Bounded Observation

A Privacy Specification for Environment-Observation Tags in the Open 6G Device Ecosystem

Companion specification to: *Density as an Availability Condition for Environment-Dependent Trust Roots (P-DENS, P-OBS)* **Contribution to IEEE Industry Connections activity IC25-009-01**

Version 2.0 | July 2026 | For IC Activity Review | Supersedes all v0.x and v1.x drafts. Incorporates as Annex A the material previously circulated as a separate annex.

Abstract

The P-OBS principle asks devices to publish signed statements about the conditions they measure locally, and derives the incentive for network participation from the fact that such statements cannot be obtained by any other route. The same property that makes the exchange valuable makes it dangerous. A device that publishes what it observes publishes something about people who published nothing, and the underlying channel measurements support identification of individuals with accuracy close to complete.

This specification bounds the publication channel. It defines the harm in two registers: the flattening of the gradient by which a person reads their own body better than the surrounding infrastructure reads it, which is measurable in decibels and requires no adversary model, and the re-identification of individuals from published statements, which requires one. It sets out a threat model centred on the person who is not a participant, and specifies eight requirements, R13 to R20, covering abstraction, a closed predicate list, publication limits, evaporation, unlinkability, opt-out, local processing and auditability.

The requirements apply in a fixed order rather than in parallel. Published evidence shows that limiting how often something is published has almost no effect while each published item remains rich, so the abstraction requirement of R13 is the precondition on which the remaining limits depend. Annex A derives that ordering and specifies the protocol by which the numerical parameters of R15 are to be obtained. Those parameters do not yet exist.

The specification also states plainly what it cannot achieve. It governs publication, not perception; an adversary who ignores the standard is unaffected by it; and two of its eight requirements are attested rather than demonstrable in the sense of R21 of the parent contribution.

Keywords: wireless sensing, privacy, non-participant, re-identification, gradient, observation tags, unlinkability, opt-out, calibration, IEEE IC25-009-01

---

1. Scope and problem statement

The parent contribution introduces environment-observation tags: signed statements about locally measured conditions, published to neighbouring devices without operator mediation, carrying an independent lifetime. Their purpose is to supply the participation incentive that a witnessing economy otherwise lacks, and to lower the barrier to reaching the density at which environment-dependent trust roots become available.

Everything the rest of this series was designed to prevent concerns what an attacker can take from a device. This specification concerns something the series has not previously had to address: what a conformant device gives away about people who are not its owner and have made no choice at all.

The asymmetry is structural. A person can decline to buy a device. A person cannot decline to be in a room measured by someone else's device. The subject of an observation tag is a volume of space, and volumes of space contain people. In an architecture deliberately built without a centre there is additionally no party who can be instructed to switch the channel off after the fact.

This document specifies what a conformant device may publish about the space around it, at what resolution, to whom, and for how long. It does not specify sensing, which is governed by the relevant radio standards, and it does not regulate perception.

1.1 Boundary with Part Eight

Part Eight of the accompanying series concerns a different exposure with a different remedy: the reflection of a person from the radio field itself, which requires no publication and no participation, and whose remedy is the removal from the architecture of the capabilities that make coherent imaging possible.

The two do not cover one another. The mechanism of Part Eight, which is to withhold from the network what the link does not need so that an eavesdropper is denied it too, does not protect against a channel the network is required to emit. This specification bounds that channel. Neither document should be read as covering the other's harm, and clause 5.4 states the interaction that must not be allowed to develop between them.

---

2. Threat model

2.1 The non-participant

The primary subject of this threat model is not the device owner and not an attacker. It is the bystander: a person within the sensing volume of a conformant device who is not its owner, has not consented, and in most cases is unaware that the volume is being measured. Every requirement in clause 3 exists to bound what conformant behaviour can reveal about this person.

This framing is not original to this specification and should not be presented as such. A body of recent work on ISAC governance identifies the same subject and the same gap. One survey states that ISAC extends observation to the surrounding physical environment, that sensing may involve collateral collection concerning bystanders and passive objects, and that personally identifiable information may be collected about individuals who carry no user equipment at all; it lists data

ownership for bystander sensing, permission revocation, third-party auditing and measurable privacy criteria among its open problems. A position paper on unconsented sensing argues that standardisation mitigations remain overwhelmingly technocentric, relying on cryptographic anonymisation and physical-layer security, and that this underestimates the legal and sociotechnical complexity, proposing purpose-bound sensing activation, citizen transparency and algorithmic accountability. A review of human-centric ISAC distinguishes the role of the human as active user, passively sensed subject, or bystander.

What this specification adds to that literature is not the diagnosis but a bounded publication channel with testable limits: a closed predicate list that excludes the interior of rooms, a quantitative ceiling on the content of a single statement, a budget attached to a volume rather than to a publisher, and a calibration protocol that produces numbers a conformance test can check. The governance work asks for measurable privacy criteria for system design; clause 3 and Annex A are an attempt at one.

