THE EYES BESIDE YOU
FROM TRACKERS TO PHONES

CLOSE-RANGE SURVEILLANCE – DETECTION AND RESPONSE

The Eyes Beside You

Detecting and Responding to Close-Range Surveillance

Odysec Technical White Paper WP-001 | **Version** 1.0 | **Published** 2026-08-30 **Domain**: Cyber × Physical

> This is a translation of the Chinese original, which is authoritative.

Table of Contents

About this white paper ... 3

Executive summary ... 3

1. Problem statement ... 4

2. A unified threat model ... 6

3. The four vectors of surveillance ... 7

4. Structural gaps in current protections ... 11

5. A layered detection methodology ... 14

6. Network-side observation: periodic reporting in encrypted traffic ... 19

7. What to do after finding surveillance ... 24

8. Recommendations for vendors and policymakers ... 27

9. Limitations and further work ... 28

Appendix A: Detection checklist ... 30

Appendix B: Terminology ... 32

Appendix C: References ... 33

Appendix D: Experimental environment and data ... 34

Revision history ... 35

About this white paper

Close-range surveillance can be achieved through four technically distinct vectors, and the person under surveillance does not know in advance which kind they are facing — they only know that "the other person always knows". **The four vectors have to be compared on one table before a reader can know what the negative check they just performed actually ruled out.** This white paper is built around that table, and adds the chapter such guidance usually lacks: **what the network side can see when the device under examination cannot itself be trusted**.

This white paper is of two kinds, which support inferences of different strength:

- **Chapters 2 to 5 and 7 to 8** are **synthesis and methodology**: they collate published academic measurement, standards documents and platform material, and build a self-executable detection procedure on top of them.
- **Can support**: by which routes surveillance can be achieved, which structural gaps exist in current protections at the level of principle, which vector each layer of checking covers, and what a negative result means.
- **Cannot support**: any claim about how a given platform or product behaves *today*. Wherever we cite someone else's measurement, we state its measurer, its subject and its date — **that is what one study observed at one point in time, not a guarantee about the present**.
- **Chapter 6** is **our own measurement**: in a reproducible controlled environment we measured how detectable periodic reporting is inside encrypted traffic. That chapter's data, its experimental environment and its analysis scripts are published with this white paper, and anyone may reproduce them.

The two must not be mixed. Chapter 6 measures a **behavioural pattern** (periodic reporting), not any real monitoring product; the gap analysis in chapters 2 to 5 is reasoning at the level of principle, not an evaluation of any vendor.

Executive summary

The technical barrier to close-range surveillance — sustained watching by a partner, an ex-partner, a family member or an acquaintance — has fallen to consumer level. The abuser needs no technical ability at all: one moment of physical access, or one password they already know.

Four principal findings:

1. **One negative check rules out at most one of the four vectors.** Surveillance can be achieved by four vectors: a physical tracker, the account layer, the device layer, and the settings layer. They are detected in **completely different places**, and three of the four are permanently negative under a

phone scan. The reassurance bought by a clean scan report is, most of the time, the product of looking in the wrong place.

2. **Current protections fail not because of implementation quality but because of their premises.** Automatic alerts assume the person under surveillance carries a compatible phone; scanning engines assume malicious and legitimate software can be told apart by behaviour; security notices assume the recipient is the person being protected. In intimate-partner surveillance — the principal abuse scenario — all three assumptions fail at once: the abuser lives in the same home, shares the bill, and is often the person who bought that phone.
3. **The act of checking is itself risky, so the order cannot be reversed.** If the phone is already monitored, every attempt to check and every request for help made on it may appear before the abuser in real time. The methodology must begin by moving to a clean device, not by scanning. By the same logic, removing surveillance the moment it is found announces to the other person that they have been detected — and that is the point of greatest danger.
4. **When the device cannot be trusted, the network side remains an independent vantage point — but its limits can be measured.** In a controlled environment we observed thirty windows for each of nine combinations of reporting period and jitter (chapter 6). **The shape of the conclusion is this: the timing criterion is defeated by jitter, and the traffic-direction criterion is not — but the direction criterion's zero false positives in this experiment are a product of how the background traffic was designed, and are not enough to support its use on a real network.** The figures and their limits are in that chapter.

