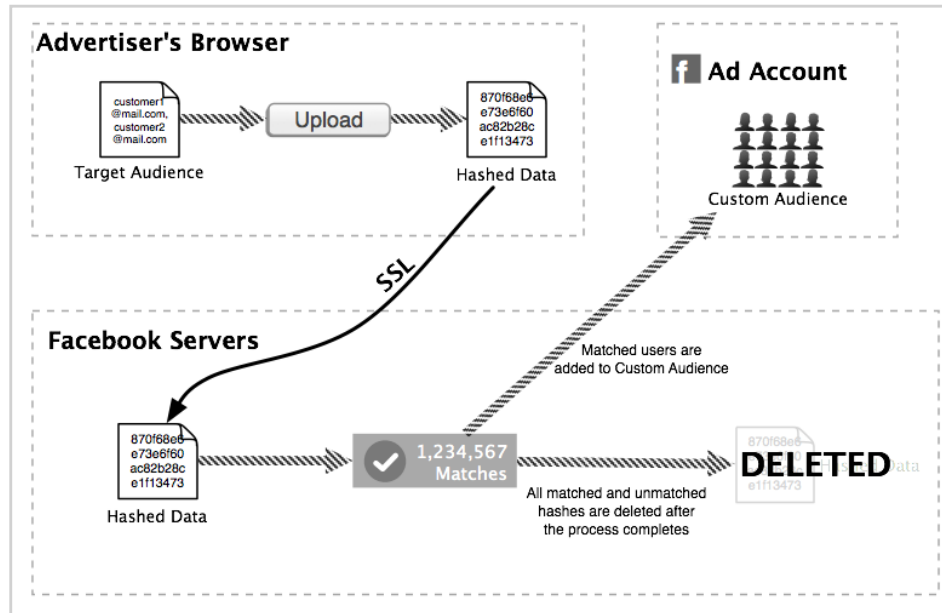


Custom Audiences: Data Security Overview

This document is designed to help address the most common questions that data security professionals ask about custom audiences. It was prepared by the Facebook information security team.

The Custom Audiences match process



- Advertiser uploads into their browser a list of email addresses or phone numbers belonging to individuals that they want to target with ads.
- Advertiser's browser hashes all of the uploaded email addresses/phone numbers locally on their computer. This process is not designed to send the plain-text email addresses or phone numbers to Facebook.
- Advertiser's browser then connects over SSL to the customer's Facebook ad account, authenticates using their Facebook account credentials, and then passes the list of hashed values to our ads API.
- On the Facebook side, we have pre-computed the hashed values for every Facebook user. We take the customer's list of hashed values and compare it with our own list of hashed values.
- For matching hashes, we add the Facebook users to a Custom Audience stored within the customer's ad account. If a hash does not match, we simply ignore it.
- Once the matching process completes, we delete all of the hashes – both matching and non-matching.
- The end result is that the customer ends up with a "custom audience" that they can target with ads. This Custom Audience is stored in the customer's ad account – only authorized account admins can target it. The customer cannot see the specific individuals who are contained in this Custom Audience, they just see the approximate number of people that this audience contains.

Frequently Asked Questions

What data is shared with Facebook?

You create a list of individuals that you want to target with ads on Facebook. For each individual, a unique non-reversible fingerprint (hash) is created based on the identifier(s) provided (e.g. email address or phone number). Only the hashes are sent to Facebook to allow us to identify matching Facebook users to populate your custom audience.

What is hashing?

A ‘hashing algorithm’ is a one-way mathematical function that creates a non-reversible fingerprint of a fixed length for a piece of text. If the same hashing algorithm is used against a piece of text, it will always produce the same “message digest” - basically a fingerprint of the original data. No mathematical function or “key” exists to reverse this fingerprint back to the original value.

