

DUMPSBOSS.

Salesforce Data Cloud Accredited Professional

Salesforce Salesforce-Data-Cloud

Version Demo

Total Demo Questions: 22

Total Premium Questions: 229

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Cloud Overview	16
Topic 2, Data Ingestion and Modeling	66
Topic 3, Identity Resolution	29
Topic 4, Segmentation and Activation	81
Topic 5, Data Cloud Administration	26
Topic 6, Data Cloud Ecosystem	11
Total	229

QUESTION NO: 1

A segment fails to refresh with the error "Segment references too many data lake objects (DLOS)".

Which two troubleshooting tips should help remedy this issue?

Choose 2 answers

- A. Split the segment into smaller segments.
- B. Use calculated insights in order to reduce the complexity of the segmentation query.
- C. Refine segmentation criteria to limit up to five custom data model objects (DMOs).
- D. Space out the segment schedules to reduce DLO load.

ANSWER: A B

Explanation:

"Split the segment into smaller segments." is an appropriate remediation because this error occurs when the segment's query reaches the platform limit for the number of Data Lake Objects it can reference. Breaking a complex audience definition into smaller, more focused segments reduces the object relationships and filtering logic evaluated in an individual segment query. The smaller results can then be used in a design that avoids concentrating all the required joins in one segment definition.

"Use calculated insights in order to reduce the complexity of the segmentation query." is also appropriate. A calculated insight preaggregates or derives measures and attributes from underlying Data Cloud data. Using the resulting insight in a segment can replace a set of detailed conditions and object-level joins with a simpler condition against an already computed value, such as a purchase total, engagement measure, or customer metric. This reduces the number of Data Lake Objects that the segment query needs to directly reference and makes the segment more maintainable. Salesforce's segment troubleshooting guidance identifies simplifying the segment design as the relevant approach for object-reference-limit errors, while calculated insights are designed to create reusable derived metrics for downstream uses such as segmentation.

See Salesforce's [Segment Troubleshooting](#) and [Calculated Insights](#) documentation.

QUESTION NO: 2

Northern Trail Outfitters unifies individuals in its Data Cloud instance.

Which three features can a consultant use to validate the data on a unified profile?

Choose 3 answers

- A. Identity Resolution
- B. Query APL
- C. Data Explorer
- D. Profile Explorer
- E. Data Actions

ANSWER: A C D

Explanation:

Identity Resolution, **Data Explorer**, and **Profile Explorer** provide complementary ways to validate unified-profile data in Data Cloud. Identity Resolution is used to review the ruleset and matching/reconciliation logic that determines whether

source records represent the same individual. This validation is essential because the identity rules directly control the composition of a unified individual and the attributes available on that profile.

Data Explorer enables a consultant to inspect the data stored in Data Cloud objects, including unified individual data and related records. It supports record-level review and filtering, helping validate that expected attributes and values are present after ingestion, harmonization, and identity resolution.

Profile Explorer provides the unified-profile-focused view needed to inspect an individual in context. A consultant can review profile details and associated information to confirm that the resulting unified profile accurately represents the linked source data. Together, these capabilities validate the identity logic, the stored unified data, and the profile-level result. Salesforce documents identity resolution and profile exploration in its [Identity Resolution documentation](#) and [Profile Explorer documentation](#).

QUESTION NO: 3

Which data stream category type should be assigned in order to use the dataset for date and time-based operations in segmentation and calculated insights?

- A. Individual
- B. Engagement
- C. Sales Order
- D. Profile

ANSWER: B

Explanation:

Engagement is the appropriate data stream category for a dataset that will be used in date- and time-based segmentation and calculated insights. Engagement data represents time-stamped customer activities or events, such as an email interaction, web visit, mobile event, or other behavioral occurrence. Assigning this category identifies the dataset as event-oriented data, allowing Data Cloud to use the event date and time context when evaluating audiences and when aggregating measures for calculated insights.

