TECHNICAL PAPER

Flood monitoring

When the list of radios stops being a list a person can read.

Off Grid Pete LLC
31 October 2026

Abstract

One address has to stay one radio long enough for a person to find it on the list. A pairing handheld, a rack of tags, and ordinary address rotation can each turn that list into a crowd of new rows in a few seconds. Treating every new address as a device, deleting the crowd, or naming a tool would throw the record away or invent a cause the advertisement does not carry. This paper counts the burst, at one loudness, and asks whether the rows stay on the live list. The sit and the log keep them. The count is the finding.

1. A list you can no longer read

The useful thing about a live list is that a person can look at a radio. That only works while one address stays one radio for long enough to find it. A room can stop working that way in a few seconds. A handheld that wants to fill a phone’s pairing screen rotates its address and sends a pairing advertisement again and again. A rack of tags in a shop, each with its own address and a name, does something that looks the same from the door. Phones and tags also change their addresses in the ordinary course of not being followed. From the list, those three situations are a crowd of new rows, and the radio the person came to look at is under them.

Counting every new address as a new device makes the list worse. Deleting every new address throws away a real crowd of tags along with the burst, and it throws away the record of the afternoon. Naming the burst as a particular tool would be a third mistake. The advertisement does not say what sent it. A shop and a handheld can cross the same line, and a warning that picked one of them would be guessing.

This paper is the gate in front of the list. It counts new addresses over a few seconds, at about the same loudness, and it asks what to do with that burst. One answer leaves the rows up. The other takes them off the live list and leaves them in the sit and in the log. The count is the finding. The rest of the paper is which advertisements share a counter, why a drive past a shop is not the same as standing in the aisle, and why a Wi-Fi burst is not confirmed until the next scan shows that the room otherwise stayed.

2. The gate in front of the list

The gate does not decide that a particular product was present. It decides that the last few seconds produced a crowd of new addresses, at one loudness, and it asks what to do with them. Continue leaves them on the list. Hide these takes this burst, and later bursts in the same hold, off the live list. Neither answer deletes them from a sit or from the log.

3. Why a new address is not a new radio

Bluetooth Low Energy has used resolvable and random addresses for years, for ordinary privacy, not only for tricks. Becker, Li, and Starobinski showed in 2019 that address randomization does not, by itself, stop a passive observer from following a device, because the payload often stays recognizable. That result cuts both ways. A stable payload with a changing address can be one device trying not to be followed. A changing address with a pairing payload, repeated faster than any shop rotates its stock, can be one device trying to look like many.

Fieldwatch does not try to re-identify the device behind the rotation. It counts. The pairing counter and the name counter are separate because the two crowds mean different things on a walk. A pairing burst is the one that fills a phone with system popups. A name burst is the one that fills the list with titles and no popup. A factory address that keeps a stable name stays out of the name count. Pairing advertisements still count every address, factory or not, because the popup does not care how the address was minted.

Wi-Fi has the same habit. Vanhoef and colleagues, and the later measurements of probe and beacon behavior in the wild, are the background for the third counter: a scan that is suddenly full of names no one has heard, all about equally loud, all gone when the radio looks again. Martin, Mayberry, Donahue, and others have documented how far MAC randomization on phones does and does not go. The flood rule does not need that literature to fire. It needs the literature to stay modest about what a fire means.

4. The pairing counter

The window is about ten seconds. An address counts once in that window. The pairing families that share the counter are the ones a phone already treats as a tap-to-pair advertisement:

- Apple proximity pairing, Continuity type 0x07.
- Nearby Action, Continuity type 0x0F.
- A three-byte Fast Pair model id.
- A Swift Pair beacon. The Windows background beacon, company 0x0006 beginning with 0x01, stays out. It is not a pairing popup.
- Samsung Easy Setup for buds or a watch. Other payloads under that company identifier, including a SmartTag, stay out.
- A LoveSpouse advertisement under company 0x00FF that carries its long fixed prefix. A shorter blob under that company identifier stays out. Those bursts have shown up as a crowd of one loudness, which is why they are on this counter and not on a general company-id rule.