2.2 Two definitions of harm

Gradient flattening. A person is at zero distance from their own body and is therefore always its closest observer. Under the two-way propagation law that governs a signal travelling out to a body and back, the advantage this confers is a function of distance and is measurable. The violation, in this register, is anything that erases the difference between how well a person can read themselves and how well the infrastructure can read them. A network can be in violation with every policy satisfied and no byte disclosed.

This definition is primary here because it requires no assumption about what an adversary knows, and because it is stated in a unit that can be checked.

Re-identification. The second register concerns published statements rather than the field: whether the tags a device emits allow an individual to be singled out or linked to earlier observations of the same individual. This definition requires an adversary model, supplied in Annex A, and governs the numerical parameters of R15.

2.3 Channel measurements are biometric material

The assumption that measurements of a space are anonymous because they contain no names does not survive contact with the evidence.

A study published in 2025 recorded 197 participants walking through a Wi-Fi field and inferred identity from beamforming feedback information, the compressed channel description that clients broadcast to an access point in the clear. Accuracy reached 99.5 per cent, held across four recording perspectives and five walking styles, and was validated against recordings of the empty room to confirm that a biometric trait rather than a session artefact was being learned.

Three features bear directly on this specification.

The adversary is weaker than had been assumed. Beamforming feedback travels unencrypted at the MAC layer, so no modified firmware and no special hardware is required, and a single passive receiver anywhere in range captures every perspective at once. The authors report that while fewer than six per cent of deployed devices support extraction of raw channel-state information, more than half already supported beamforming feedback in 2023.

The compressed representation outperformed raw channel data on roughly a ninth of the total volume, which the authors attribute to compression acting as feature extraction and to higher spatial resolution, concluding that spatial resolution matters more than temporal resolution.

The authors surveyed existing countermeasures and found none applicable in real-world conditions, and called for effective privacy protection to be built into the emerging sensing standard, noting the same concern for the joint communication and sensing approaches planned for 6G. That amendment has since been published, in September 2025, without the protections they asked for.

One limitation should be stated with equal care. The study identifies within a closed set: the adversary requires prior labelled recordings of the target, and the classifier cannot be tested on individuals absent from training. The threat is therefore linkage of a person to earlier recordings of the same person, which breaks anonymity without necessarily attaching a name. That is narrower than open-world recognition of an arbitrary passer-by, and it is sufficient for the harm this specification is written against.

Design consequence. Raw or lightly processed channel data is biometric material. It cannot be treated as environmental telemetry, it cannot be published, and it cannot be shared under a general participation agreement. This is the load-bearing constraint of this specification, and Annex A establishes that the other limits of R15 are conditional on it rather than parallel to it.

2.4 Overhearing

Analyses of the sensing standard distinguish two overhearing surfaces and both apply. In report overhearing a party collects measurement reports not intended for it and infers activity directly. In signal overhearing a party estimates the channel from transmissions it was never meant to interpret, requiring no cooperation from anyone. Candidate countermeasures discussed in that context include encrypting measurement fields so that only trusted devices can recover them, following the approach taken for ranging in an earlier amendment, and physical-layer techniques that degrade the signal available to unintended receivers.

Report overhearing is within reach of this specification, because reports are what P-OBS publishes. Signal overhearing is not, and clause 4 says so.

2.5 Aggregation

Individually harmless tags become harmful in aggregate. Occupancy of one room says little; occupancy recorded every thirty seconds for a month describes a household routine. Tags from a single publisher say little; tags from several publishers at distinct positions localise a subject.

The tension this creates with verification is real but was misdiagnosed in earlier drafts. Redundant agreement between publishers about the same predicate is what verification needs and adds little to the joint information published about a volume. Coverage of the same space from distinct vantage points is what enables triangulation. These are different quantities, and P-OBS-2 of the parent contribution now bounds them separately. The resolution adopted here is neither to bound overlap generally, which is load-bearing for integrity, nor to rely on coarsening, which Annex A shows to be weak, but to cap what a single tag may contain and how many tags concern a volume in a period.

2.6 The absent addressee

A centralised sensing deployment has an owner who can be served with an order to stop. A distributed publication channel across independently owned devices has none. This is normally the architecture's strength; here it is the reason the constraint must be structural. A rule that depends on a party being able to comply after the fact will have no party to comply.

---

3. Requirements

Numbering continues the sequence of this series: R7 to R11 were introduced with the resolution of the P-INDEP open questions, R12 and R21 with the parent contribution. The requirements below are R13 to R20.