This classification supports common analytical needs such as identifying people who performed an event within a rolling period, measuring the most recent interaction, counting events during a specified interval, or calculating trends over time. To obtain reliable results, the ingested engagement records should map the relevant event timestamp to the appropriate Data Cloud data model field and be related to the applicable individual or unified individual. The category therefore provides the intended semantic treatment for behavioral records used by segmentation and insight calculations. Salesforce describes Data Cloud's data-model concepts and data categorization in its [Data Cloud Data Modeling](#) learning module and provides implementation guidance in the [Data Streams documentation](#).

QUESTION NO: 4

Which three objects are created as a result of Identity Resolution?

- A. Unified Subscriber
- B. Unified Data Model
- C. Unified Contact Point
- D. Unified Link
- E. Unified Individual

ANSWER: C D E

Explanation:

Identity Resolution in Data Cloud reconciles records that represent the same person across source systems and produces a unified profile model for activation, analysis, and personalization. **Unified Individual** is the primary resolved profile record. It represents an individual after Data Cloud applies the configured matching and reconciliation rules to source records, bringing appropriate profile attributes together into a single unified view.

Unified Contact Point represents the person's resolved contact information, such as email addresses, phone numbers, addresses, and other contact-point data. Maintaining these contact points separately supports consistent use of channels and identifiers associated with the unified person.

Unified Link provides the lineage between source records and the resolved records created by Identity Resolution. It connects individual and contact-point source data to their corresponding unified entities, enabling Data Cloud users to understand which records contributed to a unified profile and supporting traceability of the resolution outcome. Together, these objects provide the unified profile, its associated contact information, and the record relationships needed to establish provenance. Salesforce describes Identity Resolution and the resulting unified data in its [Identity Resolution documentation](#) and [Unified Individual data model reference](#).

QUESTION NO: 5

Northern Trail Outfitters wants to implement Data Cloud and has several use cases in mind.

Which two use cases are considered a good fit for Data Cloud?

Choose 2 answers

- A. To ingest and unify data from various sources to reconcile customer identity
- B. To create and orchestrate cross-channel marketing messages
- C. To use harmonized data to more accurately understand the customer and business impact
- D. To eliminate the need for separate business intelligence and IT data management tools

ANSWER: A C

Explanation:

Data Cloud is well suited to bringing together customer data that is distributed across Salesforce applications, external systems, and both batch and streaming sources. "To ingest and unify data from various sources to reconcile customer identity" is therefore a core use case. Data Cloud can ingest data, map it to its data model, and apply identity resolution rules to link records that represent the same individual or account. This produces unified profiles that make customer information more complete, consistent, and usable across connected Salesforce experiences.

"To use harmonized data to more accurately understand the customer and business impact" is also a strong fit. After data is mapped and harmonized, organizations can query it, build calculated insights, analyze behaviors and outcomes, and create segments based on a more complete view of customer interactions. These capabilities help teams understand engagement, value, and operational or business results using connected data rather than isolated records. Data Cloud is designed to make these trusted, unified data assets available for analysis and activation across Salesforce.

Salesforce describes Data Cloud's ingestion, harmonization, identity resolution, unified-profile, segmentation, and calculated-insight capabilities in its [Data Cloud documentation](#) and its [Data Cloud product overview](#).

QUESTION NO: 6

Cumulus Financial wants its service agents to view a display of all cases associated with a Unified Individual on a contact record.

Which two features should a consultant consider for this use case?

Choose 2 answers

- A. Data Action

- B. Profile API
- C. Lightning Web Components
- D. Query APL

ANSWER: B C

Explanation:

Profile API is appropriate because it provides programmatic access to Data Cloud unified-profile data. A consultant can use the Unified Individual identity and related profile information to retrieve the data needed to identify records associated with that individual across contributing sources. This makes it suitable for supplying a service-oriented view that is based on the Data Cloud unified profile rather than being limited to a single source contact record. The implementation should ensure that the case data is modeled and related appropriately in Data Cloud, and that access is governed by the organization's data-access and consent requirements.

Lightning Web Components is appropriate because it provides the Salesforce user-interface layer for presenting this information directly on a Contact record page. A custom component can call supported Apex or API-based integration logic, render the returned associated cases in a table or list, and provide a focused experience for service agents without requiring them to navigate away from the record. Lightning Web Components are designed to be reusable components in Lightning Experience and can be placed on record pages through the Lightning App Builder. Together, Profile API access and a Lightning Web Component support both retrieval of unified-profile context and display of the associated case information where agents work. See the [Data Cloud Profile API overview](#) and the [Lightning Web Components documentation](#).

