TECHNICAL PAPER

Trusted networks

A familiar name on the Wi-Fi list, and what a scan can say before anyone joins.

Off Grid Pete LLC

3 November 2026

Abstract

The name on a Wi-Fi scan is chosen by whoever transmits the beacon, and a phone will offer to join a name that looks familiar. The scan does not say whether the access point is one this phone has heard for that name before, and joining is the step a person would rather not take to find out. For names a person has marked, this paper compares a later sighting with what was learned: weaker security, a maker that was not in the set, or a phone far from the places where that name is usually heard. A warning is that difference. A quiet watch is an empty report.

1. A familiar name is not a reason to join

People pick a Wi-Fi network by the name on the screen. At home, at a shop they use, on a team network, that name is familiar, and the phone will offer to join it. The name is a string in a beacon. Whoever transmits the beacon chooses the string. The scan result that shows the name does not say whether the access point behind it is one this phone has heard for that name before. Joining would answer a different question, and joining is the step a person would rather not take to find out.

The checks that would settle it are the ones this application does not do. It does not associate, it does not complete a handshake, and it does not time the beacon. A beacon can advertise a security suite it will not honor. An open network that was always open has no weaker setting left to notice. A copy that was learned, in the first days, as if it were the original, will then look like the original. A quiet watch is an empty report. It is not a finding that the name is safe, and it is not a finding that the name was checked.

What a receive-only scan can do is smaller, and it is only for names a person has marked. After a short window in which the access points of that name are learned, a later sighting can be shown when it does not look like that set. The security is weaker than the access points already learned, the maker was not among the makers already learned, or the phone is far from the places where that name is usually heard. A hotspot that travels with the person can be marked so the distance check stays quiet. This paper is that comparison. A warning is a difference from what this phone has heard. It is not proof of an attack, and the dialog does not name one.

2. What gets watched

Trust, on this page, means a person has marked that exact name as one to watch. It does not mean the name has been authenticated. The application does not join the network, does not ask for a password, and does not time the beacon. The watch starts off. Trust this network sits on a Wi-Fi radio's page, next to Mine, and it applies to that exact name. A hidden name, a blank name, and the placeholder “<unknown ssid>” cannot be trusted. A label can be set for display. The comparison is still the exact name.

3. What the literature already established

The copy of a known name is an old idea. By the mid-2000s it had a colloquial name, the evil twin, and a close cousin in the Karma attack, in which a radio answers for a name it has heard a client looking for. The defensive literature since then has split. One branch tries to fingerprint the access point: clock skew, beacon timing, information elements, the vendor's OUI. One branch tries to notice that the security got worse on the way from the real network to the copy, because a copy that cannot produce the real handshake often falls back to open, to WEP, or to a personal passphrase where the original was enterprise. A third branch simply asks where the phone is. A name that lives at home and appears in another city is interesting even when the beacon is well formed.

Fieldwatch uses a small slice of each idea, and only what the scan already returns. Security is parsed from the capability string Android puts on the scan result. The maker is the vendor prefix of a globally unique address, looked up the same way the rest of the application looks up a vendor, and only when the address is not locally administered. The place is this phone's GPS fix at the time of the sighting. There is no probe of the access point, no handshake, no timing analysis, and no claim that a matching beacon is the original.

802.11w, protected management frames, and WPA3 changed what a careful network can refuse. They did not change what a beacon can claim. A sighting that advertises WPA3 is recorded as WPA3. The watch does not test that the claim survives an association, because it never associates. An open network that was always open has no weaker setting left to find. That limit is structural. It is written on the settings page in plainer words.

4. What is stored when a name is trusted

Trusting a name stores the access points that were selected at that moment: address, vendor, the vendor prefix when it is a factory prefix, band, channel, the security marks parsed from the capability string, the raw capability text, and the loudest signal heard. A good GPS fix can also store where that loudest hear was. A good fix is tighter than 100 m, not paused, and inside the normal latitude and longitude bounds.

Places come later, during the learning window, from sightings of access points already known, with security that is not a downgrade. A place is a point. Points within 300 m are the same place, and the stored point moves to the newer fix. A new place is added when the phone is farther than the “away” distance, which starts at 1,000 m. Eight places are kept. A ninth replaces the one updated least recently. A network marked “This network travels” does not collect places, and the distance warning stays off for it. A hotspot in a second phone, or a shop radio that is in the truck, is the case that switch is for.

Ignored addresses are a list the person builds from a warning, by choosing not to be told about that address again. Episodes, the open and closed warnings, stay on the phone and out of the export.

