DUMPSARENA

Salesforce Customer Data Platform (CDP)

Salesforce Customer-Data-Platform

Version Demo

Total Demo Questions: 34

Total Premium Questions: 347

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Customer Data Platform (CDP)	48
Topic 2, Data Modeling	47
Topic 3, Data Ingestion and Unification	117
Topic 4, Segmentation and Activation	116
Topic 5, Data Governance	16
Topic 6, Mix Questions	3
Total	347

QUESTION NO: 1

Which data type is collected by a company that does not have any direct relationship with

the customer?

A. Partner data

B. First party data

C. Second party data

D. Third party data

ANSWER: D

Explanation:

Third party data is data obtained from organizations that do not have a direct relationship with the individual whose information is being collected. It is commonly sourced from data brokers, aggregators, advertising networks, or other external providers that compile information across multiple websites, devices, transactions, or public and commercial sources. Because the company using this information did not collect it directly from its own customers or prospects, the relevant relationship is indirect or absent.

In a Customer Data Platform context, third party data can be brought into the platform to enrich customer profiles, support audience segmentation, and improve targeting or analytics. Its use requires careful governance: the organization should confirm that the provider has a lawful basis to collect and share the information, that applicable consent and privacy preferences are honored, and that activation follows relevant privacy obligations. Salesforce's guidance on customer data emphasizes connecting and governing data to create unified profiles, while its overview of first-party data distinguishes directly collected data from externally sourced data. See Salesforce Trailhead: Data Cloud Basics and Salesforce: What Is First-Party Data?.

QUESTION NO: 2

Which of the followingcannotbe used in Segmentation? (Choose 2)

A. Numeric Measures

B. Text Measures

C. Aggregate Functions

D. Date Time Measures

ANSWER: B D

Explanation:

Text Measures and **Date Time Measures** can't be used as measures in Salesforce CDP segmentation. A measure represents a value calculated across related records for a unified individual, such as a count, sum, minimum, maximum, or average. The segmentation engine's measure processing is designed for numerical values that can be evaluated consistently in aggregate conditions. Text values don't provide a numerical result suitable for those measure calculations, and date/time values were likewise not available as supported segmentation measures in the CDP segmentation model covered by this exam. These fields can still be important in the data model and can be used in appropriate attribute-level filtering scenarios, but they are not eligible measure types for segment criteria. When building a segment that requires aggregate logic, use supported numeric measures and select the applicable aggregation and comparison conditions. Salesforce's segmentation learning material describes using profile data, related data, and measure-based criteria to define an audience, while its Data Cloud documentation explains the role of data modeling and calculated values in activation workflows. See Salesforce Trailhead: Data Cloud Segmentation and Salesforce Data Cloud Help.

QUESTION NO: 3

Records from which object can be retrieved via the Salesforce CDP Query API?

A. Segment membership

B. Segment metadata

C. Data Stream metadata

D. Data model objects

ANSWER: D

Explanation:

Data model objects is correct because the Salesforce Data Cloud Query API is the SQL-based interface for retrieving records stored in Data Cloud's harmonized data model. A client submits a SQL query to the Query API and specifies the data model object that contains the desired records in the query's `FROM` clause. This enables applications and external services to retrieve data from standard and custom data model objects after data has been ingested, mapped, and harmonized in Data Cloud. The query can select fields, filter rows, join supported objects, and return the result set asynchronously for use outside Data Cloud. Data model objects provide the business-oriented, standardized representation of customer and related data that Query API is designed to query. Salesforce documents the Query API as an API for querying Data Cloud data with SQL and documents data model objects as the objects used to organize and access harmonized Data Cloud data. See the Salesforce Data Cloud Query API overview and the Data Model Objects documentation for the applicable query and data-model concepts.

QUESTION NO: 4

Options to load company's external customer management data in nearly real time

A. CRM connector

B. Bulk load through S3 bucket

C. Ingestion API

D. MuleSoft Connector

E. Custom Connector

ANSWER: A C D

Explanation:

CRM connector, **Ingestion API**, and **MuleSoft Connector** can support near-real-time customer-data ingestion into Salesforce Data Cloud. The CRM connector is appropriate when the customer-management data resides in Salesforce CRM. Data Cloud can use Change Data Capture to receive record changes from Salesforce CRM as they occur, allowing customer profiles to be refreshed with low latency rather than relying only on periodic full extracts.

The Ingestion API is designed for programmatic ingestion from external applications and services. Its streaming capability lets an external customer-management system send records and updates to Data Cloud as events occur, making it a core integration method for near-real-time use cases. MuleSoft Connector is also suitable when the source is an external system: MuleSoft can connect to and transform data from enterprise applications, databases, and APIs, then use Data Cloud ingestion capabilities to send streaming data. This supports an integration architecture where data is validated and mapped before it reaches Data Cloud. Salesforce documents streaming ingestion through the Ingestion API and the use of connector-based data streams in its Data Cloud Ingestion API documentation and Data Cloud data streams documentation.

