The First Node

How a Network With No Owner Actually Begins

Part Seven of the series on the architecture of open 6G

A person buys the first open 6G device. It arrives, they charge it, they switch it on. And nothing happens. Not because it is broken. Because there is nobody to talk to.

Six articles had made a case: a network can exist without an owner, trust can live in a trace rather than in a vault, and the standard is where all of it is actually decided. Every one of them described how such a network behaves once it exists. The farmer's sensor reaches a chain of neighbouring devices. The doctor gets the alert through seven hops. A captured node is noticed missing by the neighbours who stop seeing its trace.

All of that holds, and all of it rests on the same unstated premise: that the neighbours are already there.

Nobody wrote down where they come from.

A network is not a gradual thing

There is a temptation to picture adoption as a slope. A few devices, a little connectivity; more devices, more connectivity. It is not a slope. It is a cliff.

The mathematics is old and unambiguous. Scatter points at random across a plane and connect any two that fall within radio range of each other. Then ask when a path exists that crosses the whole plane instead of dying inside a small island. The answer is a threshold. Below a critical average number of neighbours, only isolated clusters exist and long-range connectivity is impossible; above it, a spanning cluster appears almost at once. The best numerical work places that critical average between 4.508 and 4.515 with high confidence, while the bounds that have been proved rigorously remain several times wider and no exact value is known.

The practical reading matters more than the number. Below the threshold the network does not perform worse. It does not exist. There is no partial credit.

That has a consequence the series never faced. Several of the strongest guarantees in this architecture are the ones that need company. A trust root anchored off the device and verified by geometry needs independent vantage points before it can verify anything. A neighbour who notices that a device has gone silent has to be a neighbour first. These are not weaknesses in the design. They are the reason the design is strong. But every one of them is conditional on something the design does not state: that a neighbourhood exists.

Engineering has a word for this, and the series was not using it. When we discuss a root of trust we almost always discuss its integrity: can it be forged, can it be extracted, can a single break reach two roots at once. There is a second property, equally standard in the dependability literature and equally load-bearing. Availability: readiness to give correct service when asked. A root that cannot be forged but is not there is not a root. It is a plan.

So the question of where the first devices come from is not a question about marketing. It is a question about whether the security model has anything to stand on in year one.

The device has to be worth owning alone

The first requirement follows directly. If a device becomes valuable only after other people have bought one, nobody buys the first one, and the argument ends there.

So the device has to be worth owning when it is the only one.

That sounds like it would cost something: another radio, another chip, a larger battery. It costs almost nothing, because the capability is already in the box.

A radio is a sensor. The same antenna and the same waveform that carry data also measure the space they cross. This is not a proposal. A Wi-Fi amendment published in September 2025 standardises exactly this, using measurements the hardware already produces to estimate features of objects in a room. Sensing and communication sharing one infrastructure rather than two is one of the six headline usage scenarios of the international framework for 6G, agreed in late 2023.

And here the physics turns friendly in an unexpected direction. In the very high frequency bands where 6G is expected to live, the channel is difficult precisely because the signal interacts strongly with its surroundings. It is absorbed by water vapour and by oxygen. It is stopped by walls. It fades with distance faster than anyone would like. Every one of those problems is what makes a good sensor. The sensitivity that ruins the link is the sensitivity that reads the room.

Which brings back the person the third article was written for. A device on a bedside table that registers a fall, a change in breathing, an absence of movement where movement was expected, does all of that with zero participating neighbours. It is useful on the first day, in a house where no other open device exists. The network, when it arrives, extends the reach of the alert. It does not create it.

This is not an unprecedented idea, and it would be dishonest to present it as one. In platform economics it is a well-known strategy, usually phrased as coming for the tool and staying for the network. What is proposed here differs in exactly one respect, and it is the respect that matters for a standard. As a commercial tactic it is optional. As an answer to a threshold in physics it is a condition of existence, and conditions of existence belong in the specification.

