Visual Counter – Comments on EDPB Guidelines 02/2026 on Anonymisation

Oct 21, 2026 · @Fernando Garcia-German Benavides

1. Submitter, declaration of interest and scope

Submitted by: Identificacion y Sensorizacion SL (trading as Visual Counter), Julian Camarillo 53, 28037 Madrid, Spain. Contact: Fernando Garcia-German Benavides, CEO, info@visualcounter.com.

About the submitter: Visual Counter is a European company that designs and supplies people-counting and occupancy sensors for retail stores, transport operators and other publicly accessible spaces. Our sensors use time-of-flight (ToF) technology: they emit infrared light and measure its reflection to build a 3D depth map, with no colour or texture information, which on-board algorithms interpret as people or objects. The sensors are always mounted overhead, looking straight down.

Declaration of interest: Visual Counter competes with providers of camera-based footfall and in-store analytics. We state this openly. Our comments do not ask the EDPB to favour or exclude any technology. They ask that the final Guidelines state clearly how the anonymisation framework applies to one widespread use case, so that controllers and supervisory authorities apply it consistently.

Scope: These comments concern anonymisation carried out in real time on devices that capture images of people in publicly accessible spaces ("edge" or "on-device" anonymisation). Retail footfall analytics is the main example: entry and exit counts, occupancy, heatmaps, paths and dwell time. Paragraph numbers refer to Guidelines 02/2026 on Anonymisation, version 1.0, adopted on 7 July 2026.

2. Executive summary

The draft Guidelines already contain the principle that resolves our concern: anonymisation is itself processing of personal data and needs a legal basis (para 38). However, they do not apply that principle to the case where the source data exists only for milliseconds inside a device. That case is now common. Camera-based analytics in shops, malls and stations is widely marketed as "anonymous" or "GDPR-compliant" on the

ground that frames are processed on the device and never stored or transmitted. Controllers rely on these claims and may skip the obligations that apply at the moment of capture.

We respectfully ask the EDPB to make seven clarifications in the final text:

1. **Capture is processing.** Confirm that capturing images of identifiable people is processing of personal data, however briefly the images exist and wherever they are processed, and that only the output can be assessed for anonymity (section 2.5.2, after para 38; new Example A).
2. **Tracks are records.** Confirm that per-person tracks, temporary track identifiers, trajectories and individual dwell times are record-level data, and that matching a person across cameras, zones or visits by appearance is identification (paras 19, 25, 74, 85; new Example B).
3. **Aggregates in physical spaces.** Address aggregates produced where staff can see the people being counted: small counts, fine time and space resolution, and zones occupied by a single employee (paras 16(c), 30, 76–78, 85; new Example C).
4. **Transparency.** Extend para 40 so that systems capturing identifiable images are not described as "anonymous", "processing no personal data" or "outside the GDPR", and confirm that Article 13 information is due at the moment of capture.
5. **Legitimate interest.** Clarify in para 39 that immediate anonymisation is a mitigating factor in the balancing test, not a substitute for the necessity test.
6. **Verifiability.** State that device capabilities (video output, local storage, remote access, firmware updates) must be part of the assessment, and that vendor assurances alone are not sufficient (paras 30, 34, 35, 41, 90).
7. **Design that never captures.** Distinguish devices that never collect information relating to identifiable people from devices that collect it and then anonymise it, and say who must demonstrate which case applies (consistent with Guidelines 3/2019, section 2.1).

3. Context: on-device anonymisation of camera images

In a typical camera-based retail analytics system, personal data is processed at every stage until the final output, and often in the output too. The usual processing chain is:

1. **Capture.** A camera (often an existing security camera, sometimes a dedicated sensor) records frames of everyone in its field of view: customers, children, passers-by and staff.

2. **Detection.** A model on the device or on a local server locates each person in each frame.
3. **Tracking.** The system assigns each detected person a temporary identifier and follows them across frames, and sometimes across cameras, to measure paths, zone visits and dwell time.
4. **Optional attribute estimation.** Some systems estimate age band, gender, gaze or attention, or extract appearance features to recognise the same person in another camera or on a later visit.
5. **Output.** The device sends counts, heatmaps, paths or dwell statistics to a cloud dashboard. Frames are discarded, usually within milliseconds.