R13 (abstraction). Observation tags shall carry interpreted, quantised statements about conditions. Raw or minimally processed channel-state information, beamforming feedback, per-subcarrier amplitude or phase series, and equivalent low-level measurement data shall not be published in any tag and shall not be exported from the device. Interpretation shall occur on the device that performed the measurement.

Rationale: low-level channel data supports identification of individuals; interpreted statements about conditions do not. This mirrors the design decision taken for cooperative perception between vehicles, where messages carry object descriptions rather than sensor output. This requirement takes precedence over the remaining limits of R15 and is verified first in conformance testing, because where each statement is rich, identification saturates within a few tens of statements and no limit on their number takes effect. The quantitative ceiling that gives the requirement force is specified in R15 and derived in Annex A; the word "interpreted" is a description and not by itself a bound.

R14 (predicate allowlist). The standard shall fix a closed list of permitted observation predicates, organised by class. Predicates describing the physical state of the medium and of infrastructure are permitted by default. Predicates describing persons are permitted only in the restricted class defined below. Predicates that identify, characterise or track individuals shall not be admitted at any class. An implementation shall not publish a predicate outside the list.

Class	Examples	Publication rule
E Medium and infrastructure	Precipitation, temperature, humidity, air quality, obstruction of a passage, signal conditions, structural vibration	Permitted by default within the limits of R15
O Aggregate occupancy	Presence or absence of motion in a volume; count banded into coarse ranges	**Not published to neighbouring devices.** Permitted only within a volume under the control of the subject, and only within the limits of R15
S Safety events	Fall detected, no motion where motion was expected, distress signal	Restricted: transmitted only to a recipient designated in advance by the subject or the subject's household. Never broadcast

Class	Examples	Publication rule
Prohibited	Identity, biometric signature, gait, individual trajectory, activity classification of a trackable person, health inference about a non-participant	Not admissible at any class, by any device, under any consent obtained from a party other than the subject

Rationale: an open predicate set is an open surveillance mandate, and the boundary between conditions and persons cannot be left to implementers because the same measurement supports both readings. Class E carries the participation incentive on which P-OBS depends in its entirety, since rain on a field, a blocked road, changed air and signal conditions are all non-substitutable and all concern a place rather than a person. Class O carried the legal and ethical exposure and almost none of the incentive, and its exclusion from the exchange therefore costs the architecture nothing it needs. The exclusion additionally keeps conformant publication outside the scope of personal-data regimes of the European type, and keeps a device owner from becoming a data controller under the narrow reading of the household exemption established in Case C-212/13, where video surveillance covering even partially a public space and directed outwards from the private setting was held not to be a purely personal or household activity.

R15 (publication limits). For each predicate class and each declared environment class, the standard shall fix:

1. a maximum information content per observation tag, expressed as a maximum number of independent quantised components and a maximum number of admissible levels per component;
2. an admissible region of joint spatial and temporal resolution, expressed as a single condition on the resolution pair rather than as two independent floors;
3. a maximum number of observation statements concerning one volume that may be published in one period, counted across all publishers whose declared sensing volumes intersect that volume;
4. for any predicate whose subject may be a person, a protection radius and a tolerated distinguishability level at that radius, from which the indistinguishability parameter is derived.

The limit on information content per tag is a precondition for the others and shall be verified first in conformance testing. Where these cannot be satisfied simultaneously, the device shall publish nothing for the volume concerned. Parameters shall be established by the calibration protocol of Annex A, published together with the adversary model and environment class against which they were derived, and revised on a fixed cycle and upon any recorded new attack.

Rationale: identification saturates within a few tens of statements where each statement is rich, so a limit on the number of statements has no effect until the content of each is bounded; once it is bounded, uniqueness accumulates across statements and the number becomes the governing quantity. Resolution decays approximately as its one-tenth power and is not a control of usable strength in either regime. Independent spatial and temporal floors additionally permit the configuration easiest to attack, since data coarse in one dimension and fine in the other is weaker than data of intermediate grain in both. Derivation in Annex A.

R16 (evaporation). Observation tags shall carry a lifetime commensurate with the volatility of the condition described, and shall not be retained, re-published or archived by receiving devices beyond that lifetime. Devices shall not construct persistent histories from received tags.

Rationale: the series holds trust and memory in time rather than in stores, and applied here the same mechanism limits the accumulation identified in clause 2.5. This requirement is classified as attested under R21: it binds conformant receivers and offers nothing against a party that records the air. Clause 6.1 states the direction in which a demonstrable equivalent would lie.

R17 (unlinkability of publishers). Observation tags shall be signed under the anonymous attestation mechanism already specified for this ecosystem, establishing that the publisher belongs to a certified class without disclosing a device or user identifier. Publisher pseudonyms shall rotate on a schedule fixed by the standard, and rotation shall not be defeatable by correlation of the declared sensing volume.