Why would a self-sufficient device ever connect?

Solve the first problem well and you create the second one.

If a device is genuinely good on its own, its owner has no particular reason to relay a stranger's traffic. Relaying costs battery. It raises awkward questions about what is passing through your hardware. Life works fine without it. Multiply that by everyone and the result is not a network but a landscape of small private islands, each perfectly content. That outcome would satisfy every requirement written so far and deliver none of the point.

The way out begins by noticing a hidden assumption: that the only thing devices could trade is traffic.

Traffic is substitutable. Your neighbour can buy it from an operator, and so can you. Trading it is close to a zero-sum negotiation, which is why it always has to be paid for, policed and incentivised.

Observation is not. Your device measures the space it occupies and nothing else. What is happening two hundred metres away, whether a road is blocked, whether the air has changed, whether this field got rain, is information you cannot obtain by any route other than a device that is there. Your neighbour holds something you physically cannot generate, and you hold the same for them. That is not a negotiation. It is an exchange between parties who each have what the other lacks.

Two things have to be said about that argument, and both of them narrow it.

The first is a boundary. Two devices in the same room see the same thing, and trading identical views buys nothing. Two devices with no shared view at all can trade information but cannot check each other, which matters the moment someone has a reason to lie. What works is partial overlap: enough common ground to cross-check what a neighbour claims, enough difference to learn something new. And the property does not last. When fifty devices watch the same street, each of them is replaceable, and observation becomes as substitutable as traffic. Non-substitutability is not a property of the good. It is a property of a density regime, and it fades in the regime this architecture is designed to produce.

The second is more uncomfortable, and it changes what the argument is for. Nobody makes this decision. The transition between behaving as an endpoint and behaving as a relay is automatic, driven by what the device measures around it, and what it does on the day it leaves the factory is set by whoever configured it. Owners do not weigh substitutability. Owners do not weigh anything. The overwhelming majority of devices will do whatever they were shipped doing, for as long as they exist.

So the reasoning above is not a description of a choice. It is a justification for a requirement. It explains why a standard is entitled to insist that the default be participation rather than isolation, and it explains that in terms of the owner's own interest rather than an appeal to solidarity. That is a smaller claim than the one this section originally made, and it is the one that survives contact with how people actually behave.

Which relocates the whole question. The interesting fight is not about persuading owners. It is about who is allowed to set the factory setting.

What bounds publication, and what does not

If devices are going to publish what they observe, something has to bound what they publish, and the intuitive bounds turn out to be the weak ones.

The instinct is to make each statement vaguer. Measurements from an adjacent field, on mobility traces rather than radio observations, show how little that buys: identifiability falls roughly as the tenth root of how much you blur. Concretely, halving the resolution makes a person about nine per cent harder to pick out when the observer holds four data points about them, and about six per cent harder when the observer holds ten. Coarsening works worst against the observer who already knows most, which is the observer a specification exists to constrain.

The instinct that follows is to publish less often. That fails too, and the way it fails is instructive. Researchers who identified people from ordinary Wi-Fi signals thinned out their recordings to test exactly this defence and found their accuracy barely dented. Roughly forty observations were already more than enough, because each observation was extraordinarily rich. In the same study, a compressed channel description of about forty samples outperformed a raw one of over a thousand.

What governs the outcome is what a single published statement is allowed to contain. Bound that first, and how often you publish begins to matter. Leave it unbounded, and nothing else you do will help.

That ordering is not a detail. It is the difference between a specification that feels responsible and one that is.

The part we have not solved

An honest article names its hole rather than leaving it for a reader.

Everything above describes an exchange between neighbours. It says nothing about a third participant who publishes nothing and collects everything.

Producing an observation is stubbornly local: to know about a place you must have a device there, and covering the two hundredth place costs as much as the first. That linearity is what makes collection resistant to centralising. But the value of observations is not linear. A map of a whole city is worth vastly more than the sum of two hundred separate readings, and whoever assembles it captures the difference.