Recommended action: if you suspect you are under surveillance, do not search for help on the phone you suspect; use someone else's device or a public one. Having found any sign of surveillance, do not remove or discard anything immediately — assess the risk to your safety first, then complete evidence preservation.

---

1. Problem statement

1.1 One threat, four vectors

What the person under surveillance experiences is a single phenomenon: **the other person always knows.** Knows where they went, who they met, what they said, when they left home.

Behind that phenomenon are four technically distinct routes:

Vector	What the abuser needs	What it leaks
A: physical tracker	One opportunity to place a device in a belonging or a vehicle	Location
B: account layer	The password to a cloud account	Location, photos, backups, contacts — and on some platforms messages
C: device layer	One opportunity with the phone unlocked	Almost everything — message content, calls, the screen itself
D: settings layer and peripherals	One opportunity to "set it up for you"	Depends on the setting, and **replacing the phone usually does not help**

The four share one threat model (chapter 2) but require four different places to look. **That is precisely why this white paper treats all four together**: check only one of them and a reader easily concludes "I have checked", when the other vectors were never touched at all.

1.2 The situation in Taiwan

Taiwan's Stalking and Harassment Prevention Act took effect in 2022; the conduct it covers includes persistent monitoring, following or ascertaining the whereabouts of a specific person by electronic communications, the internet or other equipment. The Offences Against Computer Security and the Offences Against Privacy chapters of the Criminal Code likewise address unauthorised intrusion into another's computer equipment and unauthorised recording of another's non-public activities.

The legal tools exist. What is missing is the other end: **the law defines what conduct is illegal, but teaches no one how to discover that it is happening.** The practical questions victims face are:

- I suspect I am under surveillance — how do I prove it?
- My phone has raised no alert and an anti-virus scan found nothing — does that mean I am safe?
- I have found a suspicious device or application — is removing it enough?

The intuitive answer to the third question is wrong, and wrong in a way that may endanger personal safety (see section 7.1).

1.3 Questions this white paper answers

1. How the surveillance happens (a unified threat model and four technical vectors)
2. Why current protections fail in this situation (structural gaps, not implementation flaws)
3. How to detect it yourself (a layered methodology, what each layer covers, and what a negative result means)
4. What can still be observed when the device cannot be trusted (network-side measurement, including our own results)
5. What to do after finding it (personal safety, evidence preservation, legal routes, clean-up order)

2. A unified threat model

2.1 Types of attacker

Type	Motive	Characteristics
Intimate-partner abuser	Control; monitoring movements and contacts	**The largest group**: sustained physical access, and often knows or can obtain the unlock code
Party to a separation or divorce	Gathering material; continued control	Access established during the relationship often survives the separation
Surveillance within a family	Control in the name of care	Parental-monitoring tools used on adult or near-adult family members
Workplace overreach	Crossing the line in the name of management	Company devices or mobile-device-management privileges used beyond work
Pre-burglary reconnaissance	Learning occupants' routines	Targets property rather than a person; trackers usually placed on vehicles

This white paper addresses mainly the first two, because they combine **the greatest access** with **the most serious physical consequences**.

2.2 Capability assumptions: technical skill of zero

This is the point of the threat model most in need of emphasis: **the attacker's technical ability is assumed to be zero.**

No code, no jailbreaking or rooting, no understanding of Bluetooth protocols or operating systems. All they need is **any one** of the following:

- one opportunity to slip a tracker into the other person's belongings or vehicle;
- one opportunity with the phone unlocked, or knowledge of the unlock code;
- the password to a cloud account, or the ability to pass its password-reset process;
- having once handled the device under the guise of "getting your phone set up".