QUESTION NO: 7

The recruiting team at Cumulus Financial wants to identify which candidates have browsed the jobs page on its website at least twice within the last 24 hours. They want the information about these

candidates to be available for segmentation in Data Cloud and the candidates added to their recruiting system.

Which feature should a consultant recommend to achieve this goal?

- A. Streaming data transform
- B. Streaming insight
- C. Calculated insight
- D. Batch data transform

ANSWER: B

Explanation:

Streaming insight is the appropriate feature because the requirement is based on detecting an event pattern as website activity arrives: a candidate viewing a particular page at least twice during a rolling 24-hour period. Streaming Insights evaluates streaming event data in near real time and can produce an insight when the defined aggregate or threshold condition is met. This lets Cumulus Financial identify candidates promptly rather than waiting for a scheduled batch calculation.

The resulting streaming-insight output can be used in Data Cloud to support timely audience segmentation. It can also support downstream activation so that qualifying candidates are sent to the recruiting system, allowing recruiters to act while the candidate's interest is current. The rolling time window and event-count condition are central to this use case: the insight continuously evaluates behavioral events and recognizes when an individual reaches the specified threshold.

Salesforce describes Streaming Insights as a capability for deriving real-time insights from streaming data and using those insights in Data Cloud workflows. For implementation concepts and use cases, see Salesforce Trailhead's [Get Started with Streaming Insights](#) unit and the [Salesforce Help documentation for Streaming Insights](#).

QUESTION NO: 8

How should a Data Cloud consultant successfully apply consent during segmentation?

- A. Include the Consent Status from the golden record during activation for any applicable channels of engagement.
- B. Include Party Identification for any applicable channels of engagement in the filter criteria for each segment.
- C. Include the Unified Profile during segmentation for any applicable channels of engagement.
- D. Include the Consent Status for any applicable channels of engagement in the filter criteria for each segment.

ANSWER: D

Explanation:

Include the Consent Status for any applicable channels of engagement in the filter criteria for each segment. In Data Cloud, consent must be evaluated as part of the segment's audience-selection logic so that only individuals with the required permission for the intended channel and purpose qualify for activation. A consultant should use consent data modeled in Data Cloud—such as records associated with contact points and their relevant communication channels—and apply filters that represent the organization's approved consent requirements. For example, a segment intended for email outreach should include criteria that identify profiles with an appropriate email consent status, while an SMS audience requires the applicable mobile/SMS consent status. Applying these conditions directly in the segment makes consent enforcement repeatable, auditable, and aligned to the specific engagement channel before the audience is sent to a destination. Consent requirements can differ by channel, brand, purpose, or jurisdiction, so the filter criteria should reflect the policy governing that activation use case. Salesforce documents consent management as a way to capture and manage customers' communication preferences, and Data Cloud segmentation supports filtering audiences using available profile and related data. See [Salesforce Consent Management](#) and [Create Segments in Data Cloud](#).

QUESTION NO: 9

A consultant wants to confirm the Identity resolution they Just set up. Which two features can the consultant use to validate the data on a unified profile?

Choose 2 answers

- A. Identity Resolution
- B. Data Actions
- C. Data Explorer
- D. Query API

ANSWER: C D

Explanation:

Data Explorer and **Query API** provide complementary ways to validate identity-resolution results in Data Cloud. Data Explorer is the Data Cloud interface for inspecting records stored in data model objects. A consultant can examine unified-profile data and related record-level attributes to verify that the expected source data has been brought together and that the resulting values align with the configured identity-resolution rules. This makes it useful for direct, interactive spot checks after a ruleset runs.