This is not a hypothetical. Traffic layers in mapping products, crowd-sourced positioning, activity heatmaps: every one of them is built from non-substitutable local observations produced by distributed devices, and every one of them is centralised. The physical stubbornness of the raw material did not prevent it. It was simply irrelevant to it, because the concentration happened one layer up.

An architecture that removes the centre from the carriage of data and says nothing about the assembly of knowledge has moved the centre rather than removed it. There are candidate answers, and they cost something: making readability conditional on also publishing, or binding it to physical proximity so that a distant collector receives only ciphertext. Both make the channel less open than the first six articles argued for. That is a real price and it should be paid deliberately rather than discovered later.

What can be said now is only this: the problem is named, the direction is known, and it is not solved in this article.

Density has a ceiling

There is a second correction, and the next article in this series supplies it.

The argument so far has treated density as something to maximise: below the threshold nothing works, so build upwards. But a signal that must travel out to a body and back again weakens twice over, which means the advantage a person holds over the infrastructure in reading their own body is a function of distance. Doubling the number of nodes per unit area spends about six decibels of that advantage. Doubling them along a street, which is how relays are actually strung, spends twice that.

So density buys availability of trust roots and sells privacy margin, at a known exchange rate. It has an optimum, not a direction.

Three separate lines of reasoning arrive at the same place. Trust roots need density. Privacy margin is spent by density. Non-substitutability, the thing that makes neighbours want to exchange at all, is dissolved by density. A requirement that states only a floor is therefore incomplete, and the reference architecture needs a ceiling per environment as well as a threshold.

What changes

For the farmer: rainfall on this field rather than the regional average, starting on the day the first two neighbours switch their devices on.

For the doctor: monitoring that works before there is a network to carry the alert, and a network that grows because the monitoring gave people a reason to own the device at all.

For the teacher: a school whose connectivity depends on devices within a few kilometres, in a place where those devices now have an independent reason to be there.

And for everyone, a network that stops being a pipe and becomes something closer to a shared sense of the surroundings. That phrase should be handled carefully, because a shared sense of the surroundings and a distributed eye are the same object described by someone who benefits and someone who does not. Which of the two it turns out to be is decided by the bounds in the preceding sections, not by the intention of anyone building it.

One more thing is worth saying plainly, because the series has not said it. This architecture does not abolish power. It moves it. Operators lose the rent they collect on carriage. Device manufacturers gain, because sensing sells a product whether or not anything is ever relayed. Certification bodies gain a great deal, because in a system with no operator they become the party to whom a demand can be addressed. The owner of a device gains options and gains no collective instrument: declining to participate lowers the density that neighbours depend on and touches no proprietor, because there is none.

None of that is an argument against the design. It is an argument for saying where the power went, so that the answer is not left to be discovered by whoever ends up holding it.

This is checkable

Everything above is a claim about behaviour, and claims about behaviour can be wrong.

This one can be tested before it is standardised. Deploy devices in comparable areas in three configurations: communication only; communication with local sensing; communication with sensing plus exchange of observations between neighbours. Then measure how many devices are participating after a fixed period. Randomise by area rather than by person, because the effect under study spreads between neighbours by construction. Fix the metric before collecting the data, and observe again a year later, because an advantage visible after three months is exactly what a novelty effect looks like.

If the third configuration does not beat the first, the argument in this article is wrong and should be dropped. That is the point of writing it down this way.

The third configuration also publishes information about places where people who are not enrolled in anything happen to be, which means it cannot be run before the limits described above have been given numbers and its effect on those people has been assessed. Those numbers do not exist yet. Producing them is the first piece of work this argument requires, and until it is done the trial is a proposal rather than a plan.

Proposed for the reference architecture

P-DENS. Threshold density and standalone utility.

