

Salesforce (CRM) Integration Architecture and Data Flow

Version 1 · 12 October 2026 · For architecture and information security review

This document shows how Tapper connects to a customer's Salesforce org, what it reads, where the data goes inside Tapper, and what leaves Tapper. The integration answers one question: which Google ad click produced this lead or deal, and what was it worth. Tapper uploads that outcome to the customer's own Google Ads account as an offline conversion, so bidding favours clicks that produce qualified leads and sales rather than form fills.

OAuth 2.0

Authorization code flow with PKCE

0

Writes to Salesforce of any kind

5

Standard objects read, fixed field list

SHA-256

Email and phone hashed before storage

1

Active destination: the customer's Google Ads account

Belgium

Google Cloud europe-west1

The data boundary

IN / FROM SALESFORCE

A fixed list of fields on five standard objects (section 4). One backfill of the last 12 months when the connection is enabled, then every 15 minutes only the records changed since the previous read.

INSIDE TAPPER

Email and phone are hashed with SHA-256 in the sync worker and the raw values are discarded; they are never written to the database, the warehouse, logs or queue messages. OAuth tokens are encrypted with Cloud KMS before storage. All storage is in Google Cloud europe-west1 (Belgium).

OUT / TO THE CUSTOMER'S AD ACCOUNTS

For each qualified outcome: the click id, the conversion action, the value and currency, the time Tapper observed the stage, and an order id for de-duplication. Today these go to the customer's Google Ads account only. No email, phone, hash, name or free text leaves Tapper.

1. Architecture

Three parties take part: the customer's Salesforce org, Tapper's platform on Google Cloud, and the customer's Google Ads account. Every connection is outbound from Tapper over HTTPS; Salesforce never connects to Tapper except to return the user to Tapper after sign-in.




Figure 1. Components and data boundaries. Navy arrows: calls between Tapper and Salesforce. Green: the only data that leaves Tapper.

2. Sequence

The five phases of the integration, from connection to disconnection.




Figure 2. Sequence of calls. Navy: calls between Tapper and Salesforce. Green: data sent to Google Ads.

3. Technical details

AREA	HOW IT WORKS
Authentication	OAuth 2.0 authorization code flow with PKCE (S256) and a client secret, through Tapper's Salesforce External Client App. Scopes requested: api, refresh_token, offline_access. No password is ever shared with Tapper.
Who can connect	Only an owner or admin of the customer's Tapper workspace can start a connection. Each authorisation carries a random 256-bit state value that expires after 10 minutes, is single use, and must return to the same Tapper user and browser session.
Salesforce permissions	The integration user needs API Enabled and Read on Lead, Contact, Opportunity and the fields in section 4. Salesforce offers no read-only OAuth scope, so read-only access is enforced by the permission set the customer assigns. Tapper sees only the records that user can see.
Writes to Salesforce	None. The only POST requests are the token exchange, the token revoke and the creation of a Bulk API query job. There is no create, update, upsert, delete or Apex call, and Tapper creates no fields, layouts, automation or users.
API	Salesforce REST API and Bulk API 2.0, version 67.0. REST queries read changes; Bulk API 2.0 query jobs read the one-time backfill. No Streaming API, Change Data Capture, Platform Events or outbound messages.
Frequency	One backfill of the last 12 months when the connection is enabled. After that, every 15 minutes, only records changed since the previous read, up to 1,000 per page. Syncing stops automatically if Salesforce returns an authentication or API limit error.
Personal data	Email and phone are normalised and hashed with SHA-256 in the sync worker before anything is written; the raw values are discarded and never stored in the database, the warehouse, logs or queue messages. Names, addresses, notes, descriptions and activity history are never requested.
Matching	An outcome is linked to an ad click through the Google click id that the customer's web form stores on the Lead. Tapper does not match people across customers.
Data sent out	Today, to the customer's own Google Ads account only, through the Google Ads API: the click id, the conversion action, the value and currency, the time Tapper observed the stage, and an order id made of the Salesforce record Id and the stage name, used by Google to reject duplicates. Consent is sent as unspecified; uploads for people in the European Economic Area require the customer's consent attestation first. The Data Disclosure also covers uploads to Meta and TikTok; those are not active yet, and the same limits on fields apply when they are. No other party receives Salesforce data.