Rationale: a stable publisher identifier attached to a fixed location is a home address. Note that the declared sensing volume required by P-OBS-2 is itself a quasi-identifier, which is why rotation must be evaluated against it rather than in isolation.

R18 (declared and physical opt-out). The standard shall define a declared no-observation zone: a signed assertion, publishable by a premises holder, which conformant devices shall honour by suppressing publication concerning the volume it describes. Conformance testing shall verify suppression. The standard shall further define a collective declaration, by which the holders of adjacent premises may jointly declare a contiguous zone.

In addition, the architecture shall recognise physical suppression as legitimate: radio-absorbing and engineered surface treatments that attenuate the sensing signal within a bounded volume shall not be treated as interference, tampering or non-conformance, and devices shall not attempt to compensate for them.

Rationale: an opt-out that depends only on the good behaviour of others is not an opt-out. Physical attenuation is the one mechanism that does not require the cooperation of the observing party, and the standard should protect the right to use it rather than treat it as an attack. The collective form exists because an individual declaration cannot bound an aggregate assembled over a territory, and because declining to participate is not an instrument: it lowers the density on which neighbours depend and touches no proprietor, there being none. Practical limits of physical suppression are stated in clause 4.

R19 (local processing, no export). Measurement, interpretation and the decision to publish shall occur on the measuring device. Neither raw measurement data nor a derived representation from which the prohibited predicates of R14 could be reconstructed shall be transmitted to any party, including the device vendor, an operator or a cloud service, whether or not the user has agreed.

Rationale: any export path recreates the centralised repository this architecture exists to avoid, and does so with material more sensitive than anything else the ecosystem handles. Classified as attested under R21. This requirement is in direct tension with architectures that route sensing results to a core network for exposure to third-party applications, and with distributed-learning approaches that move model updates derived from measurement; clause 6.2 records both as unresolved.

R20 (auditability). The predicate classes a device is capable of publishing, the limits of R15 it applies including the budget accounting it performs for shared volumes, and its honouring of declared no-observation zones shall form part of its certified class profile and shall be attestable in operation without disclosing the instance. Violations detected by neighbouring devices shall be reportable through the accountability chain in the same manner as fabricated observations.

Rationale: a bound that cannot be checked from outside will be assumed to hold and will not hold.

---

4. What this specification cannot do

- A device that ignores the standard is unaffected by every requirement above. Conformance bounds conformant products; it does not bound purpose-built equipment.
- Signal overhearing requires no cooperation and no participation. Any party with a receiver can estimate the channel from transmissions intended for someone else. The countermeasures are physical and cryptographic rather than normative, and belong to Part Eight.
- Limiting how often something is published does not by itself protect anything. This is not a conjecture: an identification study removed entries uniformly from its recordings to reduce temporal resolution and found accuracy almost unaffected. The budget of R15 is meaningful only for statements already flattened by R13.
- Coarsening is weak in a quantified way. Halving resolution makes a subject about nine per cent harder to single out given four known data points, and about six per cent harder given ten. Its effectiveness declines as the observer gains data, which is the wrong direction.
- Single-occupancy spaces are outside the reach of calibration. Where one person is present, every occupancy statement concerns that person, and no resolution and no budget makes it anonymous. What protects such a space is suppression under R15 and the opt-out of R18.
- Physical suppression under R18 is real but narrowband and partial. Absorbing treatments are frequency-selective, and a volume shielded at one band may be transparent at another. R18 protects the right to attenuate; it does not promise that attenuation is complete.
- The limits of R15 are a statistical bound, not a guarantee. An adversary combining conformant tags with data from outside this ecosystem may still re-identify.
- Two requirements, R16 and R19, are attested rather than demonstrable. Clause 6 treats this as an open problem rather than an accepted cost.

Stating these limits is not a weakening of the specification. A privacy claim that does not name its own boundary is the kind of claim later found to have been marketing.

---

5. Relationship to other work

5.1 IEEE 802.11bf. The sensing amendment standardises the measurement procedure and was published in September 2025. Independent analyses set out the overhearing surfaces and candidate countermeasures cited in clause 2.4. This specification addresses the layer above, where measurements become published statements exchanged between independently owned devices.

5.2 Cooperative perception in the vehicular domain. The decision to broadcast interpreted objects rather than raw sensor data is precedent for R13. The plausibility rule in the parent contribution has no such precedent: the associated misbehaviour-detection standards cover cooperative awareness messages and left detectors for collective perception to a future version, and the rule was proposed in the research literature. The vehicular case also demonstrates the limit of the transfer: a road is a public space and a living room is not, and predicates acceptable in the former are not acceptable in the latter. R14 exists because the transfer is not automatic.