Suppliers commonly present stage 5 as proof that the system "processes no personal data". Under the draft Guidelines, that conclusion cannot follow. A video of a person relates to that person by its content (para 16(a)), and the glossary expressly counts images within a "record". Stages 1 to 4 are therefore processing of personal data. Only the output of stage 5 can be tested against the three criteria.

The issue matters for three reasons:

- **Scale and unavoidability.** These systems analyse every person who enters a publicly accessible space, often without any meaningful choice. Systematic monitoring of a publicly accessible area on a large scale requires a DPIA under Article 35(3)(c) GDPR.
- **Rights are hard to exercise.** The French supervisory authority has observed that people generally cannot object to the analysis of their image when algorithms do not retain images (CNIL position, July 2022). This makes accurate information at the point of capture all the more important.
- **Secondary use of security cameras.** Guidelines 3/2019 on video devices already warn against the reuse of security footage for unexpected purposes such as marketing, and note the spread of intelligent video analysis, including counting algorithms. Analytics software added to existing security cameras is a direct example.

4. Specific comments

Each comment names the paragraphs concerned, the gap we see, and proposed wording. Proposed text is offered as a starting point; we have kept to the drafting style and terminology of the Guidelines.

Comment 1 – Capture and on-device analysis are processing of personal data

Paragraphs: 10, 16(a), 37–38, glossary ("Record").

Gap. Paragraph 10 defines anonymisation as any processing that produces anonymous data, whether or not the original data is deleted. Paragraph 38 then confirms that anonymisation needs a legal basis under Article 6 GDPR. Every example in the Guidelines, however, concerns a dataset that is stored and later tested. None addresses original data that exists only briefly inside a device. In practice, this silence is read as "if nothing is stored, nothing is personal". Article 4(2) GDPR lists collection, use, erasure and destruction as processing, with no minimum duration.

Request. Add a paragraph after paragraph 38, and Example A in section 5 below.

Proposed wording (new paragraph 38a):

> The principles set out in paragraphs 37 and 38 apply regardless of how long the original personal data exists and of where it is processed. Where a device, such as a camera, captures images of identifiable individuals and converts them into information intended to be anonymous, the capture, analysis and deletion of those images are processing operations within the meaning of Article 4(2) GDPR. This remains the case even if the images are held only in volatile memory for a fraction of a second and never leave the device. The controller must therefore have a legal basis for that processing, comply with Articles 12 to 14 GDPR and assess whether a data protection impact assessment is required. Only the output of the device can be assessed for anonymity under these Guidelines. The short lifespan of the original data may be taken into account in the risk assessment and in the balancing test (see paragraph 39), but it does not take the capture outside the scope of the GDPR.

Comment 2 – Per-person tracks, temporary identifiers and re-identification by appearance

Paragraphs: 19, 25, Example 6, 50, 74, Example 15, 85–86.

Gap. To measure paths, zone visits or dwell time, a system must distinguish each person from the others in view and follow them. Under paragraph 19, this is identification: the person is distinguished within a given context in a way that makes it possible to treat them differently. Three points are not yet addressed:

- **Temporary identifiers.** A track identifier works as a pseudonym while it exists, even if it is reset after minutes or hours. It allows the system to act on that person, for example to alert staff to a customer "lingering" near high-value goods or to change digital signage. Where the data is used to treat a specific individual in a particular way, paragraph 50 already makes it a simple case of personal data.
- **Trajectories.** A per-person path through a store, with timestamps, is record-level data of high dimensionality and high resolution (paras 85–86). It is close to the location tracks in Example 15. A path that starts at the staff entrance at 07:58 and ends at the back office reveals an employee; a path between the pharmacy counter and a specific shelf may reveal sensitive information about a customer.
- **Appearance features.** Some systems extract features (body shape, clothing, gait, face) to recognise the same person in another camera or on a later visit, for example to count "returning visitors". These features function as an identifier in the same way as the device fingerprint in Example 6. Where they are used for unique identification, they may also be biometric data subject to Article 9 GDPR, as Guidelines 3/2019 already explain.

