not just sound related to data.

5.Northern Trail Outfitters uploads new customer data to an Amazon S3 bucket on a daily basis to be

ingested in Data 360. Based on this, a calculated insight is created that shows the total spend per customer in the last 30 days.

In which sequence should each process be run to ensure that freshly imported data is ready and available to use for any segment?

- A. Calculated Insight > Refresh Data Stream > Identity Resolution
- B. Refresh Data Stream > Calculated Insight > Identity Resolution
- C. Refresh Data Stream > Identity Resolution > Calculated Insight
- D. Identity Resolution > Refresh Data Stream > Calculated Insight

Answer: C

Explanation:

The design point is to preserve source fidelity while shaping data only where Data 360 processing needs it. Here, Refresh Data Stream > Identity Resolution > Calculated Insight fits because it changes the shape, keying, or refresh behavior at the Data 360 layer instead of forcing the source system to carry an analytics-specific design. In production, this keeps the upstream application simpler and gives the data team a repeatable way to prepare records for mapping, identity resolution, insights, or segmentation. The distractors fall short because they either move the problem into the wrong system, add needless duplication, ignore Data 360 object relationships, or rely on a feature built for a different lifecycle stage. In a real implementation, those choices usually create brittle pipelines, stale data, security exposure, or segments that look correct on paper but fail when activated. Thinking like an architect, the selected option places the logic where Data 360 can govern it and reuse it reliably.