5.3 Data protection regimes. Where an observation supports identification of an individual it is personal data under regimes of the European type regardless of whether a name is attached, and the classes of R14 are drawn to keep conformant publication outside that boundary. Three interactions are noted without legal argument, this document not being legal advice. Data protection by default requires that only the minimum necessary personal data be processed by default, including as to disclosure, which is why P-DENS-2 fixes a shipped state for participation in carriage and R14 excludes class O from publication rather than defaulting it off. Physiological predicates fall within the special categories whose processing is prohibited absent an exception, which is why class S is restricted to a recipient designated by the subject. And the household exemption is read narrowly where a device is directed outwards from the private setting, which is why R14 matters to whether an ordinary owner becomes a controller.

5.4 ISAC governance literature. The requirements here are intended to be complementary to the governance proposals cited in clause 2.1 rather than an alternative to them. Purpose-bound activation, transparency toward citizens and accountability for inference models operate at a layer this specification does not reach; conversely, a governance regime with no bounded technical object to govern has nothing to bind. R14 and R15 are offered as that object. Where the governance literature treats consent as the central instrument, this specification treats it as unavailable for the primary subject and relies instead on excluding the predicate, which is why R14 rather than a consent mechanism carries the weight.

5.5 P-AUTH and P-INDEP. The personal baseline protected under those principles and the observations bounded here are the same physical quantity measured for different purposes. A baseline held on the device for authentication is protected by non-storage and carrier independence; an observation about a volume is protected by abstraction and by the limits of R15. The two must not be permitted to converge, and R19 exists so that no export path allows an observation channel to become a baseline repository by accumulation.

---

6. Open questions

6.1 A demonstrable equivalent for R16 and R19. Both are absences that only a certifying body can attest to. The direction in which a demonstrable form would lie is indicated by the parent contribution's open question on aggregation: if the readability of a tag were bound to the state of the physical channel at the moment of publication, a distant collector would receive ciphertext and a recorded tag would become unreadable as the channel decorrelates. A lifetime enforced by decorrelation is demonstrable in a way that a lifetime enforced by a timer is not. Whether this is achievable given synchronisation tolerance between publisher and receiver, and what it would do to legitimate delayed readers, is unresolved.

6.2 Distributed learning. Approaches in which model updates derived from local measurement are exchanged sit in direct tension with R19, since an update is a statement about the data from which it was computed. Whether the information ceiling of R15 can be applied to such updates, and at what cost to the learning, is not addressed here.

6.3 Numerical values. The parameters of R15 do not exist. Annex A specifies how to obtain them. The information ceiling is first, because the others are conditional on it, and no accepted method of measuring that quantity in conformance testing is proposed here.

6.4 A lower bound on the publication budget. Verification of a claim requires at least two independent statements about the same subject, so a budget chosen on privacy grounds alone may leave no capacity for verification. Whether both can be satisfied at once, and in which environment classes, is the first thing calibration should test.

6.5 Authority over declared zones. Who may declare a no-observation zone under R18, over what volume, how competing declarations between neighbouring premises are resolved, and what quorum a collective declaration requires.

6.6 Class S consent. Safety events are the strongest argument for sensing and the hardest case for consent, because the subject may be unable to give it at the moment it matters. The designated-recipient model is a starting position and not a settled answer.

---

7. Conclusion

The parent contribution argues that a network with no owner begins when devices exchange what each can measure and none can obtain otherwise. That mechanism works, and it works because observation is scarce and local. Scarcity that is local is also intimate.

A distributed publication channel about physical space, built into a large number of devices, with no owner to switch it off, is either the most carefully bounded thing in the architecture or the least defensible. The difference is a closed predicate list that excludes the inside of rooms, a ceiling on what one statement may contain, a budget over a volume rather than a publisher, a lifetime on every statement, no path by which raw measurement leaves the device, and an opt-out that does not depend on anyone else's good behaviour.

None of that is expensive. All of it has to be decided before the first product ships, because a publication channel is far easier to bound while it is a document than after it is a market.

Annex A. Calibrating the limits of R15

Normative basis of R15. Previously circulated as a separate annex; incorporated here.

A.1 Why this annex exists

The first draft of R15 rested on an assumption that seemed safe: that publishing coarse statements about a volume of space, instead of fine ones, is a sufficient defence against identifying the people inside it, and that the only open work was to choose the numbers. Choosing the numbers is what this annex set out to do. The attempt failed in an informative way, twice, and this annex records why rather than presenting the current text as though it had always been obvious. A reader who cannot see what was rejected cannot judge what was kept.

A.2 What the evidence says

A.2.1 Coarsening buys very little

The reference result concerns mobility traces rather than observation tags, but the structure of the data is the same: a set of statements placing a subject in a region at a time. Studying fifteen months of data for one and a half million people, the authors found that with hourly temporal resolution and a spatial resolution equal to the operator's antenna grid, four spatio-temporal points sufficed to uniquely characterise ninety-five per cent of individuals. They then coarsened the data deliberately in both dimensions and derived uniqueness as a power function of resolution, decaying approximately as its one-tenth power.