1. No function on which a device's value to its owner depends shall require the presence of participating neighbours. Every device shall provide measurable standalone utility at zero participating neighbours, and sensing performed by the communication transceiver is recognised as a first-class capability rather than an optional extension.
2. The device shall measure participating-node density locally and shall transition between endpoint and relaying behaviour autonomously, on the basis of that measurement, according to a threshold fixed by the standard rather than by an operator or a vendor. Its behaviour as shipped is part of that specification and not a vendor prerogative.
3. The reference architecture shall state, per device class and per environment, both the density at which each distributed function becomes operative and the density above which it shall not be driven, and shall treat the availability of environment-dependent trust roots as a property to be specified rather than assumed.

Connectivity in a distributed network is a phase transition, not a gradual improvement; below the critical density the network does not degrade, it does not exist. Above a further density the same geometry that carries the link erodes the margin by which a person reads their own body better than the infrastructure does. Density is an optimum, and both ends of it belong in the specification.

P-OBS. Observation as a second exchange layer.

The device shall be capable of publishing, alongside path-quality tags, signed statements describing locally measured conditions within its declared sensing volume. Such statements shall carry a bounded information content, an independent lifetime, and a declaration of the volume they derive from; shall be readable by neighbouring devices without operator mediation; and shall impose no obligation to relay third-party traffic. They shall describe the state of the environment. The occupancy of enclosed spaces, and the presence or activity of the people in them, shall not be published to neighbouring devices at all, and shall be transmitted only to a recipient designated in advance by the subject.

Traffic is substitutable and its exchange approaches zero sum wherever an operator exists. Local observation is not substitutable at low density, and an exchange of non-substitutable goods gives participation a rationale that does not depend on solidarity or on payment. The rationale justifies a default; it does not describe a decision that owners make. Restricting the published predicate to the environment costs the exchange nothing that carries the incentive: rain on a field, a blocked road, changed air, signal conditions are all non-substitutable and all concern a place rather than a person. The same mechanism is already standardised for cooperative perception between vehicles.

What a first node is

A network with no owner is not built. It condenses, the way frost forms on a window: not everywhere at once, but around the points where something was already there. Physics can tell us the conditions are favourable. It does not supply the first speck.

The first speck is a device that was worth owning before the network existed, and a reason to turn towards the neighbours once it does.

But this article should not end on its own terms, because the sentence that follows it is not ours to write here. Everything above concerns people who choose to take part: they buy a device, it observes for them, they exchange what it sees. The next part of this series is about the person who chose nothing, carries nothing, agrees to nothing, and is measured anyway, because a network that can carry a message can measure the room the message came from.

Building the first node is the easier of the two problems. It is the one with an answer.

Part of a series accompanying contributions to IEEE Industry Connections activity IC25-009-01, A Technical Reference Architecture Framework for an Open 6G Device Ecosystem. Parts One through Six argued that the network should have no owner. This part states the condition under which such a network can begin at all. Part Eight addresses the person who is measured without taking part. The technical requirements following from these principles, together with a specification bounding what may be published and an annex deriving its limits, are the subject of separate contributions.

Sources and prior work

All links below were opened and checked in July 2026. Where a document is sold rather than published, this is stated and the link is to the official record.

Percolation. P. Balister, B. Bollobas, M. Walters, "Continuum Percolation with Steps in the Square or the Disc," *Random Structures and Algorithms* 26(4), 2005, pp. 392-403. DOI 10.1002/rsa.20064. Author copy, open: http://web0.msci.memphis.edu/~pbalistr/papers/contperc.pdf. Places the critical mean degree between 4.508 and 4.515 with 99.99 per cent confidence and notes that the rigorously proved bounds differ by roughly a factor of four. The model is E. N. Gilbert, "Random Plane Networks," *J. Soc. Indust. Appl. Math.* 9(4), December 1961, pp. 533-543, DOI 10.1137/0109045 (sold).