Request. Add the wording below after paragraph 25, and Example B in section 5.

Proposed wording (new paragraph 25a):

> Similarly, a system that assigns a temporary identifier to an individual detected by a camera and follows that individual across frames, zones or cameras distinguishes that individual from others in the given context, even if the identifier is reset after a short period and no name is known. The resulting per-person data, such as paths, zone visits or dwell times, is record-level data and should be assessed as such. Where the system relies on appearance features, such as body shape, clothing, gait or facial features, to recognise the same individual in another camera or during a later visit, those features function as an identifier in the same way as a device fingerprint (see Example 6).

Comment 3 – Aggregates produced in physical spaces

Paragraphs: 16(c), 30, 58, 76–78, Examples 16–17, 85, 89.

Gap. Paragraph 58 notes that correctly aggregated data will usually meet the No Record Isolation criterion. Retail analytics outputs are mostly aggregates, so controllers may stop there. In a physical space, however, some relevant entities can see the individuals concerned. Store staff, managers and security personnel observe customers and

colleagues directly (para 30 already lists colleagues and neighbours as possible relevant entities). The controller also often holds security recordings of the same space for several weeks, which is additional information within the meaning of paragraph 29(c). Three situations follow:

- **Small counts at fine resolution.** A heatmap or zone count by minute in a low-traffic area often shows a count of one. Combined with a staff member's observation or a security recording, it allows a specific and meaningful inference about that customer, such as how long they stayed at a particular display (paras 76–78).
- **Single-employee zones.** Counts and dwell times for a till, fitting room entrance or service desk staffed by one employee relate to that employee by reason of their effect (para 16(c), taxi example). They can be used to assess performance or breaks.
- **Outliers.** Unusual values, such as one visitor in the store before opening time, are easy to attribute (para 76).

Request. Add the wording below at the end of section 3.4.3 or in section 3.5.2, and Example C in section 5. We also suggest that the EDPB indicate the safeguards it expects in this setting: minimum counts before publication, coarser time buckets in low-traffic zones, suppression of small cells and exclusion of staff-only areas.

Proposed wording (new paragraph 78a):

> Aggregate data produced in a physical space should be assessed bearing in mind that some relevant entities, such as staff, managers or security personnel, can directly observe the individuals concerned, and that the controller may hold other recordings of the same space. In this context, low counts at a fine temporal or spatial resolution may allow specific and meaningful inferences about an individual customer or employee. Where a zone is regularly occupied by a single employee, counts and dwell times for that zone may relate to that employee by reason of their effect (see paragraph 16(c)).

Comment 4 – Transparency at the moment of capture

Paragraphs: 40–41; Articles 12–14 GDPR.

Gap. Paragraph 40 rightly says that controllers should not call data "anonymous" if individuals are still identifiable, and should explain that personal data is processed to produce anonymous data. In this market, the claims usually start with suppliers. Controllers then repeat them in their records of processing, privacy notices and signage. In practice we see two patterns: signs that mention only "video surveillance" while the

same cameras feed analytics, and notices stating that counting is "anonymous" and involves "no personal data". In both cases individuals are not told, at the time of collection, that their image is analysed for commercial purposes.

Request. Add the wording below at the end of paragraph 40.

Proposed wording (addition to paragraph 40):

> This applies equally where personal data is processed only briefly to produce anonymous data, for example where images captured by a camera are analysed on the device and then discarded. In such cases, controllers should inform individuals at the time of capture, which usually means before they enter the area covered by the device, of the purpose of the analysis, in addition to any other purpose of the same device, such as security. Controllers should not describe such systems as "anonymous", as "not processing personal data" or as falling outside the GDPR, and should not rely on such descriptions by suppliers without verifying them (see paragraph 41).

Comment 5 – Immediate anonymisation and the necessity test

Paragraphs: 39; EDPB Guidelines 1/2024 on legitimate interest (cited in footnote 34); Articles 5(1)(c), 6(1)(f) and 25 GDPR.