The authors quantify what that buys. Each halving of resolution reduces uniqueness by a constant factor, and given four known points a two-fold coarsening makes identification 9.3 per cent less likely. The exponent itself weakens as the adversary learns more, decaying linearly in the number of known points, so at ten points the same coarsening yields a reduction of only 6.2 per cent. Coarsening is least effective precisely against the adversary who already knows most, which is the adversary a specification exists to constrain. The authors state the conclusion plainly: even coarse datasets provide little anonymity.

A.2.2 For low-dimensional statements, the number is the steep variable

The same study gives the parameter that does move quickly. Four randomly chosen points uniquely characterise more than ninety-five per cent of traces; two already characterise more than half.

The scope of that result should be stated precisely, because A.2.3 narrows it. Each point is a single low-dimensional item: one region, one time. For statements of that kind the count is the steep variable, and the count is entirely within the reach of a specification, where it was not being controlled at all. It does not follow that the count is the steep variable for statements of every kind.

A.2.3 For high-dimensional statements, the number stops mattering

An identification study published in 2025 tested the opposite regime and, in effect, tested the control proposed in the first draft of this annex. Its authors removed entries uniformly from each recorded sequence, reducing temporal resolution, and retrained the classifier. Accuracy fell only slightly. They conclude that identification of this kind is extremely robust to lower sample rates and that a mitigation based on time-coarsening appears ineffective.

The comparison inside that study explains why. Their beamforming sequences average about forty points, each carrying 740 features, and yield 99.5 per cent accuracy. Their channel-state sequences average about 1241 points, each carrying 212 features, and yield 82.4 per cent. The richer representation wins on roughly a ninth of the total data, which the authors attribute to compression acting as feature extraction and to higher spatial resolution, concluding that spatial resolution matters more than temporal resolution.

Placed beside A.2.2 this is not a contradiction. The two studies describe statements of different dimensionality. A single location at a single time carries little, so uniqueness accumulates across points and the count governs. A 740-dimensional channel description carries a great deal, so uniqueness saturates within a few tens of points and the count no longer governs, because the adversary passed the threshold long before any budget could bind.

The consequence is a precedence order rather than a new control. R13, which forbids publication of raw channel data and requires interpretation on the measuring device, is what places observation tags in the first regime rather than the second. Until it holds, no budget is worth setting. Once it holds, the budget becomes the binding constraint.

It also exposes a gap. The word "interpreted" is a description, not a bound. Nothing in R13 alone distinguishes a statement that a corridor is occupied from a vector of beamforming angles presented as an interpretation. Control C0 below closes that gap.

A.2.4 Asymmetric coarsening is the weakest configuration

The mobility study also reports that a dataset coarse along one dimension and fine along the other is easier to attack than a dataset of intermediate grain along both: with four points, more than sixty per cent of traces were uniquely characterised at a fifteen-hour temporal resolution, while fewer than forty per cent were characterised at a seven-hour resolution combined with spatial clustering.

Two independent floors therefore permit exactly the configuration easiest to attack, since an implementer optimising for utility will hold one dimension at the floor and keep the other as fine as the floor allows. The constraint has to be joint.

A.3 The control set

Five controls, applying in order.

C0, information ceiling per tag. An observation tag shall carry no more than a stated maximum of information about the volume it describes, expressed per predicate class as a maximum number of independent quantised components and a maximum number of admissible levels per component. Publication of raw or minimally processed channel measurements is prohibited outright by R13 and is not subject to any ceiling.

Every control below is conditional on this one, and a qualitative requirement to publish interpretations is not enforceable without a quantitative ceiling.

C1, joint granularity condition. Spatial and temporal granularity are constrained by a single joint condition rather than two independent floors, so that fineness in one dimension must be paid for by coarseness in the other.

C2, publication budget per volume. A maximum number of observation statements concerning a given volume in a given period, counted across all publishers rather than per publisher. Publishers whose declared volumes intersect share one budget for the intersection.

This is the binding control once C0 holds, because it acts on the variable that governs uniqueness most steeply in the flattened regime. Attaching the budget to the volume rather than the publisher is also what dissolves the apparent conflict with verification: coverage may remain as high as verification requires, while the total emitted about the volume is capped.

C3, declared protection radius. For any predicate whose subject may be a person, the mechanism shall satisfy a distance-scaled indistinguishability criterion: the probability of any published output shall differ by at most a bounded factor between any two subject positions, with the bound growing in proportion to the distance between them. The parameter is the ratio of a tolerated distinguishability level to a declared protection radius.

The advantage over a floor is that the guarantee holds irrespective of what the adversary already knows.