5. Security, parsed as a small ladder

The capability string is reduced to a set of marks. Open is rank 0. WEP is rank 1. WPA personal is rank 2. WPA2 personal or enterprise is rank 3. WPA3, SAE, OWE, and Suite-B are rank 4. Enterprise is remembered as enterprise, personal as personal, OWE as its own mark. A sighting that says nothing the parser recognizes is ignored. The watch does not invent an open network out of a blank string.

A downgrade is a sighting that is genuinely weaker than the set already learned. Open or WEP, where the learned set has nothing that weak and the sighting has nothing stronger, is called open or weaker. A best rank below every learned rank is weaker. A personal sighting against a baseline that is entirely enterprise is weaker, even when the rank number would otherwise look fine. WPA2 personal where the house has only ever been WPA2-Enterprise is the case that rule exists for.

A sighting that matches, or that is stronger, is not a downgrade. During the learning window a matching sighting can add its security marks to the access point, so a network that advertises two suites is learned as both, rather than warned for being the second one.

6. The maker

A factory prefix is the first three bytes of an address whose local bit is clear. A locally administered address has no factory maker. Phones and many access points randomize, and a private address is not treated as an unknown vendor. There is nothing to disagree with.

Unknown maker requires a factory prefix, an address that is not already one of the learned access points, and a prefix that is not among the prefixes already stored for that name. It is on by default. It can be turned off.

The prefix 00:13:37 is a special case, and a narrow one. It does not, by itself, raise the unknown-maker warning. If some other reason has already fired, a second line says that 00:13:37 belongs to Orient Power Home Network. The screen does not name a tool. The line is there so a person who already has a reason to look is not left with a bare prefix that happens to be widely recognized. It is not a detector for that prefix.

7. Far from the places this name is heard

Away is the third reason. It requires the distance check to be on, the name not marked as traveling, at least one learned place, and a good fix. The nearest learned place has to be farther than the away distance. That distance starts at 1,000 m, and it can be changed. A name trusted at home and heard across town can warn. A name trusted and heard in the same house does not, even when the signal is strong.

Being merely away is a medium warning. The dialog can offer to remember this place, which adds it and closes the episode. A network that travels does not get that button. There is no place to remember.

Loudness can raise a warning one step, and only one step. If this sighting is at least 15 dB louder than the loudest stored hear of a known access point within 300 m of the same place, and the fix is good, medium becomes high. The idea is modest. A copy set down next to the phone is often louder than the access point down the hall ever was. It is not a range. RSSI is not meters.

8. The learning window

A new name should not warn on its own second access point. The learning window starts when the name is trusted and lasts three days, unless that length is changed. The short choice is 1 hour, for a name whose other radios are already in the building. The name row says how long learning has left, in days, then hours, then minutes. The choices do not list the names. Pause learning does not stop that clock. Learning can also be paused on a name, which freezes the set. Restore and export do not invent a new window. The clock is the moment of trust.

During the window, a new address is stored without a warning when the security is not a downgrade, the fix is good, and the phone is within 300 m of a place already known for that name. A private address is stored immediately. So is a factory address that shares its first five bytes with one already learned, the pattern a vendor uses for a second radio in the same box. A different address that only shares the three-byte prefix needs two fresh sightings, and a quiet stale sighting does not count toward those two. Quiet, for this purpose, is below the loudness floor, which starts at -85 dBm. After the window, new addresses are not learned by themselves. They are candidates for a warning, or they are nothing.

An address that is already known is updated: last seen, a louder signal, and, when the security still matches, any new security marks. Its place can move during the window. It is not re-warned for being itself.

9. When the dialog opens

A sighting below the loudness floor, or a sighting Android marked as not fresh, does not open a warning. It can still touch the last-heard time of an episode that is already open. An ignored address never opens one.

The reasons are collected only when the watch is on. Empty security collects none. Then the sighting has to be seen twice, fresh and loud enough, unless an episode for that address is already open. One scan result is not a warning. The second one can be.

Confidence starts at medium. Open, or a drop all the way to WEP or open against a baseline that was never that weak, starts at high. The 15 dB rule can add one step and does not go past high. Warning level, the same Low / Medium / High control used elsewhere, defaults to medium. A confidence below that floor is kept as a pending count and does not interrupt. A rise later can.

The sentence names the network and the reasons: open, weaker, an unknown maker, away. Beep and voice start on. Voice speaks that sentence. System notification starts off and does not include a coordinate. The dialog can offer Learn this access point, Ignore this address, or Remember this place, depending on whether the address is new and whether the only reason was distance. Learning writes the address into the set and ends the episode. Ignoring adds it to the ignored list. Dismiss closes the dialog and holds that address quiet until it has been gone for ten minutes. A dismissed address that is still being heard stays dismissed. The hold refreshes while it is heard.