QUESTION NO: 5

What does the ignore empty value option do in Identity Resolution?

A. Ignores empty reconciliation rules

B. Replaces the value with a null

C. Ignores empty fields when running reconciliation

D. Modified the value in the field with a predefined value

ANSWER: C

Explanation:

Ignores empty fields when running reconciliation is correct. In Data Cloud Identity Resolution, reconciliation determines which value to retain for an attribute when records that have been matched into the same unified individual contain different values. Enabling the ignore empty value setting ensures that an empty value from one contributing record is not selected as the reconciled value when another matched record supplies a populated value. This helps preserve useful profile information in the unified record and prevents blank source data from replacing a known value during the reconciliation process.

This setting is especially valuable when data comes from systems with different completeness levels. For example, one source may provide a customer email address while another source has the same customer but no email value. With empty values ignored, reconciliation can retain the available email rather than producing an empty result for that attribute. The setting applies to the treatment of empty attribute values as Identity Resolution applies configured reconciliation rules; it supports a more complete and reliable unified profile. Review Salesforce's Identity Resolution learning material at Trailhead: Identity Resolution in Data Cloud and the Salesforce Identity Resolution documentation for configuration concepts and reconciliation behavior.

QUESTION NO: 6

Which of these is present in Identity Resolution Summary?

A. Unified Individuals

B. Last Processing Status

C. Matched Individuals

D. All the above

ANSWER: D

Explanation:

All the above is correct because the Identity Resolution Summary provides an at-a-glance view of both the outcome and operational state of an identity resolution ruleset. It includes the number of Unified Individuals created after the ruleset reconciles source records into unified profiles. It also shows Matched Individuals, which indicates the individual records that have been associated through the configured matching and reconciliation logic. In addition to these result metrics, the summary displays the Last Processing Status so an administrator can verify whether the most recent identity resolution run completed successfully and understand the current processing state. Together, these details help users assess the scale of profile unification and monitor the latest execution of the ruleset without having to inspect individual records. Salesforce describes identity resolution as the process of matching records representing the same person and reconciling them into unified profiles. See Salesforce Help: Identity Resolution and Trailhead: Get to Know Identity Resolution.

QUESTION NO: 7

A developer wants to create a segment for a birthday campaign which is refreshed every day

and activateD. Which operator should he use?

A. Is Birthday

B. Is This Year

C. Is Anniversary Of

D. Is Between

ANSWER: C

Explanation:

Is Anniversary Of is the appropriate operator for a daily refreshed birthday campaign because a birth date is a recurring annual event. This operator evaluates the month and day of a stored date as an anniversary relative to the current date, rather than treating the original year of birth as part of the match. As a result, a customer whose birth date falls on the relevant day can qualify every year, regardless of age or the year recorded in the source data.

When the segment is configured to refresh daily, Data Cloud reevaluates the anniversary condition against the new current date during each refresh. The audience therefore changes automatically from one day's birthday recipients to the next day's recipients. After the segment is activated, the connected activation target receives the current eligible audience according to the configured activation schedule. This supports an always-current birthday journey or campaign without requiring the developer to manually update date ranges each day or rebuild the segment annually.

Salesforce describes segmentation as the process of defining audiences from unified data in Create Segments in Data Cloud. For additional Data Cloud segmentation concepts and activation workflow context, see Activate Data in Data Cloud.

QUESTION NO: 8

What are the two benefits of Calculate Insights compared to Segment criteria?

A. It creates simple logic on row-based operations

B. It creates reusable contents

C. Marketer friendly drag and drop interface

D. It performs complex queries on multiple objects

ANSWER: B D

Explanation:

Calculated Insights create derived metrics by querying, joining, filtering, and aggregating data across the Data Cloud data model. This makes "It performs complex queries on multiple objects" a key benefit: an insight can calculate measures such as lifetime spend, order counts, average order value, or engagement totals from related records, then make that result available for audience building. Rather than attempting to express all of that aggregation directly in an individual segment's criteria, the calculation is defined once as a durable metric.

Calculated Insights are also reusable assets. "It creates reusable contents" is correct because, once published, the resulting calculated-insight fields can be used repeatedly in segments and other downstream Data Cloud use cases. For example, a marketer can use one standardized lifetime-value or recent-purchase calculation in several audiences without rebuilding the underlying query logic for each audience. This promotes consistent metric definitions and reduces repeated configuration work as segmentation needs grow. Salesforce describes Calculated Insights as a way to define and use aggregate measures across unified individual and related data. See Salesforce's Calculated Insights documentation and the Data Cloud segmentation Trailhead module.

