

## Detailed Use Case Description: AI People Search User Query Redaction

### Use Case

AI People Search helps investigators and records staff find relevant people and related information about them inside Axon Records (RMS) faster by supporting natural-language search over structured records data — for example, “tall white male seen near the park” or “is this person connected to anyone named John.” The system interprets the user’s plain-language request, translates it into structured search filters, returns ranked record results, and can generate a short AI summary of what was found per result. This use case is intended to help Axon understand what users are trying to find with AI People Search and whether the feature is helping them reach meaningful records efficiently.

### Customer Benefit

This use case helps Axon improve AI People Search by identifying common search patterns, where the system misinterprets a user’s intent, ranking gaps, and areas where searches are abandoned or fail to surface the right records quickly enough. It supports product improvement by showing, in a privacy-protected way, what kinds of people and records users are seeking, whether the AI correctly interpreted their query, and whether results are helping them find that information with low friction — leading to a more accurate and helpful search experience over time without exposing any agency’s sensitive information.

### Data Accessed

For eligible use of this feature, Axon may access the user’s search query and limited context about the search results. This includes a one-way hashed (pseudonymized) organization identifier, a pseudonymized user or session identifier, the sanitized query text, the sanitized AI interpretation of the query (the structured filters the system derived), whether the user clicked a result or abandoned the search, the rank of returned results, the type and broad category of result returned, a system confidence value, any thumbs-up/thumbs-down feedback the user gave on the AI interpretation or summary, and the date of the search.

### Data Schema

The reviewed data is as follows:

- **Organization ID:** a one-way hashed, non-reversible identifier for the customer organization. It allows Axon to tell that a set of searches came from a single agency, but not which agency it is.
- **User / Session ID:** a pseudonymized identifier for the user or search session.
- **Sanitized Query:** the natural-language search query after PII and CJIS have been removed and replaced with synthetic stand-ins.
- **AI Interpretation:** the structured filters the system derived from the query (for example, an age or height range, sex, or relationship type), after sanitization.
- **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, using a 1-based index.

- **Result Type:** the kind of record returned, such as a person record, association, vehicle, location, or incident.
- **Result Category:** broad categories such as People, Associations, Vehicles, Locations, or Incidents & Events.
- **Confidence:** a system confidence value associated with the result.
- **User Feedback:** any thumbs-up/thumbs-down rating the user gave on the AI interpretation or AI summary.
- **Date:** the calendar date of the search.

### How Axon Uses the Data

Axon uses this information to understand how AI People Search is being used, what kinds of records-based evidence users are seeking, how accurately the system interprets natural-language queries, 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, correct query interpretation, identify missing coverage, and prioritize product improvements. Axon personnel will not perform eyes-on review of the underlying RMS records, person profiles, associations, or other source content returned by a search. 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, both personally identifiable information (PII) and criminal justice information (CJIS) are 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” or “Person name” a real address may be replaced with a synthetic address, and real plate numbers, VINs, case or report numbers, badge numbers, phone numbers, and email addresses may be replaced with synthetically valid-looking placeholders. Objective, non-identifying descriptors are preserved so the query remains useful for human review. Raw user queries are never stored.

### Examples of What May Be Reviewed

A reviewed record may include a sanitized query together with information about the ranked results returned, the type and broad category of each result, the AI’s interpretation of the query, whether the user clicked a result, a confidence value, and any feedback rating. The table below illustrates how an original user query is transformed into the sanitized record that Axon would review.

Original User Query (never stored)	Sanitized Query Reviewed by Axon
“Is Amanda Gard-Font connected to a John?”	“Is [person name 1] connected to a person named [person name 2]?”
“Tall white male linked to 123 Oak Street”	“Tall white male linked to [synthetic address]” (interpreted as height range + sex filter)

“Owner of plate ABC-1234 from report #2026-00871”	“Owner of plate [synthetic plate] from report [synthetic case number]”
“People associated with James MacArthur, DOB 4/12/1990”	“People associated with John Doe, DOB [synthetic date]”

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 or report numbers, real incident or call identifiers, real license plates, real VINs, or real badge or unit identifiers in the reviewed query text. A reviewed record will never contain raw RMS records, person profiles, associations, or other source content, and Axon will never perform eyes-on review of the records or evidence suggested by search queries.

### **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, and sanitized data is held in an ACEIP-approved platform. 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 records, source data, or live product behavior. It is analysis-only and is intended to support product improvement without changing production routing, model responses, search results, or the user experience. Original records in customer systems remain unchanged, and the original query is used only briefly to generate the sanitized summary before being discarded.

**Can I get more information about what Axon is doing and why?** Absolutely! Please write us at [aceip@axon.com](mailto:aceip@axon.com) and we'd be happy to answer your questions.

**Am I able to withdraw my agency from this use case and from ACEIP altogether?** If at any time you'd like to withdraw, please write us at [aceip@axon.com](mailto:aceip@axon.com).