In an intimate relationship these are everyday occurrences. **Knowing a partner's unlock code is treated in most relationships as a sign of trust — the threat model has to acknowledge that social reality** rather than assume the code is secret.

Precisely because the bar is that low, this class of threat is far more prevalent than attacks requiring skill. Defensive thinking must not assume a sophisticated adversary — the correct assumption is an adversary who understands nothing technical and has abundant access.

2.3 The victim's position

- **Long unaware period**: surveillance may have run for weeks or months before an alert or a suspicion.
- **Knowledge asymmetry**: the abuser knows what was planted and what was configured; the victim does not even know what to look for.
- **The channel for checking is itself monitored**: the most natural reaction is to search for what to do on one's phone — and that phone may be the monitored device.
- **Consent and coercion are indistinguishable**: in a controlling relationship, "voluntarily" handing over a password or "voluntarily" installing a location app may happen under pressure. The technical mechanism sees one lawful authorisation; it cannot see the coercion behind it.
- **Several points lost at once**: an abuser may hold the phone, the cloud account, the email account and the car's services simultaneously. Cleaning up any single point may be noticed and reversed from another.

2.4 One principle running through the whole document

That position converges on a single principle, and it determines the order of every procedure in this white paper:

> **The device under examination cannot be assumed to be a trustworthy tool.** In the tracker case the phone is usually still usable as a tool; in the other three vectors that phone is itself the subject of the investigation. Any check performed on the device under examination is untrustworthy whether it comes back positive or negative — and the act of checking may itself be seen.

---

3. The four vectors of surveillance

3.1 Vector A: physical trackers

Low-cost Bluetooth trackers were designed to find lost belongings, but three of their properties — negligible price, battery life measured in years, and a global crowd-sourced network that removes any need for their own connectivity — also make them ideal stalking tools.

A tracker does not know where it is. It has no GPS and no cellular connection. It does exactly one thing: continuously broadcast an identifier over Bluetooth Low Energy. Locating happens on other people's devices:

1. the tracker broadcasts its identifier;
2. any passing phone belonging to the same ecosystem hears it;
3. that phone uses its own GPS to determine "I am here, and I heard this identifier";
4. the pairing is encrypted and uploaded to the vendor's servers;

5. the tracker's owner queries and obtains the decrypted location.

The relaying phone cannot know what it relayed — location reports are end-to-end encrypted and only the owner can read them. For ordinary users' privacy, this is good design.

The double edge of rotating identifiers: to prevent a tracker becoming a long-lived trackable beacon, devices in each ecosystem rotate their broadcast identifiers periodically. This protects ordinary users — a bystander cannot log your movements by following one fixed identifier. But for a **person being tracked**, the same design creates difficulty: you cannot conclude "this identifier keeps following me", because the identifier changes; and a generic Bluetooth scanner sees a stream of ever-changing anonymous devices, hard to distinguish from the sea of other Bluetooth gear around you.

This is a classic privacy trade-off: a mechanism that protects the many from being tracked makes self-rescue harder for the few who are. Understanding it explains why self-detection cannot rely on generic scanning apps alone.

Network density differs enormously between ecosystems, and **density directly determines precision and update frequency**: in a device-dense city, tracking is far closer to real time than in sparsely covered areas.

3.2 Vector B: the account layer (nothing installed on the device)

With the cloud-account password, an abuser can sign in from their own device and obtain everything the account synchronises: location, photos, backup contents, contacts, calendar, and depending on the platform possibly messages too.

Three properties make this vector extremely favourable to the abuser:

- **A device scan will always be negative**: there is nothing extra on the phone;
- **It persists**: as long as the password is unchanged and the session is not signed out, the access continues;
- **The source is hard to notice**: few people check the list of signed-in devices on their accounts.

In addition, a platform's built-in **location sharing** and **family sharing** features, switched on during the relationship and never switched off as it deteriorated, amount to a forgotten surveillance channel.