Availability as a dependability attribute. A. Avizienis, J.-C. Laprie, B. Randell, C. Landwehr, "Basic Concepts and Taxonomy of Dependable and Secure Computing," *IEEE Transactions on Dependable and Secure Computing* 1(1), January-March 2004, pp. 11-33. DOI 10.1109/TDSC.2004.2 (sold; author copies are in several institutional repositories). That the attribute has not been carried across to roots of trust is not an impression: a 2026 survey of the field, "Root of trust: survey, taxonomy, and open challenges," *Telecommunication Systems*, DOI 10.1007/s11235-025-01398-5, categorises roots of trust by function and evaluates them against confidentiality, integrity and authentication, with availability absent from the list.

Adjacent work on percolation and security, addressing minimum node density for connected secrecy in the presence of eavesdroppers rather than the availability of a trust root: A. Sarkar, M. Haenggi, "Percolation in the Secrecy Graph," https://arxiv.org/abs/1107.4613; R. Vaze, "Super Critical and Sub Critical Regimes of Percolation with Secure Communication," https://arxiv.org/abs/1104.1075. Both open.

Sensing by communication hardware. IEEE Std 802.11bf-2025, Amendment 4 to IEEE Std 802.11: Enhancements for Wireless LAN Sensing, published 26 September 2025; working group record at https://www.ieee802.org/11/, and the text is free of charge through the IEEE GET programme. Recommendation ITU-R M.2160-0 (11/2023), "Framework and overall objectives of the future development of IMT for 2030 and beyond," approved by the ITU Radiocommunication Assembly, Dubai, November 2023; free at https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.2160-0-202311-I!!PDF-E.pdf. Integrated sensing and communication is one of its six usage scenarios. Channel modelling was finalised in the 3GPP Release 19 study item on integrated sensing and communication in May 2025; official record for TR 38.901 at https://www.3gpp.org/dynareport/38901.htm.

The weakness of coarsening. 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, 25 March 2013. DOI 10.1038/srep01376, open: https://www.nature.com/articles/srep01376. Source of the tenth-power relation, of the figures of 9.3 and 6.2 per cent, and of the finding that the exponent weakens as the observer gains data points.

Saturation. 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 via the KIT repository at https://publikationen.bibliothek.kit.edu/1000185756. Source of the finding that thinning recordings barely reduces identification accuracy, and of the comparison between roughly forty rich samples and over a thousand raw ones.

Cooperative perception. 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, "V2X Sensor-Sharing for Cooperative and Automated Driving," August 2022, DOI 10.4271/J3224_202208, official record at https://www.sae.org/standards/j3224_202208-v2x-sensor-sharing-cooperative-automated-driving (sold).

The plausibility check adopted here is a research proposal rather than a standardised detector, and is due to M. R. Ansari, J.-P. Monteuuis, J. Petit, C. Chen, "V2X Misbehavior and Collective Perception Service: Considerations for Standardization," IEEE CSCN 2021, DOI 10.1109/CSCN53733.2021.9686156; preprint open at https://arxiv.org/abs/2112.02184. The corresponding ETSI documents, TR 103 460 V2.1.1 (2020-10) at https://www.etsi.org/deliver/etsi_tr/103400_103499/103460/02.01.01_60/tr_103460v020101p.pdf and TS 103 759 V2.2.1 (2026-01) at https://www.etsi.org/deliver/etsi_ts/103700_103799/103759/02.02.01_60/ts_103759v020201p.pdf, both free, specify detectors for cooperative awareness messages and leave collective perception to a future version.

Attracting users with a single-player tool before a network exists is a recognised strategy in the platform-strategy literature, commonly phrased as coming for the tool and staying for the network. No novelty is claimed for it and no single origin is asserted.

The exchange rate between node density and privacy margin, and the argument that a dense ownerless network is a better observer than a sparse operator network, are from Part Eight of this series. The underlying relation is the two-way radar law, in which received power falls as the fourth power of range: doubling the distance costs 12 dB, so halving the mean spacing between nodes costs the same, and doubling areal node density, which shrinks that spacing by the square root of two, costs about 6 dB.