QUESTION NO: 9

What role in a company was Salesforce CDP designed for?

A. Developer

B. Marketer

C. Salesperson

D. Data Analyst

ANSWER: B

Explanation:

Salesforce CDP, now called Data Cloud, was designed primarily for marketers who need to unify customer data from many sources and use it to deliver relevant, personalized engagement. Marketers can use the platform's unified customer profiles, calculated insights, segmentation, and activation capabilities to define audiences and send those audiences to engagement and advertising destinations. This supports activities such as targeting customers based on behavior or attributes, suppressing audiences, personalizing campaigns, and measuring marketing outcomes using a more complete view of each customer. Although technical teams and data specialists commonly help configure data ingestion, identity resolution, governance, and integrations, the core business value and original positioning of Salesforce CDP centers on enabling marketing teams to use customer data without depending on custom engineering for every audience-building task. Salesforce describes Data Cloud as a platform that connects and harmonizes customer data to create unified profiles and power personalized experiences across Salesforce applications, including Marketing. See Salesforce's Data Cloud overview and the Data Cloud Basics Trailhead module for its customer-data and audience-activation capabilities.

QUESTION NO: 10

Which API would you use to surface additional information from the CDP into Service Cloud?

A. Query API

B. Unified Profile API

C. Calculated Insights API

D. Profile API

ANSWER: D

Explanation:

The **Profile API** is designed to retrieve information about a unified customer profile from Salesforce CDP (now Data Cloud) so that it can be displayed and used in connected Salesforce experiences, including Service Cloud. It gives Service Cloud users access to a customer's consolidated, identity-resolved profile and related profile attributes at the point where they are assisting the customer. This helps agents work with a more complete customer context rather than relying only on data stored in the Service Cloud org.

In practice, an integration can use the Profile API to request profile information for an identified individual and surface that context in a Service Cloud agent experience, such as a custom Lightning component or other integrated service workflow. The API supports the core purpose of CDP: bringing together data from multiple sources, resolving it into a unified profile, and making that profile actionable in engagement and service use cases. Salesforce's Data Cloud API documentation describes the Profile API as the interface for accessing unified profile data, while Salesforce's Service Cloud and Data Cloud guidance highlights using unified customer context to improve agent interactions. See Salesforce Profile API documentation and Salesforce Data Cloud documentation.

QUESTION NO: 11

What are the two refresh models available for a cloud storage data stream?

A. Update

B. Upsert

C. Append

D. Full Refresh

ANSWER: B D

Explanation:

Cloud storage data streams support the **Upsert** and **Full Refresh** refresh models. Upsert is designed for incremental file delivery: Data Cloud uses the configured record identifier to update an existing record when a matching record is received and to insert a record when no match exists. This model is appropriate when a source provides changed or newly available records over time and the target data should retain its current state without reloading every record in each run.

Full Refresh is intended for files that represent a complete replacement snapshot of the source dataset. Each refresh reloads the incoming full data set so that the resulting data reflects the latest complete extract. Selecting the proper model is important for accurate record maintenance, identity processing, and downstream calculated insights or activations. Salesforce describes cloud-storage ingestion and data-stream configuration in its Data Cloud documentation and training materials, including Data Streams in Salesforce Help and the Data Cloud Ingestion Trailhead module.

QUESTION NO: 12

What type of data is collected by a company or entity that does not have a direct relationship with

the visitor and customer?

A. Partner Data

B. First Party Data

C. Third Party Data

D. Second Party Data

ANSWER: C

Explanation:

Third Party Data is data gathered by an organization that has no direct relationship with the individual whose information is being collected. It is commonly obtained from external data providers, aggregators, publishers, advertising networks, or other entities that compile information across websites, devices, and sources. The organization using the data did not collect it directly through its own customer interactions, transactions, forms, website visits, service usage, or other first-hand touchpoints.

In a customer data platform strategy, third-party data can be used to supplement a company's existing customer profiles, such as adding demographic, interest, or audience attributes. However, organizations must govern its use carefully: they need to understand the source, permissions, quality, contractual terms, and applicable privacy requirements before activating it. This distinction is important because the directness of the relationship between the collector and the individual determines whether data is classified as first-party, second-party, or third-party data. Salesforce describes third-party data as information collected by an entity that does not have a direct relationship with the customer in its overview of first-, second-, and third-party data. See also Salesforce's guidance on customer data platforms for the role of unified, governed customer data.

QUESTION NO: 13

When activating on Unified Individual which three attributes are automatically included?