C4, suppression instead of degradation. Where C0 to C3 cannot be simultaneously satisfied for a volume during a period, the device shall publish nothing for that volume. Silence is a permitted and expected output, and for class O predicates in single-occupancy environments it is the normal one rather than an edge case.

A.4 Adversary model

Calibration is meaningless without a stated adversary. A candidate parameter set must be evaluated against all four, and the operating point set by the strongest.

Adversary	Capability	Measured quantity
A1 Passive collector	Receives all published tags within radio range over a period. No side information	Uniqueness of subject traces reconstructed from tags, as a function of statement count and resolution pair
A2 Multi-publisher aggregator	Collects tags from all publishers whose volumes intersect the target volume, and fuses them	Localisation error of a single subject; whether the joint budget holds under fusion
A3 Side-information join	Holds an external dataset containing a small number of known spatio-temporal points about the target	Probability of correct linkage as a function of known points, tested at two, four and ten
A4 Colluding publishers	Controls a minority of publishers in the region and may publish truthfully in order to increase coverage of a target volume	Increase in exposure attributable to collusion; whether the volume budget is enforceable against a bad-faith publisher

A4 is the case in which the volume budget is hardest to enforce, because a colluding publisher gains by exhausting the budget with statements of its own choosing. Whether the budget can be enforced without a coordinator is the principal unresolved question of this annex.

A.5 Calibration protocol

- **Environments.** Calibrate separately for at least four occupancy regimes: single-occupancy dwelling, multi-occupancy dwelling, workplace, public thoroughfare. The single-occupancy dwelling is the binding case.
- **Data.** Recorded sensing data with informed consent from all subjects present, or synthetic traces validated against recorded data. Calibration shall not be performed on data collected from non-participants, since that would reproduce the harm under study.
- **Sweep.** For each environment, sweep three quantities together: information content per tag, the resolution pair, and the publication budget. Evaluate each combination against A1 to A4. Sweeping resolution alone reproduces the first error this annex corrects; sweeping resolution and budget without per-tag content reproduces the second.
- **Privacy metric.** Uniqueness at two, four and ten adversary points, and re-identification accuracy expressed as the margin above the accuracy obtainable by guessing given the population of the environment. Absolute accuracy without the chance baseline is not interpretable.
- **Utility metric.** The fraction of the intended function of each predicate class that survives, measured on the same data: detection latency for safety events, error of aggregate occupancy where permitted, and residual cross-checking power available to a verifier. A parameter set that satisfies privacy and destroys utility will not be implemented and therefore protects nobody.
- **Operating point.** Among parameter sets meeting the declared privacy bound against the strongest adversary, select the one maximising utility. Where no set meets the bound, the predicate class is not publishable in that environment and is removed from the allowlist for that environment rather than published with weaker parameters.
- **Reporting.** Publish per environment and predicate class: information ceiling, resolution pair, budget, protection radius, adversary, privacy metric, utility metric, together with the data description and the code used. A parameter without its adversary and environment is not a result.
- **Revision.** Re-identification methods improve. Repeat on a fixed cycle, and off-cycle upon any new attack recorded in the accountability chain.

A.6 What calibration cannot settle

- Single-occupancy spaces, per C4.
- External side information: A3 is bounded by assumption, and a richer real adversary will do better.
- Non-conforming devices, and signal overhearing, which requires no publication at all.
- Method drift. Parameters are calibrated against attacks known at the time, which is the ordinary condition of security parameters and the reason for the revision cycle.

A.7 Open questions of this annex

- **Measurement of information content per tag.** C0 requires a ceiling, and expressing it in components and levels is easy to state and does not by itself bound what a determined implementer can encode. Whether an entropy-based measure over the published predicate set is workable in conformance testing, and what it would cost to verify, is unresolved. This is the first parameter calibration should produce.
- **A lower bound on the budget**, per clause 6.4 of the parent document.
- **Enforcement without a coordinator**, per A.4.
- **Composition across predicate classes.** Parameters are calibrated per class; an adversary combining classes about one volume may exceed the bound each satisfies separately.
- **The unit of the period.** A long period admits bursts and a short one admits a steady high rate. The shape of the limit, not only its magnitude, requires specification.

A.8 Conclusion of the annex

The requirement this annex was written to complete rested first on a variable of the wrong kind, and the first attempt to replace it reached for a variable that was right in only one regime. Resolution is what a designer naturally reaches for and is close to useless at usable strengths. The number of statements governs, but only once each statement has been flattened. The order is fixed and not a matter of preference: bound what a statement may contain, then bound how many are published.

Invitation. This specification is submitted for technical review with its parent contribution. The authors invite criticism in particular of the predicate classes of R14, of the enforceability of the volume budget in the absence of a coordinator, and of whether a demonstrable equivalent for R16 and R19 exists.