AREA	HOW IT WORKS
Storage and region	Google Cloud, europe-west1 (Belgium): Cloud SQL for PostgreSQL (encrypted tokens and lead outcomes, keyed to the customer's connection, SSL-only connections) and BigQuery (copies of the records read, in a dataset for that Salesforce org, with email and phone keys refused at write).
Encryption	In transit: HTTPS with TLS 1.2 or higher on every call. At rest: Google Cloud encryption (AES-256) on all storage. OAuth tokens and the client secret are additionally encrypted with a Cloud KMS key in europe-west1, each bound to its own database column. Tokens are excluded from every analytics copy.
Disconnection	The customer disconnects in Tapper at any time. Tapper revokes the token with Salesforce and deletes the stored tokens and client secret. The customer can also remove the app in Salesforce, which ends access immediately.
Visibility in Tapper	The Tapper web app shows the connection's status and settings only. It never displays Salesforce records, hashes or tokens.
Network	Every call is outbound from Tapper's Google Cloud project through Cloud NAT. Salesforce reaches Tapper only through the browser redirect at the end of sign-in.

4. Fields read

The field list is fixed in Tapper's source code. Changing it requires a code change, review and release.

OBJECT	FIELDS
Lead	Id, Email*, Phone*, Status, LeadSource, CreatedDate, LastModifiedDate, IsConverted, ConvertedOpportunityId, ConvertedDate, the click id field; if present: RecordTypeId, Rating, CurrencyIsoCode
Contact	Id, Email*, Phone*, CreatedDate, LastModifiedDate
Opportunity	Id, StageName, Amount, CloseDate, CreatedDate, LastModifiedDate, IsWon, IsClosed, the click id field; if present: CampaignId, RecordTypeId, CurrencyIsoCode
OpportunityContactRole	OpportunityId, ContactId, IsPrimary; if present: Role
OpportunityHistory	Id, OpportunityId, StageName, CreatedDate, Amount

* Hashed with SHA-256 on arrival; the raw value is never stored. The only financial figure read is the Opportunity Amount and its currency. No Account, Case, Task, Event, Campaign member, Order, Quote, Invoice, Payment or custom object is ever queried.

5. Setup by the customer's Salesforce admin

1. Create an integration user (Salesforce's Integration User licence is suited to this) or choose an existing user for Tapper.
2. Assign a permission set with API Enabled and Read on Lead, Contact (Salesforce also requires Read on Account for this), and Opportunity, with read access to the fields in section 4. Grant no edit, create or delete permissions.
3. Make sure the web forms that create Leads store the Google click id (gclid) in a custom field on the Lead. Use the same field API name on the Opportunity, mapped through lead conversion, so the click id follows a converted lead.
4. In Tapper, a workspace owner or admin clicks Connect Salesforce and signs in as the integration user. Sync starts when it is enabled after the click id field is found.
5. Agree with Tapper which lead and opportunity stages count as qualified, which mean cancelled, and the value of each, so uploads reflect real outcomes.

Related documents on the Tapper Security, Data & Privacy page: the Salesforce (CRM) Integration Data Disclosure; Security, Data Protection & Platform Architecture Overview; Tapper Platform Architecture Description; Subprocessors. Questions: security@tapper.ai.

6. As built

The diagrams below show the integration as built on Tapper's production platform, checked on 12 October 2026: components and data stores, then the call sequence.




Figure 3. Components and data stores as built. Solid lines: built and running in production. Dashed lines: built, not yet run for a CRM outcome.




Figure 4. Connect, sync, value and upload sequence as built. Solid: built and running in production. Dashed: built, not yet run for a CRM outcome. Red: not built.