3.3 Vector C: the device layer (requires one moment of physical access)

That is, installing monitoring software on the phone. The two platforms differ:

- **Android**: permits installation from outside the store. Once installed, monitoring software typically requests high privileges such as **accessibility services** and **device administrator** — mechanisms designed for assistive use and enterprise management, turned around to read screen contents, intercept notifications and prevent removal. After installation such software usually hides its icon or disguises itself as a system component.
- **iOS**: without jailbreaking it is difficult to sideload software that persists, so surveillance in the iOS case usually takes vector B or vector D instead. A jailbroken device is as exposed as Android, and

the jailbreak itself usually leaves identifiable traces: an unofficial store interface appears, system updates keep failing, or the device behaves noticeably unlike others of the same model.

Device-layer software falls into two classes whose legal and detection positions differ sharply:

Class	How it is obtained	Detection position
Dual-use applications	Lawfully listed in the official store as parental monitoring, anti-theft or employee management	**The mainstream case**: the software is lawful, and a scanning engine has no legitimate reason to flag it
Dedicated monitoring suites	Sideloaded, marketed explicitly or implicitly for watching a partner	More readily flagged by security products, but hides itself once installed

This white paper deliberately does not name any monitoring product, its source, or how it is installed. Such information helps abusers far more than victims: what a victim needs is how to find something wrong on their own device, not a catalogue of what is on the market. This choice follows the treatment of offensive detail set out in our Research Ethics Code.

3.4 Vector D: the settings layer and peripherals

The third route is neither a leaked password nor installed software, but **the misuse of the device's own legitimate settings**. It is the easiest to overlook, because each setting on its own is a reasonable feature:

- **Configuration profiles and work profiles**: mechanisms designed for enterprise device management which, once installed, can reach traffic routing, application delivery and some usage records. Installing one under the guise of "getting your phone set up" is unremarkable within a relationship.
- **Message and call forwarding rules**: automatic forwarding in email, call diversion with the carrier, or syncing text messages to another device. Once established, such settings **persist even if the phone is replaced**.
- **Paired peripherals**: watches, in-car systems, headphones, tablets. Pairings often survive the deterioration of a relationship, and in-car systems deserve particular attention — they hold both location history and contacts.
- **Desktop and browser sync**: phone messages synced to a shared computer, or a browser account syncing bookmarks and passwords, can each form a bypass that never touches the phone.
- **Built-in sharing features**: location sharing, family groups, shared albums. These are designed on the assumption of continuing mutual consent, but **no mechanism prompts you to reconfirm when the relationship changes**.

3.5 The four vectors compared

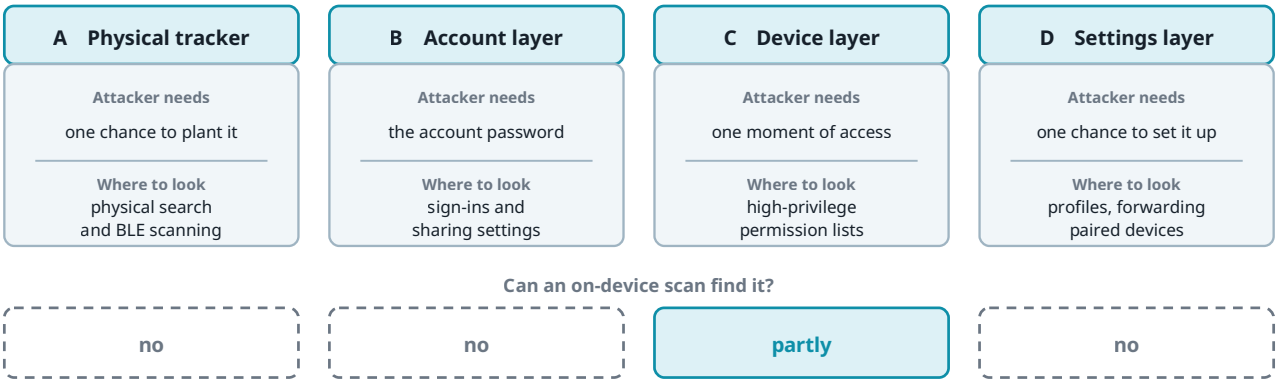

Figure 1: Four surveillance vectors and where each is detected. For three of the four, a phone scan cannot see it in principle.