The Query API provides programmatic access to Data Cloud data through SQL queries. A consultant can query unified profile objects and relevant related objects, filter for known records or identifiers, and compare returned results with expected match and reconciliation outcomes. This is especially helpful when validation needs to cover many records, be repeatable, or be incorporated into a technical test process. Together, Data Explorer supports hands-on record inspection, while Query API supports targeted and scalable validation of the unified data produced by identity resolution. See Salesforce's [Data Cloud Query API overview](#) and [Data Explorer documentation](#).

QUESTION NO: 10

How does Data Cloud ensure high availability and fault tolerance for customer data?

- A. By distributing data across multiple regions and data centers
- B. By using a data center with robust backups
- C. By Implementing automatic data recovery procedures
- D. By limiting data access to essential personnel

ANSWER: A

Explanation:

By distributing data across multiple regions and data centers is correct because resilient cloud services are designed so that the loss of an individual infrastructure component, availability zone, or data-center location does not make the service unavailable. Data Cloud runs on Salesforce's cloud infrastructure, whose architecture is designed for redundancy and continuity. Distributing service components and data-storage capabilities across independent locations reduces single points of failure and allows traffic and processing to continue through healthy infrastructure when an underlying component has an outage. This supports high availability, meaning that Data Cloud remains accessible and operational, as well as fault tolerance, meaning that it can withstand infrastructure failures without losing its ability to serve customer-data workloads. Salesforce's Hyperforce infrastructure is built on public-cloud regions and availability zones, which provide the physical separation and redundant infrastructure underpinning this resilience model. Salesforce also publishes real-time service availability and incident information through Salesforce Trust. See the [Salesforce Hyperforce overview](#) and [Salesforce Trust status site](#) for details about Salesforce infrastructure and service reliability.

QUESTION NO: 11

A consultant has an activation that is set to publish every 12 hours, but has discovered that updates to the data prior to activation are delayed by up to 24 hours.

Which two areas should a consultant review to troubleshoot this issue?

Choose 2 answers

- A. Review data transformations to ensure they're run after calculated insights.
- B. Review calculated insights to make sure they're run before segments are refreshed.
- C. Review segments to ensure they're refreshed after the data is ingested.
- D. Review calculated insights to make sure they're run after the segments are refreshed.

ANSWER: B C

Explanation:

Review calculated insights to make sure they're run before segments are refreshed, and review segments to ensure they're refreshed after the data is ingested. An activation can publish only the audience membership available when its scheduled run begins. That membership is produced through an upstream processing chain: source data must first be ingested and harmonized into the data model; any calculated insight used as a segment attribute must then be recalculated; and the segment must be refreshed to evaluate its criteria against the latest available profile and insight data. If either the calculated-insight schedule or segment-refresh schedule runs before its upstream dependency has completed, the activation can publish an older audience even though the activation itself runs every 12 hours. Reviewing these schedules and their completion history identifies whether the apparent 24-hour delay results from a dependency running in the next processing cycle. Salesforce documents that calculated insights derive metrics from Data Cloud data model objects and can be used for segmentation, while segments are refreshed to recalculate membership before activation. The relevant processing order is therefore ingestion, calculated insights where applicable, segment refresh, and activation. See Salesforce's [Calculated Insights reference](#) and [Data Cloud Segments documentation](#).

QUESTION NO: 12

A customer has unified records from Salesforce CRM and an ecommerce platform. Both systems provide an email address, but the CRM email has been verified by the customer and must always be displayed on the Unified Individual when values differ.

Which identity resolution configuration should the consultant recommend?

- A. Create a source priority reconciliation rule that ranks Salesforce CRM first.
- B. Create an exact matching rule on the Salesforce CRM email address.
- C. Use a most occurring reconciliation rule for the email attribute.
- D. Use a fuzzy matching rule on the ecommerce email address.

ANSWER: A

Explanation:

A **source priority reconciliation rule that ranks Salesforce CRM first** is the appropriate configuration. Reconciliation rules determine which attribute value is retained on a unified profile after records have been matched. Ranking the verified CRM source first ensures that its email value is selected when multiple matched records contain different emails.

A matching rule determines whether records represent the same person; it does not determine which email value is displayed after unification. A most occurring or most recent reconciliation strategy could select an ecommerce value and would not satisfy the requirement that the verified CRM value always take precedence.

