FAQ - Detailed Explanation of Use Case - Redaction Assistant Model Training

1. Use Case: Redaction Assistant Model Training

What is this use case?

This use case allows Axon to improve Redaction Assistant AI models by learning from how users interact with redaction masks in customer videos. Redaction Assistant models detect privacy-sensitive objects such as heads, screens, license plates, notebooks, cell phones, and ID cards so customers spend less time manually redacting video. In this use case, Axon uses user-created, user-modified, and user-deleted masks as training signal to improve the quality of those models over time.

2. What are Redaction Assistant masks and user edits?

Redaction Assistant masks are the boxes placed over objects in video that need to be redacted. Those masks may be generated by the product or created manually by a user. Users can leave an AI-generated mask unchanged, modify it, delete it, or create a new mask themselves. Those actions provide signal about where the model performed well and where it missed, overreached, or drew an imprecise box.

3. What data will Axon access?

For this use case, Axon will access customer videos together with the associated redaction mask information needed to train and evaluate the object detection models. This includes the set of masks that users created, updated, or deleted, and the associated frame-level information needed to interpret those masks. The data processed for model training and evaluation may include:

- frame number
- bounding box coordinates
- an internal bounding box ID
- model confidence score
- whether a box was AI-generated and left as-is, AI-generated and modified, AI-generated and deleted, or user-created
- the object type, where known, such as notebook or screen
- top-level evaluation metrics such as precision and recall

4. How will Axon use your data?

Axon will use this data to train and evaluate discriminative Redaction Assistant AI models within the customer boundary. The goal is to improve detection accuracy for privacy-sensitive objects and better align the models with real customer usage.

This work supports several concrete outcomes:

a. Improve model accuracy.

User edits help teach the model where its existing detections are incomplete, inaccurate, or too loosely or tightly drawn.

b. Cover blind spots in existing training data.

Customer environments contain real-world variation that may not be fully represented in Axon's existing datasets. For example, certain notebooks, screens, or other object appearances may differ materially from what the model has previously seen.

c. Measure product performance on real customer workflows.

The training and evaluation process helps the team understand where the model deviates most from customer expectations so improvement efforts can be targeted.

5. What is the customer benefit?

The customer benefit is a more accurate Redaction Assistant. Better models reduce the amount of manual redaction required, improve the tightness and completeness of redaction boxes, reduce missed objects, and help Axon prioritize improvements on the object types and scenarios that matter most in real customer workflows.

6. How much data and for how long?

This use case is expected to rely on ACEIP-participating customer environments and on the minimum set of videos and associated mask-edit data needed to support training and evaluation. Access is expected to be ongoing and may occur daily as the team iterates on the models. Extracted or generated data will be retained according to business need and applicable retention controls for the project.

7. What Privacy Preserving Technique will be used and is this eyes-off processing?

The privacy-preserving design for this use case is that the resulting model is a discriminative object detection model, not a generative model. The model learns boundaries and detection patterns for objects, rather than learning to recreate customer content. The primary outputs used outside the immediate training context are privacy-preserving model artifacts and structured detection information such as bounding box coordinates, confidence values, categorical source labels, and top-level evaluation metrics.

Axon will avoid logging PII or case-specific details and will limit outputs to the statistical patterns and content characteristics necessary for product improvement. The extracted content will not be linked back to the customer. Access and processing are limited to the model training and evaluation use case within the customer boundary.

8. Does this processing preserve the original content, and are temporary copies created?

This processing does not modify the original customer videos. The original content remains preserved in place. To perform training and evaluation, temporary working copies or intermediate processing artifacts may be created within the controlled customer-boundary environment, but the intended exported or retained outputs are limited to the trained object detection model, structured bounding box data, and evaluation metrics needed for product improvement.

9. Can I get more information about what Axon is doing and why?

Yes. Customers can contact the ACEIP or Privacy point of contact identified through the program for additional information about the purpose of the use case, the data involved, and the privacy-preserving controls that apply.

10. Am I able to withdraw my agency from this use case and from ACEIP altogether? What will you do with my data if I withdraw after the fact?

If a customer withdraws from ACEIP or this use case, Axon would stop using that customer's data for the use case going forward, subject to the program's applicable withdrawal and retention processes. Original customer content would remain governed by the customer's product and service relationship, while extracted or derived use-case data would be handled in accordance with ACEIP withdrawal, deletion, and retention rules then in effect.

11. Do you have examples of what data Axon may or may not extract for this use case?

Yes. Examples of data Axon may extract or generate for this use case include:

May extract or generate

- frame number
- bounding box ID
- bounding box coordinates
- model confidence score
- source category such as user-created, AI-generated and modified, AI-generated and deleted, or AI-generated and left unchanged
- known object type, such as notebook, screen, head, or plate
- aggregate evaluation metrics such as precision and recall
- trained object detection model artifacts

Will not extract as a reporting output

- transcripts or other potentially revealing raw information
- raw images or videos
- customer-identifying linkage fields

June 11, 2026 Addendum

As of June 11, 2026, Axon is updating the categories of data that may be extracted for this use case to include:

- Failure reasons, such as unsupported category, blurry object, occlusion, dark lighting, etc.

All existing exclusions regarding data that Axon will not extract remain unchanged.