A. Contact address

B. Email SK

C. Contact phone number

D. Contact ID

E. Contact email address

ANSWER: B C E

Explanation:

When an activation is based on the Unified Individual, Data Cloud automatically carries key contactability attributes needed to make the resulting audience actionable in downstream systems. **Email SK** provides the email contact-point key used to identify the selected email record. **Contact phone number** provides the phone contact detail associated with the unified profile, and **Contact email address** provides the address required for email-based engagement. Including these fields by default supports activation use cases without requiring a segment builder to manually select the core email and phone contact-point data each time. This behavior reflects Data Cloud's unified-profile model: a Unified Individual can be related to multiple contact points, and activations need the relevant contact data and identifier to deliver audience members appropriately. Review Salesforce's Data Cloud activation guidance at Activate Data Cloud Segments and the Data Cloud developer documentation at Data Cloud Reference Guide.

QUESTION NO: 14

Which permission set needs to be added to the salesforce CRM object to beavailable in Salesforce CDP

A. Salesforce CDP external connector

B. Salesforce External Integration

C. Salesforce CDP Salesforce Connector Integration

D. Salesforce External Connector

ANSWER: C

Explanation:

Salesforce CDP Salesforce Connector Integration is the required permission set for the integration user used by the Salesforce CRM connector. This connector securely makes CRM data available to Salesforce CDP (now called Data Cloud) so that selected standard and custom CRM objects can be ingested and mapped into the data model. The permission set grants the connector integration the access it needs to read the CRM metadata and records that are authorized for use in Data Cloud. It should be assigned to the Salesforce user identified as the connector's integration user, following the principle of least privilege and with any required object- and field-level access configured for the CRM data being brought into Data Cloud. After the connector is authenticated and the integration user has this permission set, an administrator can select eligible CRM objects and fields while configuring data streams. Salesforce's setup guidance for the CRM connector and its integration-user permissions is available in the Salesforce CRM connector documentation. For broader Data Cloud access and setup concepts, see the Data Cloud Basics Trailhead module.

QUESTION NO: 15

Which two dependences can prevent a Data Stream for being deleted?

A. A data stream attribute is used in Segmentation

B. A data stream attribute is used in Calculated Insight

C. A data stream attribute is mapped to a Data Model object

D. A data stream attribute is used in Activation

ANSWER: B C

Explanation:

A data stream attribute being used in Calculated Insight is a deletion-blocking dependency because calculated insights are persisted metric definitions that reference attributes from the Data Cloud data model. Deleting the underlying stream without first removing or updating that reference could invalidate the insight's logic and its resulting calculated-insight data. Salesforce therefore requires dependent calculated insights to be addressed before the associated data stream can be removed.

A data stream attribute being mapped to a Data Model object is also a deletion-blocking dependency. Mapping is the configuration that connects incoming source fields to Data Model Objects and makes the ingested data available for harmonization and downstream Data Cloud processing. Since this mapping is part of the stream's active data-model configuration, it must be removed before the data stream can be deleted. This safeguard preserves metadata integrity and prevents an active model configuration from referring to an ingestion source that no longer exists. Salesforce's Data Cloud learning materials describe how data streams use mappings to populate Data Model Objects and how calculated insights are defined from Data Cloud data: Salesforce Trailhead: Get Started with Data Cloud and Salesforce Help: Calculated Insights.

QUESTION NO: 16

Which CDP permission set manages the overall segmentation strategy and identifiesthe target campaigns?

A. Marketing specialist

B. IT manager

C. Marketing manager

D. Data Aware Specialist

ANSWER: C

Explanation:

Marketing manager is the appropriate CDP permission set for the person who owns the high-level audience and campaign strategy. This role is focused on translating marketing objectives into a segmentation approach: determining the relevant customer groups, defining the intended campaign audiences, and directing how those audiences support business goals. In Salesforce CDP, now called Data Cloud, access and responsibilities are commonly organized around personas so that strategic marketing users can work with audiences and activation planning without needing to perform the underlying data-engineering administration. A marketing manager therefore provides the strategic direction for segmentation and identifies the campaigns that should use the resulting target audiences. This responsibility connects audience design to marketing outcomes, helping ensure segmentation is purposeful, aligned to campaign goals, and ready for activation in supported engagement destinations. Salesforce's learning material on Data Cloud audience capabilities describes how marketers use unified customer data to create audiences for activation, while Salesforce Help documents the permission-set model used to grant Data Cloud capabilities. See Salesforce Trailhead: Customer 360 Audiences for Marketers and Salesforce Help: Data Cloud Permission Sets.

QUESTION NO: 17

Which two activities can a marketer perform after a segment is published?

A. Create an activation

B. Review segment membership

C. Modify the source system schema

D. Create a new Data Lake Object

ANSWER: A B

Explanation:

Create an activation is correct because a published segment can be sent to an activation target, such as Marketing Cloud Engagement, Marketing Cloud Personalization, advertising destinations, or cloud storage. The activation defines the destination and attributes to deliver.