Gap. Paragraph 39 treats immediate anonymisation and deletion as a positive factor in the legitimate interest balancing test. We agree. The risk is that it is read as making image-based analytics lawful in itself. Under Article 6(1)(f) GDPR, the necessity of the processing is assessed before the balancing. Processing is not necessary if the purpose can reasonably be achieved by less intrusive means. For counting visitors and measuring occupancy, means that do not capture identifiable images are available on the market.

Request. Add the wording below at the end of paragraph 39.

Proposed wording (addition to paragraph 39):

> However, immediate anonymisation does not replace the assessment of whether the preceding processing is necessary. Where the purpose pursued, such as counting visitors or measuring occupancy, can reasonably be achieved by means that do not involve the collection of personal data, or that involve less of it, the capture of identifiable images may not meet the necessity requirement of Article 6(1)(f) GDPR, even if the images are anonymised immediately. Controllers should also consider such means as part of their obligations under Article 25 GDPR.

Comment 6 – Verifying what a device actually does

Paragraphs: 15, 30, 34–35, 41–42, 89–90.

Gap. An on-device anonymisation claim depends entirely on how the device behaves. Many analytics functions run on standard IP cameras or video recorders that can also stream and record video. The same camera often serves security and analytics at once. Installation and calibration interfaces commonly display live, identifiable images. Firmware can be updated, sometimes remotely by the supplier, and can change what is stored or sent. The Guidelines already provide the elements needed: rogue employees and cybercriminals as relevant entities (para 30), contractual prohibitions as a complement only (para 34), reassessment over time (para 35), documentation of testing (para 41), and the limits of access controls (para 90). Paragraph 15 also confirms that a supplier processing on the controller's behalf takes the controller's perspective. These elements are not yet brought together for device-based anonymisation.

Request. Add the wording below after paragraph 41.

Proposed wording (new paragraph 41a):

> Where anonymity depends on the behaviour of a device, such as a camera configured to discard images after analysis, the documentation should demonstrate that behaviour, not only describe it. Relevant factors include: whether the device can display, transmit or store identifiable images, including during installation or calibration; whether the same device also supplies images to another system, such as a security system; who can change its configuration or software, including remotely, and whether such changes are logged; and how the controller verifies the behaviour of the device over time, including after software updates. Contractual assurances from the supplier may complement, but cannot replace, such technical verification (see paragraph 34).

Comment 7 – Devices designed never to collect identifiable information

Paragraphs: 1, 6, 10, 29, 51; EDPB Guidelines 3/2019, section 2.1.

Gap. Guidelines 3/2019 already recognise that the GDPR does not apply where a camera is built or adjusted so that it collects no information relating to an identifiable person. The draft Guidelines on anonymisation do not address the boundary between that situation and the one where personal data is collected and then anonymised. The difference is decisive for controllers: in the first case the obligations described in Comments 1, 4 and 5 do not arise; in the second they do. Both suppliers and controllers need a test to tell them apart, and that test should apply to every technology, including ours.

Request. Add the wording below to section 2.1 or section 3.3.

Proposed wording (new paragraph 10a):

> These Guidelines concern information that is personal data before it is anonymised. This should be distinguished from situations where a device is designed so that it never collects information relating to an identifiable individual, for example because the resolution or the nature of the signal it captures does not allow individuals to be distinguished with means reasonably likely to be used (see also EDPB Guidelines 3/2019, section 2.1). In such cases, the GDPR may not apply to the information collected by the device. Whether this is the case should be assessed on the raw signal captured by the device, using the criteria and factors set out in these Guidelines, and taking into account whether the device can be reconfigured to collect identifiable information. The entity claiming that a device collects no personal data should be able to demonstrate it.

5. Proposed new examples

The three examples below follow the format of the existing examples and could be numbered accordingly. Each applies paragraphs already in the draft; none creates a new legal standard.