QUESTION NO: 13

What is the primary purpose of Data Cloud?

- A. Providing a golden record of a customer
- B. Managing sales cycles and opportunities
- C. Analyzing marketing data results
- D. Integrating and unifying customer data

ANSWER: D

Explanation:

Integrating and unifying customer data is the primary purpose of Salesforce Data Cloud. Data Cloud connects data from Salesforce applications, external systems, and other sources, then harmonizes that information into a common data model. This enables organizations to resolve identities across disparate records and build unified customer profiles that can be used throughout Salesforce for more relevant, real-time engagement, analytics, and automation.

The platform is designed to make customer data usable across teams and channels rather than leaving it isolated in separate systems. Its capabilities include data ingestion, mapping data to the Customer 360 Data Model, identity resolution, calculated insights, segmentation, and activation to Salesforce and external destinations. A unified profile is therefore a foundational outcome of Data Cloud: it provides the connected customer-data basis needed for downstream business processes and personalized experiences. Salesforce describes Data Cloud as a platform that unifies customer data and makes it actionable across the Salesforce ecosystem. See the [Salesforce Data Cloud overview](#) and the [Salesforce Data Cloud documentation](#).

QUESTION NO: 14

How many calculated insights can be created per tenant?

- A. 5
- B. 10
- C. 50
- D. 100

ANSWER: D

Explanation:

100 is correct. A Data Cloud tenant can have up to 100 calculated insights. Calculated insights are reusable, derived metrics that aggregate and measure data stored in Data Cloud, such as customer lifetime value, total purchases, average order value, or engagement counts. They can include dimensions, measures, filters, and time-based calculations, allowing an organization to create business-ready metrics from harmonized data without requiring each downstream consumer to rebuild the same logic.

This tenant-level limit matters during solution design because calculated insights are shared assets that can support segmentation, activation, analytics, and other Data Cloud use cases. Teams should therefore define calculated insights at an appropriate level of reuse and govern their creation so that the available capacity is reserved for meaningful, broadly useful measures. Salesforce documents calculated insights and their applicable platform limits in its Data Cloud documentation. For implementation details, see [Salesforce Help: Calculated Insights](#) and the [Salesforce Trailhead module on creating calculated insights](#).

QUESTION NO: 15

The marketing manager at Cloud Kicks plans to bring in corporate phone numbers for its accounts into Data Cloud. They plan to use a custom field with data set to Phone to store these phone numbers.

Which statement is true when ingesting phone numbers?

- A. Text value can be accepted for ingestion into a Phone data type field.
- B. Data Cloud validates the format of the phone number at the time of Ingestion.
- C. The phone number field can only accept 10-digit values.
- D. The phone number field should be used as a primary key.

ANSWER: A

Explanation:

Text value can be accepted for ingestion into a Phone data type field. Phone numbers should be treated as character-based values rather than numeric measures because they can contain a leading plus sign, leading zeroes, spaces, parentheses, hyphens, and country or area codes. Treating a phone number as text preserves those meaningful characters and avoids numeric transformations such as stripping leading zeroes or interpreting the value as a number for calculation.

In Data Cloud, the Phone field type represents the intended semantic use of the attribute while allowing phone-number values supplied by a source to be ingested as text. This is especially important for corporate account phone numbers, which can use international dialing conventions and extensions rather than a single fixed-length format. During data-model mapping, Cloud Kicks can map the source phone column to the custom Phone field so the attribute is consistently represented in the Data Cloud data model and available for downstream activation or segmentation use cases.

For Data Cloud data-model field types and their intended use, see the [Salesforce Data Cloud data types reference](#) and Salesforce's [Data Cloud data modeling Trailhead module](#).

QUESTION NO: 16

To import campaign members into a campaign in Salesforce CRM, a user wants to export

the segment to Amazon S3. The resulting file needs to include the Salesforce CRM Campaign ID in the name.

What are two ways to achieve this outcome?

Choose 2 answers

- A. Include campaign identifier in the activation name.
- B. Hard code the campaign identifier as a new attribute in the campaign activation.
- C. Include campaign identifier in the filename specification.
- D. Include campaign identifier in the segment name.