Review segment membership is also correct because publishing evaluates the segment criteria and generates the audience membership. Reviewing the resulting membership helps a marketer validate that the intended profiles qualify before using the audience downstream. Publishing does not modify the underlying Data Lake Object schema or create an

identity resolution ruleset. See Salesforce's Data Cloud Segments documentation and the Data Cloud Segmentation and Activation Trailhead module.

QUESTION NO: 18

What are two ways a marketing user can incorporate a campaign code from an external (…) bucket?

A. Include campaign code into the filename specification

B. Include campaign code into the activation name

C. Include campaign code into the segment name

D. Include campaign code into the activation description

ANSWER: C D

Explanation:

Including the campaign code into the segment name and including the campaign code into the activation description are appropriate ways to retain a campaign identifier as part of the Data Cloud audience and activation configuration. A segment name is a meaningful business label for the audience definition, so adding the code enables marketers to identify the campaign context when locating, reviewing, reusing, and managing the segment. The activation description provides supporting metadata for the outbound activation and can document the applicable campaign code for operational users who need to understand the activation's intended business purpose. Together, these fields make the campaign identifier visible within Data Cloud's segmentation and activation workflow, helping with governance, handoffs, auditing, and campaign-level organization without requiring the code to be a customer attribute. Salesforce describes segments as reusable audiences built from customer data, and activations as the process for sending a segment's members to a target. Naming and descriptive metadata are therefore useful places to preserve business context such as a campaign code. See Salesforce's segment creation guidance and activation creation guidance.

QUESTION NO: 19

Salesforce CDP supportsthese content file formats?

A. .txt

B. .csv

C. .tsv

D. .json

ANSWER: B C

Explanation:

`.csv` and `.tsv` are supported content file formats for loading delimited file data into Salesforce CDP, now called Data Cloud. A CSV file represents tabular records with fields separated by commas, while a TSV file represents the same tabular structure using tab characters as delimiters. These formats are appropriate for file-based ingestion because Data Cloud can interpret the header row as field names and map the resulting source fields into a data stream and its associated data lake object. After ingestion, those fields can be mapped to the Data Cloud data model for use in identity resolution, calculated insights, segmentation, and activation. File-upload workflows require consistently structured rows and a stable delimiter, so CSV and TSV are the supported delimited-content choices covered by this question. Salesforce's Data Cloud documentation describes creating data streams to bring source data into Data Cloud, including file-based ingestion and subsequent mapping. See the Salesforce Data Cloud data streams documentation and the Data Cloud Basics Trailhead module for the overall ingestion and harmonization process.

QUESTION NO: 20

Which user rolemanages the data model and creates Identity Resolution rulesets for unified

profiles?

A. Admin

B. Marketing Specialist

C. Data Aware Specialist

D. Marketing Manager

ANSWER: C

Explanation:

Data Aware Specialist is correct because this Data Cloud persona is responsible for the foundational data work required to make customer data usable across the platform. That work includes understanding and managing the data model, mapping and harmonizing incoming data into Data Model Objects, and configuring the identity matching and reconciliation logic used to create unified individual profiles. Identity Resolution rulesets define how Data Cloud recognizes records that represent the same person across sources, using match rules and reconciliation rules to produce a unified profile. Because these rules depend on an accurate understanding of source data, object relationships, identifiers, and data quality, they align directly with the Data Aware Specialist's responsibilities. This role provides the trusted, unified customer-data foundation that downstream teams can use for segmentation, activation, analytics, and personalization. Salesforce describes Identity Resolution as the capability that matches records from disparate sources and reconciles them into unified profiles; its configuration is therefore a core data-management activity. See Salesforce's Identity Resolution developer documentation and the Identity Resolution in Data Cloud Trailhead module for additional detail.

QUESTION NO: 21

How can a marketer change attribute names to match personalization in an Activation Target?

A. Update field names in data model

B. Set preferred attribute names when configuring activation target

C. Update attributes names in the data stream configuration

D. Set preferred attribute names when configuring activation

ANSWER: D

Explanation:

Set preferred attribute names when configuring activation is correct because attribute naming for an outbound audience is controlled at the activation level. When a marketer creates or configures an activation, they select the attributes to send to the chosen activation target and can define the preferred names used for those attributes in the target payload. This lets the marketer align Data Cloud attribute names with the field names, merge fields, or personalization attributes expected by the destination system without changing the underlying Data Cloud data model. The configured names apply to that activation's output, so the same source attribute can be represented differently when required by different downstream destinations. This approach preserves the source schema while ensuring the activation delivers an audience in the format needed for target-side personalization. Salesforce's activation workflow includes selecting and mapping attributes as part of configuring an activation; see Create an Activation. For additional Data Cloud activation concepts and destination configuration guidance, see Salesforce Help: Create an Activation.

QUESTION NO: 22

What should a Marketer do first when implementing Salesforce CDP?