This table is the page of this white paper most worth remembering. The crucial difference between the four is not technical sophistication but **where you should look**:

	A: physical tracker	B: account layer	C: device layer	D: settings layer
Physical access needed	Yes (once)	No	Yes (once)	Usually
Findable by a phone scan	**No**	**No**	Partly	**No**
Does the system alert on its own	Partly (depends on ecosystem)	Partly (sign-in notices)	Rarely	**No**
Where it is detected	Physical search and Bluetooth scanning	Sign-in records and sharing settings	The device's high-privilege lists	Device settings and external service rules
What it leaks	Location	Location, backups, photo library, contacts	Almost everything	Depends on the setting
How it is removed	Take the device out (preserve evidence first)	Revoke sessions and change credentials	Remove the software and revoke privileges	Review and delete settings one by one
Does replacing the phone help	No (the tracker is on the object)	**No** (it follows the account if untouched)	Yes	**No** (most rules live outside the device)

The last row is the most common misconception in practice: **a new phone does not end three of the four vectors.** If only the device is replaced and the account is left untouched, the new device is exposed the moment it signs in.

The second row explains why "I scanned it" cannot be a conclusion: **for three of the four vectors, a phone scan cannot see anything, as a matter of principle.**

---

4. Structural gaps in current protections

The conclusion first: **current mechanisms work for typical users in typical situations — but the close-range abuser is not a typical situation.**

4.1 Current mechanisms follow three lines

1. **Automatic alerts for unknown trackers**: a notification when a tracker that does not belong to you has been travelling with you for an extended period. A cross-platform detection standard (DULT) is being advanced jointly by the major vendors at the IETF. Most trackers also emit a sound after prolonged separation from their owner.
2. **Controlling software provenance and making permissions visible**: store review, restrictions on installation from outside the store, explicit authorisation for high-risk permissions, and an interface for reviewing permission history.
3. **Account activity notices**: notifications on sign-in from a new device, and a list of signed-in devices.

All three are effective against the threat model of a stranger planting malware. The difficulty is that the threat model in this white paper is not a stranger.

> **How this white paper handles concrete figures**: trigger conditions, delays and coverage change with OS version, region and vendor policy. Our approach is not to omit numbers but to **quote only figures backed by published measurement, always with the measurer, the subject and the date attached**. The reasoning: omitting numbers leaves the reader unable to estimate the order of magnitude of the risk — "there is a delay" and "the delay is eight hours" mean entirely different things for a safety plan — while an undated number gets read as the present state. Both are harmful; the dated number is the lesser harm. **If what you need is current behaviour, the only reliable source is your own device.**

4.2 Five structural gaps

The following five are not implementation flaws but distances between a mechanism's premises and the abuse it meets. No parameter adjustment resolves them.

Gap one: the premises exclude the people most in need of protection

Automatic alerts require three premises at once: the person under surveillance **carries a smartphone**, **Bluetooth is on**, and **the OS version supports the feature**. Those outside the premises receive effectively no protection: elderly feature-phone users, children without phones, people who disable Bluetooth to save power or by privacy habit, and users on old OS versions.

These groups overlap substantially with the groups at highest risk of stalking.

The same class of problem appears in another form at the device layer: where the phone was bought, configured or "sorted out" by the abuser, its chain of trust was incomplete from the outset. No mechanism today lets a user establish after the fact what was done to a phone before it reached their hands.

Gap two: the attacker chooses the coverage of detection

A standard's existence does not equal full coverage. DULT now has an IETF working group producing a threat model and an accessory-protocol draft, but a standard describes what ought to be; actual detection capability depends on what each vendor has implemented and shipped on each platform.