ANSWER: A C

Explanation:

Including the campaign identifier in the activation name is a valid approach because the activation name can be used as part of the exported file's naming convention for an Amazon S3 activation. Setting the activation name to a value that contains the Salesforce CRM Campaign ID provides a clear, static identifier that helps associate the delivered file with the intended CRM campaign. This is particularly useful when each activation is created for a specific campaign.

Including the campaign identifier in the filename specification is also correct because the Amazon S3 activation configuration supports defining how output files are named. A filename specification can include a literal Campaign ID, enabling the delivered file name to carry the identifier required for the downstream campaign-member import process. This approach directly controls the S3 output file name and can be combined with other naming elements, such as dates or segment-related values, to make delivered files easier to identify and manage.

Both methods are configured as part of the activation setup and ensure that the Campaign ID is present in the exported file name before the file is retrieved from Amazon S3 for import into Salesforce CRM. See Salesforce's [Data Cloud activation documentation](#) and the [Data Cloud activation Trailhead module](#).

QUESTION NO: 17

A customer notices that their consolidation rate has recently increased. They contact the consultant to ask why.

What are two likely explanations for the increase?

Choose 2 answers

- A. New data sources have been added to Data Cloud that largely overlap with the existing profiles.
- B. Duplicates have been removed from source system data streams.
- C. Identity resolution rules have been removed to reduce the number of matched profiles.
- D. Identity resolution rules have been added to the ruleset to increase the number of matched profiles.

ANSWER: A D

Explanation:

In Data Cloud Identity Resolution, the consolidation rate indicates the degree to which source profiles are combined into unified individual profiles. It is calculated as one minus the number of unified profiles divided by the number of source profiles. Consequently, an increase means that, relative to the total population of source profiles, more records are being recognized as belonging to the same person and consolidated into fewer unified profiles.

“New data sources have been added to Data Cloud that largely overlap with the existing profiles.” is a likely cause because the added source records increase the source-profile population while identity resolution identifies records shared with the existing sources and combines them into the applicable unified profiles. This increases the proportion of consolidated source records.

“Identity resolution rules have been added to the ruleset to increase the number of matched profiles.” is also a likely cause. Adding appropriate match rules or match criteria can allow Identity Resolution to identify additional records as the same individual. More successful matches result in more source profiles rolling up to unified profiles, thereby increasing the consolidation rate after the ruleset runs.

Salesforce documents consolidation rate as an Identity Resolution metric and explains how rulesets use match and reconciliation rules to build unified profiles. See [Salesforce Identity Resolution documentation](#) and [Configure Identity Resolution Rulesets](#).

QUESTION NO: 18

What are the two minimum requirements needed when using the Visual Insights Builder to create a calculated insight?

Choose 2 answers

- A. At least one measure
- B. At least one dimension
- C. At least two objects to Join
- D. A WHERE clause

ANSWER: A B

Explanation:

At least one measure and **At least one dimension** are the minimum components for a calculated insight built with Visual Insights Builder. A measure defines the quantitative result that Data Cloud calculates, such as a count of records, a sum of transaction amounts, or an average value. This provides the actual metric that the calculated insight will store and make available for activation, segmentation, and analysis.

A dimension supplies the business context used to organize that metric, such as an individual, product, location, or time-related attribute. Dimensions establish the grouping level for the result, allowing Data Cloud to produce meaningful aggregated insight records rather than an uncontextualized calculation. Together, the measure and dimension define both what is being calculated and the entity or attribute by which it is analyzed. Visual Insights Builder is designed to let users construct these aggregate calculations through a guided interface using Data Model Objects and their relationships. Salesforce describes calculated insights as metrics based on dimensions and measures, and documents Visual Insights Builder as the no-code experience for creating them. See [Salesforce Help: Calculated Insights](#) and [Salesforce Help: Create a Calculated Insight](#).

QUESTION NO: 19

A consultant wants to ensure that every segment managed by multiple brand teams adheres to the same set of exclusion criteria, that are updated on a monthly basis.

What is the most efficient option to allow for this capability?