A. Manage Consent

B. Identify Data Sources

C. Obtain Cross-organizational buy-in

D. Identify Activation Targets

ANSWER: B

Explanation:

Identify Data Sources is the appropriate first step because Salesforce CDP, now Salesforce Data Cloud, depends on bringing together customer data from the systems where it is created and maintained. A marketer must first understand which sources contain the information needed for intended customer insights and use cases, such as CRM records, commerce activity, web engagement, service interactions, loyalty data, and consent-related data. This inventory establishes the practical scope of the implementation: what data is available, who owns it, how it can be connected, how often it should be refreshed, and whether it contains the identifiers needed to unify customer profiles.

Once sources are known, the implementation team can configure connectors and ingestion, map incoming data to the Data Cloud model, establish identity resolution rules, and create audiences for downstream activation. Identifying sources early also helps marketers recognize data-quality gaps and coordinate with the business and technical owners responsible for resolving them. This is the foundation for later decisions about activation destinations and consent-aware engagement. Salesforce describes Data Cloud as a platform that unifies data from multiple sources to create a harmonized customer profile; see the Salesforce Data Cloud Basics Trailhead module and the Salesforce Data Streams documentation.

QUESTION NO: 23

Which two features can be used to validate the data in the unified profile object? [Answer need confirmation]

A. Query API

B. Data Explorer

C. Segmentation

D. Identity Reconciliation

ANSWER: A B

Explanation:

Query API and **Data Explorer** can both be used to validate data in the unified profile object. Data Explorer provides an in-product way to inspect Data Cloud objects and their records. An administrator or data practitioner can use it to review unified profile attributes, verify that expected source data is present, examine record values, and check whether identity resolution and mapping results are producing the expected unified data. This makes it useful for interactive validation during configuration and troubleshooting.

Query API provides programmatic access to Data Cloud data through SQL queries. It can be used to retrieve and assess unified profile records at scale, validate attribute population and values, compare outcomes against expected criteria, and support automated quality checks or external validation workflows. Because the API queries Data Cloud objects directly, it is appropriate when validation must be repeatable, integrated into another system, or performed over larger result sets. Salesforce describes Data Cloud's unified data model and tools in the Data Cloud Basics Trailhead module, and documents API access in the Data Cloud Developer Guide.

QUESTION NO: 24

Which of the following functions returns the first value from a list that isn't empty?

A. CONCAT

B. COALESCE

C. REPLACE

D. FIND

ANSWER: B

Explanation:

`COALESCE` returns the first non-null value from the supplied list of expressions. In Salesforce Data Cloud, this is useful when customer data can arrive from multiple sources or when different fields may be populated depending on the record. For example, a calculated insight or data transformation can use `COALESCE` to select the first available email address, phone number, or other attribute from an ordered set of possible values. The function evaluates the expressions in sequence and returns the first value that is not null; if every supplied expression is null, its result is null. This provides a concise, predictable way to establish a fallback sequence while retaining the preferred value whenever it exists. In the wording of the question, "isn't empty" corresponds to the practical intent of finding the first available value, although the documented function behavior is specifically based on null values. Salesforce documents `COALESCE` as a conditional expression that returns the first non-null expression. See the Salesforce Data Cloud SQL functions documentation and the Data Cloud SQL COALESCE reference.

QUESTION NO: 25

How do the Data streams that use the Marketing Cloud Connector refresh data?

A. Manually refreshed by CDP Admin

B. Marketing Cloud's Automation Studio handles this process.

C. APIs are used for refresh

D. None of the above

ANSWER: B

Explanation:

Marketing Cloud's Automation Studio handles this process. For Marketing Cloud Engagement data streams, Data Cloud relies on an Automation Studio automation in the connected Marketing Cloud Engagement account to prepare and transfer the selected source data for ingestion. The automation performs the recurring export and delivery activity that makes updated Marketing Cloud Engagement data available to Data Cloud, so refresh behavior is controlled through the configured automation and its schedule. This approach supports repeatable, scheduled ingestion of Marketing Cloud Engagement data without requiring an administrator to manually initiate each refresh. When configuring the connector, the implementation should ensure that the relevant Automation Studio process is active and scheduled appropriately for the organization's data-latency requirements. Salesforce describes Automation Studio as the Marketing Cloud Engagement capability for creating and scheduling multi-step automations, including activities that import, extract, and move data. The Data Cloud Marketing Cloud Engagement connector documentation further describes the connector's use for bringing Marketing Cloud Engagement data into Data Cloud. See Salesforce Marketing Cloud Engagement Automation Studio documentation and Salesforce Data Cloud Marketing Cloud Engagement data stream documentation.

QUESTION NO: 26

Which two frequency options are available when ingesting MarketingCloud Data Extensions?