From the attacker's perspective this implies an evasion requiring no technical knowledge at all: **choose a tracker from outside the victim's phone ecosystem.** One academic response to this gap is cross-ecosystem third-party detection (AirGuard, published by Heinrich et al. in 2022, is one example: it reverse-engineered iOS's tracking-protection logic so that Android users could detect trackers in the Apple ecosystem). **The existence of such tools is itself evidence of the gap** — were first-party coverage sufficient, they would be unnecessary.

At a deeper level: detection relies on a device behaving as the standard prescribes, and **whether it does so is the device's own choice**. Devices with their speakers removed circulate in the market, as do self-built tags participating in no standard; and Shafqat et al. demonstrated in 2023 a cloned tracker that kept reporting location while evading victim-side alerts. **The mechanism's effectiveness rests on the attacker's cooperation, which is not an assumption a threat model may make.**

Gap three: lawful and coerced authorisation are technically indistinguishable

The feature list of a parental-monitoring application is almost identical to that of a monitoring application. The difference lies not in what the software does but in **whether the device's holder knows and consents** — a fact outside the code, which a scanning engine cannot determine in principle. A vendor that flags all of them harms legitimate use; a vendor that flags none of them admits the primary tool used in stalking cases.

This is not a marginal case. Chatterjee et al.'s 2018 study, the first systematic measurement of this ecosystem, found that **the great majority of applications usable for intimate-partner surveillance are dual-use tools with legitimate purposes, rather than overt spyware**; Roundy et al. in 2020 further showed that the software used for interpersonal attacks extends well beyond surveillance alone. In other words, what a scanning engine faces is not a handful of hard-to-classify exceptions but the principal form of the problem.

The system sees a lawful authorisation: a prompt appeared, the user pressed accept. What the system cannot see is that another person was standing beside them at the time. **In a controlling relationship, "voluntary" is a concept the technical layer cannot verify.**

The same dilemma appears on the tracker side as the co-travel problem: judging too loosely produces floods of false positives (family in the same car, strangers on long public-transport rides), and users quickly learn to ignore alerts; judging too strictly produces missed detections, **above all where victim and abuser share the same space — which is exactly the usual shape of intimate-partner abuse. When the two live together or share a car, "abnormal co-travel" loses its discriminating power entirely.**

Gap four: the recipient of the alert may be the abuser

If the phone number, the bill or the family group is controlled by the abuser, then security notices, sign-in alerts and subscription statements arrive in their hands. More fundamentally, the device itself may have been bought and set up by the abuser. **At that point every protection premised on "notify the user" has its protected party inverted.**

Gap five: protections assume that discovery should be followed by removal

Current safety prompts almost invariably advise removing suspicious software or devices at once. Under a stranger threat model that is correct; in close-range surveillance, a sudden halt in reporting announces to the abuser that they have been detected, and may trigger retaliation directly.

Between detection and removal there needs to be an intermediate step that current mechanisms do not provide at all: safety assessment and evidence preservation. Chapter 7 of this white paper is that step.

4.3 Gaps mapped onto vectors

Gap	A tracker	B account	C device	D settings
One: premises exclude people	●	○	●	○
Two: coverage chosen by the attacker	●	○	●	●
Three: authorisation context invisible	●	●	●	●
Four: alert recipient inverted	●	●	●	●
Five: assumes discovery means removal	●	●	●	●

● = applies directly ○ = applies partly

Gaps three, four and five span all four vectors, because their root is not technical: **the threat model assumed the wrong attacker.**