Find My, Continuity type 0x12, stays out, including a later proximity type inside that same payload. Counting it would turn a pocket full of ordinary tags into a flood.

Ten new pairing addresses in the window, at about the same loudness, is the line. Loudness is a band 12 dB either side of the median. A spread-out crowd, some in the room and some down the street, does not trip it. The families that actually fired are named on the dialog, as families, not as a product that must have sent them.

5. The name counter

The second counter is a new randomized address that advertises a name and is not already on the pairing counter. Fifteen such addresses, again within the 12 dB band, is the line. A factory address with a stable name is left out. Tags on a shelf can cross this line honestly. So can one radio that changes its name and its address on every packet. The dialog says both, and then it says the advertisement does not name the tool.

The on-phone text names handhelds that are able to do the second thing: a Flipper Zero, or an ESP32 running Marauder or Bruce. Naming them is a hint about a shape the person might recognize. It is not an identification. A warning that said “this is a Flipper” would be a different product, and a worse one. The same paragraph says a store can do it.

6. Driving through someone else’s crowd

A walk through a shop should still be able to warn. A drive past a shop should not warn on every aisle. Above 5 m/s, about 11 mph, a burst whose median is -80 dBm or weaker stays quiet. A louder burst still warns. The threshold is deliberately above a walking pace, so a person standing in the aisle is not treated as a car. The weak floor is there because a highway full of distant tags is not the event the dialog was written for.

7. The Wi-Fi beacon flood

Wi-Fi is not counted in a ten-second smear. It is counted on full scans. The first full scan of the process only fills a memory of names and addresses already heard. It cannot warn, because everything in it is new to the process. A partial scan, a cached scan, or a second callback inside three seconds is ignored.

A candidate is a set of brand-new names, each seen on a single address, on one channel, at least fifteen of them, each within 6 dB of the median loudness of that set. A repeated name does not count. A trailing guest, extender, or band suffix is stripped before the name is compared, so a mesh and its guest network are not fifteen networks. The whole name is kept when that suffix is all it is.

The candidate waits for the next full scan. It is confirmed only when two things are true. None of those new addresses are still present, and some other address that was already known is still present. The second test is the anchor. If the whole scan changed, the phone may have moved, or the scan may have failed, and the burst is dropped rather than confirmed. A real beacon flood is a set of names that were not there, then were, then were not, while the room otherwise stayed.

The dialog uses the same two buttons. The explanation names the same class of handheld, and the same limit: the advertisement does not name the tool. A mesh, an extender, and a guest network are called out as things this counter is built to ignore.

8. Hide, Continue, the sit, and the map

With no sit open, Continue or Hide these holds for fifteen minutes from the tap. The clock does not slide. When it ends, the burst already on screen stays quiet, and the next burst asks again. With a sit open, the answer holds until that sit ends. A fifteen-minute answer that is still running when a sit starts is carried into the sit.

Hide these removes the counted addresses from the live list, including addresses that join the same episode later. It does not remove them from the sit, and it does not remove them from the log. The sit stores a flood note: the time, which counter, the count, the families when the pairing line was the one crossed, and the median loudness. Debrief repeats that note and then leaves those addresses out of the presence counts and the lists below, so a flood does not become six hundred radios that “were here.” Compare names which sit the burst was in and applies the same omission. Older sits, saved before addresses were stored on the note, can only omit what they have, which may be nothing.

On a TAK map the flood is a line on this phone’s own marker, for as long as the notice lasts. A Wi-Fi beacon flood uses that same line. There is no flood pin, no per-address spot, and no ellipse. A pin would pretend the burst had a place to stand. It has a loudness and a count. BLE flood warnings are a switch on the TAK page. Off sends

nothing. The switch does not turn the red line on Live on or off.

The red line on Live is the flood’s line, and it has Hide this burst. A GNSS hit has its own red line and does not. The two notices can be up in the same hour. They are not the same event, and hiding one does not hide the other.