A. 15minutes

B. 24 hours

C. 60 minutes

D. 12 hours

ANSWER: B C

Explanation:

When configuring a Marketing Cloud Engagement data stream for Data Cloud, the available recurring ingestion schedules for a selected Data Extension include **24 hours** and **60 minutes**. The 60-minute schedule supports hourly synchronization when business processes need reasonably current Marketing Cloud data in Data Cloud, such as for audience preparation and downstream activation. The 24-hour schedule supports a daily refresh pattern for Data Extensions whose data does not need to be synchronized throughout the day. The ingestion schedule is configured at the data-stream level and determines how often Data Cloud retrieves updated records from the Marketing Cloud source after the stream is deployed. Selecting the appropriate supported cadence should reflect the required data freshness while accounting for source-data update patterns and operational needs. Salesforce's Data Cloud documentation describes data streams as the mechanism used to connect, ingest, and refresh source data, while the Marketing Cloud Engagement connector documentation covers bringing Marketing Cloud data, including Data Extensions, into Data Cloud. See Salesforce Help: Data Streams and Salesforce Developer Documentation: Data Cloud Connectors.

QUESTION NO: 27

ActivationTarget auto configured after connected with the platform

A. SFMC

B. S3 Bucket

C. Personalization through Interaction Studio

D. Commerce Cloud

ANSWER: A C

Explanation:

SFMC and **Personalization through Interaction Studio** are automatically provisioned activation destinations when their respective Salesforce platform connections are established. This built-in provisioning reduces setup work: after the connection is authorized, Data Cloud recognizes the destination's native Salesforce integration and makes the corresponding activation target available for publishing audiences. Administrators can then select that target while creating an activation, choose the appropriate audience, map the supported attributes, and schedule or publish the activation to the connected Salesforce product.

This behavior supports a unified Salesforce workflow in which customer profiles and calculated audiences built in Data Cloud can be operationalized in Marketing Cloud Engagement for messaging and in Personalization for real-time, individualized experiences. The automatically created target still relies on the connection's configured credentials, permissions, and destination-specific configuration, so those details should be validated before an activation is published. Salesforce's activation learning material describes using activation targets to send Data Cloud audiences to connected engagement systems, while the Data Cloud implementation guidance covers connecting and using downstream Salesforce applications. See Salesforce Trailhead: Activate Audiences and Salesforce Help: Activation Targets.

QUESTION NO: 28

Which two steps required when setting up a Marketing Cloud activation?

A. Set publish schedule

B. Choose an Activation Target

C. Set subscriber key field

D. Select the Email Contact Point

ANSWER: B D

Explanation:

Choosing an Activation Target is required because an activation must identify the destination system that will receive the segment audience. For a Marketing Cloud activation, this selection establishes Marketing Cloud Engagement as the endpoint and enables Data Cloud to prepare the audience data for delivery to that destination. The target selection is therefore a foundational part of configuring the activation.

Selecting the Email Contact Point is also required for an email-oriented Marketing Cloud activation because Data Cloud needs the appropriate contact-point data to identify the email address associated with each unified individual in the audience. This ensures the activated audience contains the contact information Marketing Cloud Engagement needs to address email communications correctly. Contact points are part of the customer data model and allow Data Cloud to use the appropriate channel-specific identity and address data when an audience is activated. Salesforce describes activation as the process of sending segment data to a selected destination, and its Data Cloud learning materials explain the role of profile and contact-point data in creating actionable audiences. See Salesforce Trailhead: Data Cloud Basics and Salesforce Help.

QUESTION NO: 29

Which of the following are characteristics of Formulas? (Choose 2)

A. Simple logic on a row-based operation

B. Ease of use, self-service

C. Highly reusablecontent

D. Attribute updated regularly

ANSWER: A B

Explanation:

Formulas are intended for **simple logic on a row-based operation**. In Data Cloud, a formula evaluates values available on an individual record, such as combining fields, applying conditional logic, changing text or date formats, or deriving a new value from existing attributes. This makes formulas appropriate when the calculation does not require grouping records or aggregating data across a population. Formula capabilities are therefore commonly used as part of data preparation and transformation to produce a clean, usable attribute at the record level.

Ease of use, self-service is also a key characteristic. Formula builders provide a declarative way for business users and administrators to create derived fields using supported functions and operators, without requiring custom code. This supports faster iteration and lets teams implement common transformation logic directly in the Data Cloud user interface. Salesforce describes data transforms as tools for preparing and transforming ingested data, including the creation of derived values through transformation logic. For further guidance, see Salesforce's Data Transforms documentation and the Data Cloud data modeling and transformations Trailhead module.

QUESTION NO: 30

Which data model entity do you map the Name field from Salesforce CRM Contact object?