5. A layered detection methodology

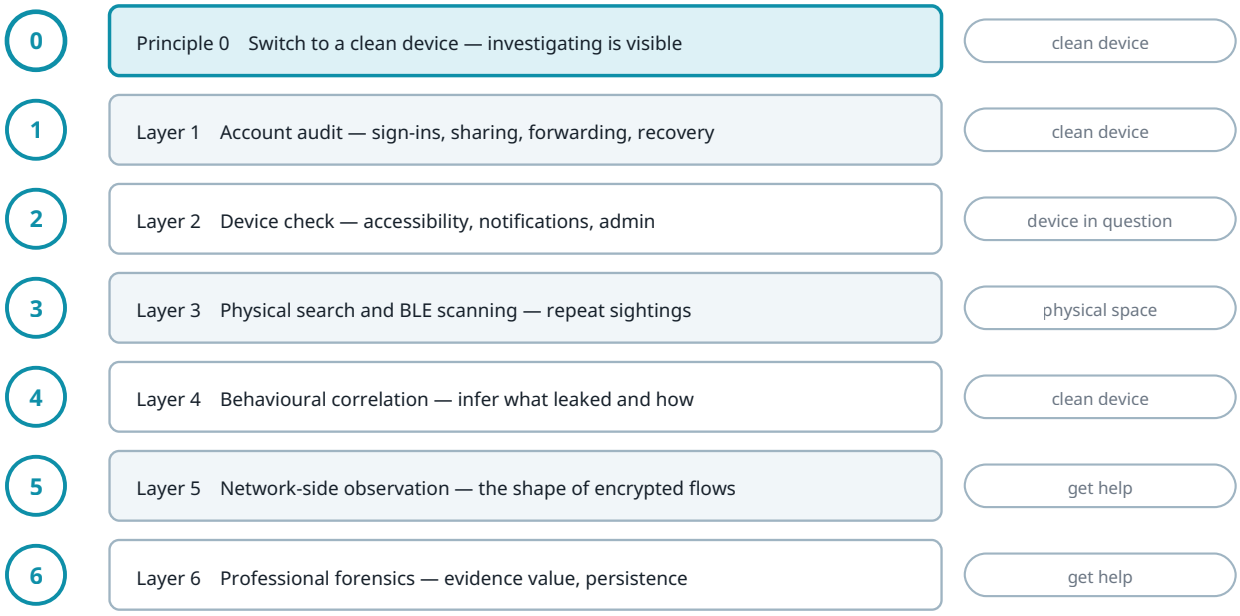

Figure 2: Six layers of detection. Principle 0 is not preparation — it is what makes every other layer trustworthy.

Principle 0: obtain a clean device first

This is the one step that cannot be skipped.

If surveillance is present, then every search, every message asking for help and every call you make on the device under examination may appear before the abuser in real time — **including the fact that you are reading this white paper**.

Therefore:

- Use another person's device, a public computer, or a device with no connection to any of your existing accounts;
- On that clean device, **do not sign in** to any of your existing accounts (signing in may itself trigger a sign-in notice and expose you);
- If you must keep using the original phone normally to avoid arousing suspicion, keep its pattern of use consistent with the past.

Decision criteria: what counts as a finding, and what does not

Without an explicit decision criterion, a check is only "having a look" — after which you can neither confirm nor rule anything out.

Layer	Positive criterion (on reaching it, go to chapter 7)	Vectors covered	Where it is carried out
One: account audit	**Any entry whose origin you cannot account for** — a device, a sharing recipient, a forwarding rule, an authorisation	B, plus the service-side part of D	Clean device
Two: device inspection	An entry you do not recognise in a high-privilege list	C, plus the device-side part of D	The phone in question (highest risk)
Three: physical search and Bluetooth scanning	Finding an electronic device that is not yours; or **the same target reappearing in three or more distinct locations, separated in time**	A	The physical environment
Four: behavioural cross-reference	The possible source of what the other party knows has narrowed to a specific vector	All four, but it can only point, not confirm	Clean device
Five: network-side observation	A persistent connection observed at the network boundary whose upstream is markedly greater than its downstream (periodicity can corroborate, but is unreliable on its own — see chapter 6)	B and C, part of D	Requires professional help
Six: professional forensics	Determined by the examiner	All	Professional help

Layers one and two share a single criterion: **"you cannot account for its origin" is a positive**, not "it looks suspicious". The distinction matters — monitoring software disguises itself as a system component, and sharing relationships display plausible names; if the criterion were "looks suspicious", every successfully disguised entry would be excluded by the reader themselves. **Whether you can say when, by whom, and why an entry was created is the reliable criterion.**

Layer three's criterion deserves emphasis: **seeing an unknown Bluetooth device in a single scan means nothing** — an urban environment has dozens of devices advertising at any moment. The meaningful signal is **the same target recurring across locations**, because only that rules out "someone else's device that happened to be nearby". This is also precisely why identifier rotation weakens generic scanners (see section 3.1).