One network can have several episodes. They are capped, and a GNSS dialog or another notice still owns the screen first.

10. The sit, the export, and the pin

A sit report lists a trusted-network warning that overlapped the sit: the name, the reason, the time, and the address. Coordinates stay off that row. Compare does not grow a second copy of the feature. The log is unchanged.

Export trusted networks writes a JSON file. The suggested name is fieldwatch-trusted.json. The file keeps names, access points, security, the travels switch, and ignored addresses. It drops places, episodes, and the coordinates of the loudest hear. Import adds names and access points onto this phone. A name that is already here keeps its places. New access points and ignored addresses are merged in. A broken file is refused and the phone's current set is left as it was. Export settings, the general settings pack, includes the switches and does not include the names. The two files are on purpose. A settings pack moving to a new phone should not silently carry someone's home name, and a trusted-network file should not silently carry someone's places.

Network warnings on the TAK page can drop a pin where this phone was when it heard the sighting. The pin and its hold stay dim until Watch trusted networks is on. The callsign is “Possible evil twin:” plus the name, cut at 32 characters. The pin moves when a later hear is louder. A weaker hear keeps the old point and can refresh in place. Turning the pin switch off removes pins this phone sent. Turning the watch off does not, by itself, take a pin down. Turning Network warnings off does. There is no ellipse. The pin is a hear, marked as a hear, with the same honesty problem as every other hear-here pin: it is where this phone was, not a surveyed location of the access point.

> **A quiet watch**
>
> A quiet watch is not a finding that the name is safe. A matching copy of an access point already learned stays quiet. An unmarked name stays quiet. A radio that was never heard stays quiet. A private-address copy, sitting at the table, often stays quiet too, because there is no factory maker to disagree and the security may look the same. An open shop has no weaker security left to find. A dialog is not proof of an attack. It is proof that this sighting differed from what the phone had stored, in one of the three ways above.

11. One sighting, two passes

A trusted name is an exact SSID. Trusting it stores the access points that were checked at that moment, and the clock starts then. A hidden name, a blank name, and the placeholder for an unknown name are refused before a record exists. Later, each Wi-Fi sighting whose name matches a stored record is handed to one function. That function ages the book first. An open episode that has not been heard for ten minutes is closed. An ignored address returns immediately. Nothing is learned and nothing is tested.

What remains goes through learn, and then through detect, in that order. Detect sees the record after learn has already written. The record for one name holds the access points, the places, the ignored addresses, two pending maps, the episodes, and the travels and pause flags.

12. How an access point is learned

An address already in the set updates its last-seen time. The capability string is parsed, and the new marks are added to the stored marks only when the parse is non-empty and is not a downgrade against every mark already stored for the name. A louder hear updates the stored loudest value. A good fix on that louder hear also stores the latitude and longitude. A good fix is not paused, has an accuracy under 100 m, and has a finite latitude and longitude. During the learning window, a known access point with a good fix and a security parse that is not a downgrade can also move a place that is within 300 m, or add one when the hear is farther than the away distance from every stored place.

An address that is not in the set is eligible only while the window is open and learning is not paused on that name. The window is the number of days on the control, counted from the moment the name was trusted, at 86,400,000 milliseconds a day. The sighting has to parse, it has to not be a downgrade, the fix has to be good, and the hear has to be within 300 m of some stored place. A private address is then stored at once. So is a factory address that shares its first five bytes with an address already stored and is not itself locally administered. A different address that only shares the three-byte prefix is counted in a pending map. A fresh sighting increments the count. A quiet stale sighting does not. Two fresh sightings store it and clear the count. A different prefix is not stored. After the window, an unknown address leaves the record unchanged.

13. How the capability string becomes a rank

The parser scans the capability text for a fixed set of tokens. WPA3, SAE, OWE, and SUITE-B are the top rank. WPA2 and RSN are the rank under that. WPA without those is the rank under that. 802.1X or EAP marks the result enterprise. PSK marks it personal. SAE is stored as personal at the top rank. SUITE-B is stored as enterprise at the top rank. OWE is stored on its own at the top rank. WEP is the rank above open. OPEN, NONE, ESS, or IBSS, and only when no crypto token was present, is open. A string that matches none of this is an empty set. An empty set is not treated as open. It is ignored.

A downgrade is computed against the set of every mark stored on the name, not against one access point. It returns a reason in three cases only. The sighting contains open or WEP, the stored set has nothing that weak, and the sighting itself has nothing of rank 2 or higher. Or the sighting's best rank is below every stored rank. Or the stored set is entirely enterprise and the sighting is personal. Any other combination, including a blank sighting or a blank baseline, is not a downgrade.