A. Account

B. Contact

C. Party Identifier

D. Individual

ANSWER: D

Explanation:

Individual is the appropriate Data Cloud data model entity for the Name value from a Salesforce CRM Contact record. In Data Cloud's standard data model, an Individual represents a person and holds person-level attributes used to create a unified customer profile. A CRM Contact is likewise a person record, so its identifying name information belongs on the Individual entity when mapping the source object into Data Cloud. This mapping supports identity resolution and profile unification by associating the person's CRM details with the corresponding individual record across data sources. The Account entity instead represents an organization or business relationship, while a party identifier is used to store identifiers that help recognize and reconcile a person rather than the person's name itself. Salesforce documents the standard Data Cloud data model and its entities in its Data Model Objects reference. For an overview of modeling and mapping source data to Data Cloud's standard model, see the Data Cloud Data Modeling Trailhead module.

QUESTION NO: 31

Which feature determines the attribute value in the file when activating on the Unified Individual?

A. Einstein Activation

B. Reconciliation Rules

C. Data Stream Sequence

D. Random Selection

ANSWER: B

Explanation:

Reconciliation Rules determine which source value is retained for an attribute on a Unified Individual. Data Cloud can ingest multiple records that represent the same person from different sources, and identity resolution links those records into a unified profile. When the linked source records contain different values for the same attribute, reconciliation applies the configured rule to select the value that appears on the unified record.

Activation uses the attributes available on that unified profile when it prepares records for a target. Therefore, the value delivered in an activation file for a Unified Individual reflects the outcome of the applicable reconciliation rule, such as a rule based on source priority or most recently updated information. This makes reconciliation an essential configuration step when activation needs a consistent, trusted value for fields such as email address, phone number, or loyalty status.

For more detail, review Salesforce Trailhead's discussion of configuring reconciliation rules and the Salesforce Help documentation on reconciliation rules in Data Cloud.

QUESTION NO: 32

Which two options are available for automated batch publish and activation?

A. 2 − 1 hour

B. − 24 hours −

C. 15 minutes −

D. 12 hours

ANSWER: B D

Explanation:

Automated batch publishing and activation supports schedules of **24 hours** and **12 hours**. These intervals allow an organization to refresh a segment's membership, publish the resulting audience, and send the updated audience to its activation targets on a predictable recurring cadence. A 12-hour schedule is useful when audiences need to be refreshed twice daily, while a 24-hour schedule supports a daily operational refresh. Selecting an automated cadence is appropriate

when the underlying data and calculated audience membership are updated regularly and downstream systems need timely, repeatable audience delivery without a user manually initiating each publish.

In Data Cloud, publishing makes segment results available for activation, and activation delivers those results to a configured destination. The schedule therefore governs the recurring batch process used to keep activated audiences aligned with the latest eligible customer and prospect data. Administrators should select the cadence that meets the business's freshness requirements while allowing sufficient time for source-data ingestion, identity resolution, calculations, and segment processing. See Salesforce Trailhead's Activate Segments unit and the Salesforce Data Cloud segment publishing documentation for additional guidance.

QUESTION NO: 33

Which two statements describe a Data Cloud calculated insight?

A. It contains dimensions and measures

B. It can be used in segment criteria

C. It replaces Data Lake Objects

D. It is configured with SOQL

ANSWER: A B

Explanation:

It contains dimensions and measures is correct because a calculated insight defines the dimensions used to group the result and the measures that calculate aggregate values. For example, an insight can group by unified individual and calculate the sum of purchase amounts.

It can be used in segment criteria is also correct. Once a calculated insight is published, its resulting metrics can be used to qualify an audience, such as identifying customers whose lifetime spend exceeds a specified amount. Calculated Insights provide reusable aggregate metrics rather than a one-time segment filter. See Salesforce's Calculated Insights documentation.

QUESTION NO: 34

Which two options are available when configuring a segment publish schedule?

A. Manual refresh

B. Don't refresh

C. Hourly

D. 24 hours

ANSWER: B D

Explanation:

"Don't refresh" and "24 hours" are the available segment publish-schedule choices in the Salesforce Customer Data Platform scheduling experience covered by this question. Selecting "Don't refresh" publishes the segment without setting up a recurring refresh. This is appropriate when the audience should remain based on the results at publication time until a user intentionally republishes it. Selecting "24 hours" establishes a daily recurring refresh, so Salesforce reevaluates the segment membership and republishes the resulting audience on that cadence. A recurring schedule is important when segment criteria depend on customer data that changes over time, because the published audience can stay aligned with the latest eligible individuals without requiring a separate publication action every day.

Segment publishing is the step that makes a defined audience available for downstream activation and related use cases. The schedule controls the ongoing refresh behavior of that published audience rather than changing the segment's defining

criteria. Salesforce's Data Cloud learning materials describe segment creation, publishing, and activation concepts in the Customer 360 Audiences Segmentation module. Additional Salesforce product documentation is available through the Salesforce Help Center.