

Detailed Use Case Description - Supervisor Feedback Loop

Use Case

Axon's AI Behavioral Signals feature helps supervisors review body-worn camera video by flagging moments in the audio transcript that may be worth attention - for example, signals of professionalism, compassion, or escalation. When a supervisor reviews a video, they can mark each AI-flagged moment, confirm or reject a detection, or suggest a signal the AI missed, and optionally add a short reason for their rating.

Customer Benefit

The value of AI Behavioral Signals depends on accuracy — noisy or wrong signals erode supervisors' trust and waste review time. Today Axon has no structured, at-scale way to know which signals the model gets right versus wrong in real reviews; customer feedback and usage analytics alone do not reveal per-signal accuracy. Using supervisors' de-identified accept / reject / suggest feedback, Axon can find the signals the model gets wrong most often and refine how it detects them — so supervisors get more reliable, less noisy signals and spend less time on false flags.

Data Accessed

For eligible reviews, Axon's automated systems access the following within the customer environment. No Axon individual reviews transcript text, listens to audio, or views video as part of this use case:

- The review's behavioral-signal annotations — the signal type, the supervisor's verdict (Confirmed / Rejected — binary, no "not sure") or suggested signal, the AI's 1–10 confidence, and the transcript turn position number (a number only, never the words).
- Identifiers needed to locate the related evidence (organization and evidence identifiers), which are one-way hashed before anything is exported.
- The call type (from dispatch / CAD, e.g. "traffic stop").
- The AI model version that produced the detection.
- The supervisor's reason for their rating — a preset reason or a short typed note — which is automatically repopulated to replace personal information with synthetic stand-ins before it is stored.

How Axon Uses the Data

When a supervisor completes a review, an automated process runs entirely inside the customer environment. At every stage, customer content is handled only by automation; Axon personnel do not view or listen to it:

- Selecting eligible reviews. Only reviews that contain supervisor feedback are considered. Evidence marked restricted/confidential, or tagged as involving juveniles or CSAM/sensitive categories, is discarded and never exported.
- Building the de-identified record. For each piece of feedback, automation assembles a small structured record, looks up the call type, repopulates the reason text, and one-way hashes the identifiers.
- Export for evaluation. The de-identified record is written to Axon's AI evaluation environment. Data scientists review only aggregate, de-identified feedback to measure per-signal accuracy and improve the model. They do not see identifiers, transcript text, audio, or video.

Privacy-Preserving Technique

- Eyes-off / ears-off. No Axon individual views customer transcript text, listens to audio, or views video for this use case. Only automation processes the review.
- One-way hashing of identifiers. Organization and evidence identifiers are hashed before the data leaves the review system, using a key held separately by Axon's AI platform (not the Performance product). The hash cannot be reversed.
- Automated repopulation of the reason. The supervisor's typed reason is passed through an automated repopulation step that detects personal information — names, addresses, license plates, badge numbers, dates of birth, phone numbers, case numbers, and similar — and replaces it with a plausible synthetic stand-in before storage. The original text is never retained. Repopulation effectiveness is verified by an offline PII-leakage evaluation.
- Exclusion of sensitive evidence. Restricted/confidential evidence and juvenile/CSAM-tagged evidence are excluded before export.
- No transcript content. Only a transcript turn position number is used — never the spoken words.
- Coarsened timestamps (reduced to the hour) and aggregate use only. The result is de-identified, not linked to any individual, and used only in aggregate.

Data Schema

Only the de-identified export (right column) leaves the review system. The internal record (left column) exists only inside the customer environment and is not exported in that form.

Internal record (inside the customer environment)	De-identified export (leaves to AI evaluation)
organization_id: 550e8400-e29b-41d4-a716-446655440000	organization_id: one-way hash
evidence_id: 123e4567-e89b-12d3-a456-426614174000	evidence_id: one-way hash
feedback_id: opaque unique id (UUID) — used for deduplication only; carries no content meaning	same
model_version: prva-2026.3	same
reviewed_at: 2026-08-11T14:37:02Z	reviewed_at: 2026-08-11T14:00 (coarsened to hour)
annotation_type: AccuracyFeedback / HumanSuggestedSignal	same
signal_type: Compassionate (one of: Escalatory, Compassionate, Heightened, Grateful — EA-ready signals; Respectful, Dismissive, Resistant in development)	same
turn_index: 7 (position number, no text)	same
verdict: Rejected (one of: Confirmed, Rejected — binary, null = Not reviewed; for AccuracyFeedback type; absent for HumanSuggestedSignal type)	same
confidence: 4 (model 1–10)	same
call_type: Traffic Stop	same
reason (raw): "Officer Bennett wasn't escalating"	reason (repopulated): "Officer Jonathan Harris wasn't escalating"

Examples of What May or May Not Be Extracted

- Reason repopulation: the supervisor's reason is stored and exported only after personal information is automatically replaced with synthetic stand-ins.

Original reason (as typed by the supervisor)	Repopulated reason (stored and exported)
Officer Bennett at 5th & Main wasn't escalating	Officer Daniel Hughes at 302 Birchwood Lane wasn't escalating
This wasn't dismissive — he helped the woman find the shelter	This wasn't dismissive — he helped the woman find the shelter (unchanged; no personal information)
Suspect John Doe, DOB 08/03/1992, was the aggressor	Suspect Anthony Ramirez, DOB June 14, 1987, was the aggressor
Preset reason: Tone did not match the situation	Tone did not match the situation (unchanged; no personal information)

- Evidence eligibility: Sensitive evidence is excluded before any data is exported:

Evidence in the review	Included in this use case?
Standard evidence, no sensitive tags, with supervisor feedback	Yes — de-identified feedback is exported
Evidence marked Restricted or Confidential	No — excluded before export
Evidence tagged juvenile / CSAM / sensitive category	No — excluded before export
Review with no supervisor feedback	No — only reviews with feedback are exported

Preservation of Original Content

This use case does not modify the customer's original evidence, transcripts, review annotations, or live product behavior. Automated systems read only a de-identified derivative of a completed review; the original data in customer systems remains unchanged.

Withdrawal

Customers participating in this use case may request withdrawal at any time by emailing aceip@axon.com.