Example A – Footfall analytics on existing security cameras (proposed after paragraph 38a)

> **Example A:** A clothing retailer installs analytics software on the video recorder of its existing security cameras. The software detects people in each frame, counts entries and exits, and builds a heatmap of the sales floor. Frames are analysed in the recorder's memory and discarded within one second. Only hourly counts per zone are sent to the supplier's cloud dashboard. The sign at the entrance mentions video surveillance for security. The retailer's privacy notice states that the analytics are "anonymous" and involve no personal data.
>
> The images captured by the cameras relate to the individuals shown in them by reason of their content (paragraph 16(a)), and those individuals are identifiable. Capturing and analysing the images to produce counts and heatmaps is therefore processing of personal data by the retailer, as controller, for a purpose distinct from security. The supplier, which runs the software on the retailer's behalf, should assess the data from the retailer's perspective (paragraph 15). The fact that the images exist only for a fraction of a second does not change this. The retailer needs a legal basis for the analytics, must inform individuals of this purpose before they enter the monitored area, and should assess whether a data protection impact assessment is required. The hourly counts per zone may be anonymous if they pass the three criteria, taking into account the considerations in paragraph 78a. Describing the system as "anonymous" without further explanation is, however, misleading (paragraph 40).

Example B – Dwell time and returning visitors in a shopping centre (proposed after paragraph 25a)

> **Example B:** A shopping centre uses ceiling cameras to measure how long visitors stay in each area and how many come back within 30 days. The system assigns each detected person a temporary identifier and extracts appearance features, such as clothing colour, body shape and gait, to follow that person from one camera to the next and to recognise them on a later visit. The supplier receives one record per visit, without the temporary identifier, listing each zone entered with entry and exit times.
>
> The temporary identifier and the appearance features allow the system to distinguish each visitor from all others in the centre and to treat them differently, for example by counting them as "returning". The visitors are therefore identified within that context (paragraph 19), and the appearance features function as an identifier in the same way as the device fingerprint in Example 6. Where they are used to uniquely identify a person, they may also be biometric data subject to Article 9 GDPR. The per-visit records sent to the supplier are record-level data of high dimensionality and high resolution (paragraph 85). Many records will be unique, failing the No Record Isolation criterion. Some will allow specific and meaningful inferences, such as a path that starts at the staff entrance before opening time and reveals an employee, failing the No Inference criterion. The shopping centre also holds security recordings of the same corridors, which can be linked to the records by time and place (paragraph 29(c)). The records are therefore not anonymous for the shopping centre, nor for the supplier processing them on its behalf (paragraph 15).

Example C – Fine-grained aggregates in a supermarket (proposed after paragraph 78a)

> **Example C:** A supermarket receives, for each zone of the store, the number of people present and their average dwell time for every minute of the day. No individual records are kept. The data shows, among other things, that one person stayed for twelve minutes in front of a low-traffic shelf of pharmacy products, and that the area behind the service counter, staffed by one employee per shift, was empty for 35 minutes during a particular afternoon.
>
> The data is aggregated and meets the No Record Isolation criterion. However, the store manager, who is on the sales floor and has access to the security recordings, can tell who the person at the pharmacy shelf was. The minute-by-minute counts then allow a specific and meaningful inference about that customer, failing the No Inference criterion (paragraphs 76–78). The counts and dwell times for the service counter relate to the employee on shift by reason of their effect, since they can be used to assess that employee's breaks (paragraph 16(c)). The data is therefore not anonymous for the supermarket. After reviewing the results, the supermarket decides to receive hourly figures only, to suppress any figure below a minimum count, and to exclude staff-only areas, before testing the anonymity of the data again.

6. Conclusion

The draft Guidelines give controllers a sound framework for testing whether data is anonymous. Our seven requests ask the EDPB to apply that framework explicitly to devices that anonymise images in real time, so that the framework is not read as exempting the capture itself. The proposed wording relies only on principles and paragraphs already in the draft.

Visual Counter would welcome the opportunity to contribute technical input, for example at a stakeholder event, and is available to answer any questions on these comments.

Madrid, 21 October 2026

Fernando Garcia-German Benavides, CEO, Visual Counter