- A. Create, publish, and deploy a data kit.
- B. Create a reusable container block with common criteria.
- C. Create a nested segment.
- D. Create a segment and copy it for each brand.

ANSWER: B

Explanation:

Create a reusable container block with common criteria is the most efficient approach because a container block lets teams define a common set of segment logic once and reuse it in multiple segments. The shared exclusion criteria can be placed in the block and included wherever each brand team needs it. When the criteria are revised during the monthly maintenance cycle, the consultant updates the reusable block rather than manually changing each individual segment. This supports consistent audience governance across brands, reduces the risk of different teams applying slightly different suppression logic, and substantially lowers recurring administrative effort. Container blocks can organize and encapsulate segment criteria, including exclusions and other filtering logic, so they are well suited to centrally managed rules that must be applied repeatedly. The teams can continue to own and refine their brand-specific segment criteria while inheriting the centrally maintained exclusion requirements through the shared block. Salesforce's Data Cloud segmentation documentation describes how segments are built from reusable logic and filters, and Salesforce Trailhead provides additional guidance on designing audiences in Data Cloud. See [Salesforce Help: Create a Segment](#) and [Trailhead: Data Cloud Segmentation](#).

QUESTION NO: 20

A financial institution wants to identify individuals whose total deposits exceed \$100,000 during the previous calendar month. The audience can refresh overnight and does not require event-by-event processing as deposits arrive.

Which Data Cloud feature should the consultant recommend to calculate the total before segmentation?

- A. Calculated insight
- B. Streaming insight
- C. Data Action
- D. Data transform

ANSWER: A

Explanation:

Calculated insight is the appropriate feature because the requirement is an aggregate measure over a defined period that can be computed on a schedule. A calculated insight can sum deposit transactions by individual for the prior calendar month, and the resulting measure can then be used as a segment criterion.

A streaming insight is designed for continuous evaluation of events as they arrive and is more appropriate when an immediate response is required. A data action initiates a downstream action based on qualifying data but does not perform the scheduled aggregation. A data transform can prepare data, but calculated insights are intended for reusable aggregated metrics used in segmentation and analysis.

QUESTION NO: 21

A consultant is configuring an Amazon S3 data stream. Data Cloud must discover timestamped files in a bucket prefix and read the contents of each matching file.

Which two Amazon S3 permissions should be granted to the credentials used by the data stream?

Choose 2 answers

- A. s3:GetObject
- B. s3:PutObject
- C. s3:ListBucket
- D. s3:DeleteObject

ANSWER: A C

Explanation:

s3:ListBucket is required so that Data Cloud can list objects in the configured bucket or prefix and identify files that match the configured filename pattern. This is particularly important when incoming files have timestamped names.

s3:GetObject is required so that Data Cloud can retrieve and ingest the contents of the matching S3 objects. Write-oriented permissions, such as PutObject and DeleteObject, are not required for an inbound S3 data stream and would grant unnecessary access.

QUESTION NO: 22

If a data source does not have a field that can be designated as a primary key, what should the consultant do?

- A. Use the default primary key recommended by Data Cloud.
- B. Create a composite key by combining two or more source fields through a formula field.
- C. Select a field as a primary key and then add a key qualifier.
- D. Remove duplicates from the data source and then select a primary key.

ANSWER: B

Explanation:

Create a composite key by combining two or more source fields through a formula field. Data Cloud requires every record ingested into a data model object to have a primary key value that uniquely identifies that record. This key is essential for reliable ingestion, updates, and identity-related processing because Data Cloud must be able to distinguish one source record from another. When the source system has no single naturally unique field, the consultant should derive a stable unique identifier from a combination of fields whose combined values uniquely identify each record, such as a source-system identifier together with a business or record attribute. A formula field can concatenate or otherwise combine those source values into one key field, which can then be mapped as the primary key. The derived value should be deterministic and consistently populated, so the same source record produces the same key on subsequent ingestions. This approach preserves record-level uniqueness without requiring the source to supply a standalone primary-key column. Salesforce's guidance on mapping data streams and data model objects emphasizes mapping source fields appropriately, including identifying record keys used by Data Cloud. See [Map Data Stream Data](#) and [Data Model Objects](#).