14. How a reason becomes a dialog

Detect looks first for a dismissed episode on this address whose last hear is under ten minutes old. That episode's last-heard time is refreshed and the function returns with no alert. A hear below the loudness floor, or a hear the scan marked as not fresh, updates the last-seen time of a known access point and returns. The floor starts at -85 dBm.

The reason list is built only when the watch is on and the parse is non-empty. Weaker security is the downgrade result, when that switch is on. Unknown maker requires its switch, an address that is not already stored, a factory prefix, a prefix other than 00:13:37, and a prefix that is not already among the stored prefixes. A locally administered address has no factory prefix, so it cannot take this reason. The prefix 00:13:37 does not take it either. If some other reason has already been collected, the episode records that the prefix was 00:13:37, and the dialog adds the second line. Away requires its switch, a name that is not marked as traveling, at least one place, a good fix, and a distance from the nearest place greater than the away setting. The setting starts at 1,000 m. The 300 m radius used for learning is a different test and is not this one.

An empty reason list updates the last-seen time and returns. Otherwise a fresh hear increments a pending count for the address. The dialog needs two such hears, unless an episode for the address is already open. One scan is not enough. Confidence starts at medium. Open, or a weaker result whose sighting includes WEP or open against a baseline that has neither, starts at high. A hear 15 dB louder than the loudest stored hear of a known access point within 300 m of the nearest place adds one step, and the result is capped at high. Below the alert floor, the pending count is kept and no dialog is built.

When the dialog is built, the episode stores the reasons, the confidence, the louder of the hears, and the fix from the louder good hear. Learn this access point and Ignore this address are offered when the address is not in the set. Remember this place is offered only when the address is already known, the only reason is away, the name does not travel, and the episode has a latitude and longitude. Dismiss marks the episode dismissed and closed. A later hear inside ten minutes refreshes that hold instead of opening again. Learn writes the episode back through the same access-point constructor and closes it. Ignore adds the address to the ignored list, which is the check at the top of the next sighting.

15. Parameters

Parameter	What the implementation uses
Watch trusted networks	Off until turned on. With it off, nothing alerts and no new pin is sent.
Reasons, all on by default	Weaker security. Unknown factory maker. Away from learned places.
Away distance	1,000 m. Ignored when This network travels is on.
Here	300 m. A hear inside this radius is at a place already known for that name.
Learning window	3 days from the moment the name was trusted. Also 1 hour, or 1, 7, or 14 days. Can be paused per name.
Loudness floor	-85 dBm. Below that, no new warning. A stale scan result does not count.
Two sightings	A new address needs two fresh sightings above the floor before the dialog, unless an episode is already open.
Confidence	Medium by default. Open, or a fall to WEP or open, starts high. 15 dB louder than the stored hear at that place can add one step.
Dismiss	Closed until that address has been gone for 10 minutes.
Places	Up to 8. A ninth replaces the least recently updated. Not included in the export.
00:13:37	Does not alert as an unknown maker by itself. If another reason fired, a second line names Orient Power Home Network. The screen does not name a tool.
TAK pin	Off until Network warnings is on, and dim until the watch is on. Callsign begins “Possible evil twin:”. Moves only when the hear is louder.

References

1. Dai Zovi, D. A., and Macaulay, S. A. “Attacking Automatic Wireless Network Selection.” *Proceedings from the Sixth Annual IEEE SMC Information Assurance Workshop*, West Point, June 2005. doi:10.1109/IAW.2005.1495975. The paper behind the Karma attack: a radio answers for a name a client is willing to join. Fieldwatch does not transmit, and it does not test association.

2. IEEE 802.11-2020, and the Wi-Fi Alliance WPA3 and Enhanced Open (OWE) specifications. The security ladder in this watch is a reading of the capability string, not a test that the named suite is real.

3. IEEE 802.11w. Protected management frames. Relevant because a beacon can advertise a suite the watch never challenges.

4. Android Open Source Project. Wi-Fi scan results, including the capabilities string, BSSID, frequency, and RSSI, as returned to an application that does not associate.

5. Orient Power Home Network, OUI 00:13:37, in the public IEEE registration. Named on the Fieldwatch dialog only when another reason has already fired. The dialog does not name a tool.

A note on this paper

Trusted networks as described here are the ones in Fieldwatch 1.1.39. The 00:13:37 line shipped in 1.1.36. The one-hour learning choice shipped in 1.1.38. The feature does not protect the network, and it does not protect the person standing near a copy that was learned, during the window, as if it were the original. That last case is the one the learning window makes possible, and it is why the window is short and why learning can be paused.