So, for example, if two computers both hash ‘email.address@fb.com’, both will end up with a fingerprint like ‘7f759cccof730cbdb4e297010b8ec5e5’. However, given the fingerprint ‘7f759cccof730cbdb4e297010b8ec5e5’, there is no mathematical function to reverse that back to ‘email.address@fb.com’.

It’s important to note that the fingerprint does not actually contain the original data - it’s just a fingerprint. It’s possible to hash all of the text in Tolstoy’s ‘War and Peace’ to the value ‘07fd2dd465a882756e89f330d5880118’, but clearly that value is not reversible back to the full text of War and Peace.

Facebook uses an industry standard secure hashing algorithm called ‘SHA-256’, which is commonly used in technologies such as online banking.

Why is the data hashed before it is sent to Facebook?

Before any email addresses or phone numbers from your list are transmitted to Facebook, they are ‘hashed’ locally on your computer, which means that Facebook only receives the hashed values and not the original list. This is done to ensure that we only learn which Facebook users you want to target with ads – if a customer in your list is not a Facebook user, we will not find a match of their hashed value to a Facebook user.

How is the data transmitted to Facebook?

The hashed data is transmitted to Facebook using TLS (Transport Layer Security), which means that it is encrypted during transmission using the same technologies that are used for online banking.

How does the matching process work?

Facebook pre-computes the ‘hash’ values for every Facebook user. When your data is sent to our servers, your list of hashes is compared with our pre-computed hashes. If a matching hash is found, that Facebook user is placed into a custom audience, which is stored in your ads account. If no match is found for a given hash, we simply ignore it. Facebook knows that you want to target the matched users with ads, but does not know the relationship between you and these matched users (for example whether they are prospects or customers).

What happens with my hashed data once the matching process completes?

After the matching process, all of the matching and non-matching hashes are deleted from Facebook’s servers. The matching process can take up to two days to complete. No further processing of the hashed values is performed beyond the match process.

Where is my Custom Audience stored?

Your Custom Audiences are stored in your ads account. The infrastructure that physically hosts your ads account is located in Facebook's production data centers, located in Oregon and North Carolina in the US, and Lulea in Sweden.

What information can I learn about the members of my Custom Audience?

For privacy reasons, the list of individual matched users is not accessible to the advertiser. In the Power Editor interface, you will see only an approximate number of matched users. We apply a level of approximation to the numbers to prevent individual Facebook users from being identified via statistical sampling.

Can other advertisers access my Custom Audiences?

No. Only your ad account has access to your Custom Audiences. Your Custom Audience is not targetable by other advertisers without your explicit consent. Facebook does not share personally identifiable information with advertisers.

Are Facebook's datacenters secure?

Facebook's advertising platform is delivered from our production datacenters in Oregon, North Carolina and Sweden. We invest heavily in technologies, processes and teams of dedicated security staff to ensure the security of our production datacenters. The information stored for each advertiser's Custom Audiences is a list of the Facebook issued IDs for our users. These IDs allow us to deliver the ads you request when you use one of your Custom Audiences.

Is my Facebook advertising account secure?

Facebook invests heavily in technologies to protect our customers' accounts. For added security on your advertising account, we offer free, optional security technologies such as login approvals, which is a 'two factor authentication' solution similar to that used by many online banks.

What changes are made to Facebook user accounts who are added to a Custom Audience?

No data (e.g. flag) is added to individual Facebook user accounts as a result of being added to a Custom Audience.

Has the security of the Custom Audiences product been audited by a third party?

Yes. The major components of the Custom Audiences process have been audited by a third party audit firm. The details of the confidential audit are available upon request. In addition to the Custom Audiences audit, our security and privacy practices are audited regularly by privacy regulators in the US and EU to ensure that we are providing industry leading levels of protection. We regularly audit our infrastructure for application and network vulnerabilities, and utilize a bug bounty program for encouraging responsible disclosure of security issues from security researchers.

Does Facebook have a security incident response process?

Facebook has a detailed process for handling any security incidents, which includes a complete internal review of any incidents.