> **What a flood is not**
>
> It is not a finding that one radio did it. It is not a finding that a named handheld did it. A shop, a rack of tags, or a phone that is pairing in good faith can cross the same line. Hiding the burst does not delete it from the record. A quiet stretch is not a clean room. Addresses that were already known, a factory address with a stable name, Find My, the Windows background beacon, a mesh, and a guest network are outside the counters on purpose.

9. How a Bluetooth address enters the window

The scan loop hands the flood counter the first sighting of a Bluetooth Low Energy address. A later advertisement from the same address is not passed in. The counter ignores anything that is not BLE. What remains is classified, and a classification of nothing never enters the window.

Classification is a set of bits, read from the manufacturer data. Company 0x004C is walked as Continuity length-and-type fields. A type of 0x07 sets the proximity bit, unless a type of 0x12 has already been seen in that same payload, in which case the 0x07 belongs to Find My and is left unset. A type of 0x0F sets Nearby Action. A type of 0x12 only latches Find My. It is not itself a hit. Company 0x0006 is a Swift Pair beacon only when the payload is at least three bytes, the first is 0x03, the second is 0x00, 0x01, or 0x02, and the third is 0x80. A payload that starts with 0x01 does not match, which is how the Windows background beacon stays out. Company 0x0075 matches only when the payload starts with the prefix used for Easy Setup buds or the prefix used for an Easy Setup watch. Other payloads under that company, including a SmartTag, return false. Company 0x00FF matches only when the payload starts with one long fixed prefix. A shorter blob returns false. A Fast Pair model id already parsed on the observation sets its own bit. Any of those bits makes the sighting a pairing hit, and the family labels on the dialog are the bits that were set.

If no pairing bit is set, a non-blank name on a randomized address is a name hit. A factory address that carries a stable name is neither, and it is dropped. Pairing hits still count every address, factory or randomized, because the pairing test does not look at the address type.

A hit that survives is appended to a deque, and its address is added to a set, so the same address cannot be counted twice while it remains. The deque holds at most 96 addresses. The oldest is dropped when a 97th arrives. On every change, addresses older than ten seconds are removed from the front and from the set. The window is then split. Pairing hits are banded on their own. Name hits are banded on their own.

10. The band, the speed gate, and the episode

A band sorts the loudness values the radio actually measured and takes the middle one as the median. A hit more than 12 dB from that median is removed, and the count is what remains. If fewer than half the hits have a measured loudness, there is no median and the band is the whole list. A pairing band of 10 or more is hot. A name band of 15 or more is hot. Either one is then checked against speed. If the phone’s GPS speed is at least 5 m/s and the median is -80 dBm or weaker, the band is not hot. A missing speed does not suppress the band. A median louder than -80 dBm still warns while the phone is moving.

The first hot window opens an episode and appends one flood record: the time, the two counts, the family labels when the pairing line was the one crossed, the median, and the set of address keys. A later hot window in the same episode does not append another record. It rewrites the open one, keeping the earlier start, the higher of the two counts, the union of the family labels, and the growing key set. Keys are stored as BLE followed by the address. The live-hide set and the report set are capped, the report set at the sit’s radio cap and the episode set at 900, with the oldest key dropped past the cap. When the window is no longer hot, the episode ends, the notice is cleared, and the next hot window can ask again, subject to the hold taken from Continue or Hide these.

11. How a Wi-Fi scan becomes a confirmed burst

The Wi-Fi counter takes a whole scan, not a ten-second smear. Rows that are not fresh Wi-Fi, or whose address does not normalize, are dropped. Two observations of the same address keep the one with a measured loudness, and the louder one when both are measured. An empty parse returns without touching state. A callback inside three seconds of the last accepted scan returns the same way. The first accepted scan only records the addresses and the base names it saw. It cannot warn, because every name in it is new to the process.