References

All links checked in July 2026. Entries [1] to [3] are documents of this activity and of the accompanying series: contributions are held by IEEE SA and available to activity participants, articles are published by the authors, and neither category has a public URL of the kind given for the external sources that follow.

[1] *Density as an Availability Condition for Environment-Dependent Trust Roots: The P-DENS and P-OBS Principles*, contribution to IEEE IC25-009-01, v2.0, July 2026.

[2] *Presence Protocol* (P-AUTH), *Accountability Chain Protocol* (ACP), *The Independence-of-Foundations Principle* (P-INDEP), and *Resolution of the P-INDEP Open Questions and the P-DELEG Principle*, contributions to IEEE IC25-009-01, 2026.

[3] *A Network With No Owner, a Reflection With One*, Part Eight of the accompanying series, July 2026. Source of the gradient definition of harm in clause 2.2.

[4] J. Todt, F. Morsbach, T. Strufe, "BFId: Identity Inference Attacks Utilizing Beamforming Feedback Information," *Proc. ACM CCS 2025*, Taipei, 13-17 October 2025. DOI 10.1145/3719027.3765062. Open access: https://publikationen.bibliothek.kit.edu/1000185756. Source of clause 2.3 and of A.2.3.

[5] Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, V. D. Blondel, "Unique in the Crowd: The privacy bounds of human mobility," *Scientific Reports* 3, article 1376, 2013. DOI 10.1038/srep01376, open: https://www.nature.com/articles/srep01376. Source of A.2.1, A.2.2 and A.2.4.

[6] M. E. Andres, N. E. Bordenabe, K. Chatzikokolakis, C. Palamidessi, "Geo-indistinguishability: differential privacy for location-based systems," *Proc. ACM CCS 2013*, Berlin, pp. 901-914. DOI 10.1145/2508859.2516735. Open access: https://inria.hal.science/hal-00766821. Source of control C3.

[7] IEEE Std 802.11bf-2025, Amendment 4: Enhancements for Wireless LAN Sensing, published 26 September 2025. Working group record: https://www.ieee802.org/11/. Free through the IEEE GET programme.

[8] ETSI TS 103 324 V2.1.1 (2023-06), *Collective Perception Service; Release 2*, free: https://www.etsi.org/deliver/etsi_ts/103300_103399/103324/02.01.01_60/ts_103324v020101p.pdf. SAE J3224_202208, August 2022, DOI 10.4271/J3224_202208 (sold).

[9] ETSI TR 103 460 V2.1.1 (2020-10), *Pre-standardization study on Misbehaviour Detection*, free: https://www.etsi.org/deliver/etsi_tr/103400_103499/103460/02.01.01_60/tr_103460v020101p.pdf; and ETSI TS 103 759 V2.2.1 (2026-01), *Misbehaviour Reporting service*, free: https://www.etsi.org/deliver/etsi_ts/103700_103799/103759/02.02.01_60/ts_103759v020201p.pdf.

[10] Case C-212/13, *František Ryneš v Úřad pro ochranu osobních údajů*, Judgment of the Court (Fourth Chamber), 11 December 2014, ECLI:EU:C:2014:2428. Full text in English at https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:62013CJ0212. Party names appear in the original Czech in the official English-language record; the respondent is the Czech data protection authority. Cited in R14 and clause 5.3.

[11] ISO/IEC 20008-1:2013, *Anonymous digital signatures; Part 1: General*, https://www.iso.org/standard/57018.html, and ISO/IEC 20008-2:2013, *Part 2: Mechanisms using a group public key*, https://www.iso.org/standard/56916.html; both sold, links to the official records. Trusted Computing Group, *TPM 2.0 Library Specification*, free: https://trustedcomputinggroup.org/resource/tpm-library-specification/, published as International Standard ISO/IEC 11889:2015 Parts 1 to 4. Direct anonymous attestation is specified in ISO/IEC 20008-2 and supported by the TPM 2.0 library. Cited in R17.

[12] *ISAC Privacy: Challenges and Solutions for 6G*, 2026, https://arxiv.org/abs/2605.28325. Source of the bystander framing acknowledged in clause 2.1 and of the open problem of measurable privacy criteria for system design.

[13] A. Sedrati, *Unconsented Sensing: A Sociotechnical Governance Framework for 6G ISAC*, KTH Royal Institute of Technology, 2026, https://arxiv.org/abs/2605.07328. Source of the argument that technocentric mitigations are insufficient, acknowledged in clause 2.1.

[14] *Seeing Without Being Seen: A Review of Ethical and Human-Centric ISAC in 6G*, Telecom 7(3), article 52, 2026. DOI 10.3390/telecom7030052, open access: https://doi.org/10.3390/telecom7030052. Source of the distinction between active user, passively sensed subject and bystander.