

Detailed Use Case Description: Brief One User Query Redaction

Use Case

Brief One helps investigators and prosecutors find relevant evidence inside a case faster by supporting natural-language search across images, documents, and transcript-backed audio and video. This use case is intended to help Axon understand what users are trying to find with AI Search and whether the feature is helping them get to meaningful evidence efficiently, while minimizing privacy risk. For eligible customer tenants, Axon may review a privacy-protected version of search activity in which the original search query is sanitized before review. The original query is not stored in raw form.

Customer Benefit

This use case helps Axon improve Brief One by identifying common search patterns, ranking gaps, and areas where the product may not yet return the most useful evidence quickly enough. It supports product improvement by showing, in a privacy-protected way, what kinds of evidence users are seeking and whether search results are helping them find that evidence with low friction.

Data Accessed

For eligible use of this feature, Axon may access the user's search query and limited context about the search results. This may include whether the user clicked a result or abandoned the search, the rank of returned results, the type of result returned, broad category labels associated with the result, a confidence value, and the date of the search.

How Axon Uses the Data

Axon uses this information to understand how Brief One is being used, what kinds of evidence users are seeking, and where the search experience succeeds or falls short. Axon personnel may review sanitized query records and associated search-outcome context to improve search quality, identify missing coverage, and prioritize product improvements. Axon personnel will not review evidence files (audio, video, transcripts, etc.) suggested in searches. This use case is for product improvement and analytics only. It does not change the customer's live search experience.

Privacy-Preserving Technique

Before a query is reviewed, personally identifiable information is removed and replaced with synthetic stand-ins so the query remains understandable without preserving the original identifiers. For example, a real person's name may be replaced with "John Doe" or "Jane Doe," a real address may be replaced with a synthetic address, and real plate numbers, VINs, case numbers, badge numbers, phone numbers, and email addresses may be replaced with synthetically valid-looking placeholders. Raw user queries are never stored. Organization and user identifiers are pseudonymized before review.

Data Schema

The reviewed data is as follows:

- Organization ID: a pseudonymized identifier for the customer organization.
- User ID: a pseudonymized identifier for the user.
- Sanitized Query: the search query after personal information has been removed and replaced with synthetic stand-ins.
- Search Outcome: whether the user clicked a result, and if so which one, or whether the search was abandoned.
- Result Rank: the position of a returned result in the list.
- Result Type: the kind of item returned, such as image, document, audio transcript, video with transcript, or transcript moment.
- Result Category: broad categories such as People, Weapons, Vehicles, Locations, Objects, Behaviors, Conversation, Legal Milestone, or Time and Events.
- Confidence: a system confidence value associated with the result.
- Date: the calendar date of the search.

Examples of What May Be Reviewed

A reviewed record may include a sanitized query such as “Show video of the defendant holding a firearm” together with information about the ranked results returned, the type of each result, broad categories such as Weapons or Behaviors & Actions, whether the user clicked a result, and a confidence value. It may also include sanitized queries such as a search for a blue sedan with a damaged rear bumper or a search for video outside an apartment building entrance, along with similar search-outcome context.

Examples of What Is Not Retained in Original Form

This use case is not intended to retain raw user queries. It is also not intended to retain real names, real addresses, real phone numbers, real email addresses, real dates of birth, real case numbers, real license plates, real VINs, or real badge or unit identifiers in the reviewed query text. A reviewed record will never contain raw video, audio, transcripts, or other data that may potentially be personally identifiable.

Retention Period

Axon retains only the sanitized, extracted information for a limited period no longer than necessary to support product improvement for this use case. Raw user queries are not stored. This mirrors the retention approach used in similar customer-facing use case descriptions, where Axon keeps only the minimum derivative data needed for product improvement and does not retain the original sensitive content in raw form.

Preservation of Original Content

This use case does not modify the customer's original evidence, source records, or live product behavior. It is analysis-only and is intended to support product improvement without changing production routing, responses, or the user experience. Original evidence in customer systems remains unchanged.

Withdrawal

Customers participating in this use case may request withdrawal. After withdrawal, Axon would stop collecting new data for this use case from that customer. Previously generated de-identified or sanitized derivative insights may continue to be retained where permitted and where they no longer identify the customer, user, or original query.

Summary

This use case helps Axon improve Brief One by reviewing privacy-protected search activity from eligible customers. Before review, search queries are sanitized so that personal information is removed and replaced with synthetic stand-ins. Axon reviews only the sanitized query text together with limited search-outcome context, using pseudonymized identifiers and without storing raw user queries.