On the next accepted scan, a candidate left by the previous scan is resolved before a new one is built. An anchor is any row whose address was already known and was not one of the candidate’s addresses. No anchor means the scan changed as a whole. The remembered addresses and names are replaced with this scan, the candidate is discarded, and nothing is confirmed. With an anchor, the candidate is confirmed only when none of its addresses are still present. If any of them stayed, it is discarded.

A new candidate is built from what this scan has that the memory does not. An address already seen is not fresh. A base name already seen is not fresh. The base name is the trimmed name with one trailing guest, guests, extender, ext, or band suffix removed, and only when that suffix is not the whole name. A blank name, and the placeholder for an unknown name, are not names. Names that appear on more than one address in this scan are removed, so one mesh does not become fifteen networks. What remains is grouped by channel. The largest group is kept, and a tie goes to the lower channel number. Inside that group, a loudness more than 6 dB from the median is removed. Fifteen or more survivors are stored as a pending candidate. They are not a warning. The warning is the confirmation on the following full scan, by the rule above. The dialog, the hold, and the hide set are the same objects the Bluetooth counter uses.

12. Parameters

Parameter	What the implementation uses
Bluetooth window	About 10 seconds. An address counts once.
Pairing line	10 new addresses in the pairing families, within 12 dB of the median.
Name line	15 new randomized addresses that advertise a name, within 12 dB. A factory address with a stable name stays out.
Left out of pairing	Find My. The Windows background beacon. SmartTag payloads that are not Easy Setup buds or a watch. A short 0x00FF advertisement.
While driving	At or above 5 m/s, a burst at or below -80 dBm stays quiet. Louder still warns.
Wi-Fi line	15 new single-address names, one channel, within 6 dB of the median, gone on the next full scan, while a previously heard address remains.
Wi-Fi ignored	The first full scan. A partial or cached scan. A repeat inside 3 seconds. A repeated name, a mesh, an extender, a guest suffix.
Hold	15 minutes with no sit, from the tap, not sliding. Until the sit ends when a sit is open.
Record	The sit and the log keep the addresses. Debrief and Compare leave a burst’s addresses out of the presence lists when the keys were stored.
TAK	A line on this phone’s marker. No pin.

References

1. Becker, J. K., Li, D., and Starobinski, D. “Tracking Anonymized Bluetooth Devices.” *Proceedings on Privacy Enhancing Technologies*, 2019. doi:10.2478/popets-2019-0036. Address randomization does not by itself make one device look like many, or many like one. The flood counters do not try to undo it.

2. Martin, J., Mayberry, T., Donahue, C., Foppe, L., Brown, L., Riggins, C., Rye, E. C., and Brown, D. “A Study of MAC Address Randomization in Mobile Devices and When it Fails.” *Proceedings on Privacy Enhancing Technologies*, 2017. doi:10.1515/popets-2017-0054. Background for why a new Wi-Fi address is not, on its own, a new network.

3. Vanhoef, M., Matte, C., Cunche, M., Cardoso, L., and Piessens, F. “Why MAC Address Randomization is not Enough: An Analysis of Wi-Fi Network Discovery Mechanisms.” *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security (ASIACCS)*, 2016. doi:10.1145/2897845.2897883.

4. Bluetooth Core Specification, Supplement to the Bluetooth Core Specification, and the public Fast Pair and Microsoft Swift Pair beacon formats. The pairing counter matches advertisement shapes a phone already treats as tap-to-pair. It does not decode a vendor’s private protocol beyond those shapes.

5. IEEE 802.11. Beacon frames and the scan results Android returns to an application. Fieldwatch uses a full fresh scan. It does not put the radio into monitor mode, and it does not capture traffic from clients.

A note on this paper

Flood monitoring as described here shipped with Fieldwatch well before the 1.1.37 release and remains in that release. The Wi-Fi beacon case uses the same dialog, the same hold, and the same TAK line as the Bluetooth cases. Customer-facing sentences stay with the count and the loudness. They do not promote a flood to a named tool, and this paper should not be read as doing so either.