Layer four is different in kind: it can only **narrow the range**, never confirm. Its value is that when every technical check comes back negative, it is often the only way left to make progress.

The three sources of a negative result

This is the most important passage in this methodology. When you finish a check and find nothing, there are three possible reasons, and **they mean entirely different things for what you do next**:

1. **A true negative**: there really is no surveillance.
2. **Looking in the wrong place**: surveillance exists, but not in the layer you checked. In this subject that is **the principal case, not the exception** — three of the four vectors are permanently negative under a device scan (see section 3.5). A clean phone scan does not rule out account-layer or settings-layer surveillance at all.
3. **The check itself failed**: the check never actually ran. Bluetooth was off, so scanning was not really operating; location permission was never granted, so the alerting mechanism never activated; an audit that belongs on a clean device was performed on the phone under investigation; or the application list was reviewed but not the high-privilege authorisation lists.

The third is the most dangerous, because **it looks exactly like a true negative**: both are "nothing found". So before trusting any negative result, confirm that the check itself is in working order — that is what the preparation list in Appendix A is for. **It is not preparatory work; it is the precondition for a negative result being believable at all.**

> This principle comes from our measurement experience: in WP-003 we three times hit a situation where the instrument had already broken while its output was indistinguishable from a genuine negative — had it gone unnoticed, "no leak" would have become the conclusion. The measurements in chapter 6 therefore make a positive control a precondition of every run. **Verify that the tool is in a measurable state before reading the result** — the same principle holds for self-detection.

Layer one: account audit (on the clean device; lowest risk)

The account layer is both the most common and the most frequently skipped. Review each of the following:

- **The list of signed-in devices**: is there a device or location you do not recognise? Watch for entries whose names look plausible but are not in fact yours.
- **Location sharing and family groups**: who can currently see your location? When was the sharing turned on?
- **Automatic forwarding and filter rules in email**: is there a forwarding rule you did not create? This is one of the most common and most persistent bypasses.
- **Account recovery options**: do the recovery email address and phone number still belong to the other person? **If recovery options are not cleaned up, a password change can be undone within minutes.**
- **Two-factor settings**: is the second factor bound to a device or number the other person can reach?

- **App-specific passwords and authorised third-party applications**: is there anything you do not remember authorising?
- **Backup settings**: does the backup destination account genuinely belong to you?

Layer two: device inspection (on the phone in question; highest risk)

Read section 7.1 before carrying out this layer. Everything you do in this layer may be recorded by the monitoring software.

To keep running, device-layer surveillance must hold a small number of high privileges. The point of the check is therefore not to spot a suspicious icon but to **go through these lists one by one and confirm that you recognise every entry**:

- **The accessibility services list**: this is the key privilege for reading screen contents and simulating input, and very few legitimate applications need it.
- **The notification access list**: this privilege allows reading every application's notifications, message previews included.
- **Device administrator or device owner privileges**: software holding this cannot be removed in the ordinary way.
- **Configuration profiles and work profiles**: a company you never worked for, or an organisation name you do not recognise, is an anomaly.
- **Permission to install from unofficial sources**: which applications are allowed to install other applications?
- **Battery and mobile-data usage rankings**: software continuously reporting in the background leaves traces in these two lists even when its icon is hidden.
- **Paired Bluetooth devices and known wireless networks**: is there anything that is not yours?

If you find an entry you do not recognise in any of these lists, **record its full name and stop**, then go to chapter 7.

Layer three: physical search and Bluetooth scanning (for vector A)

Active scanning (about ten minutes): use the manual-scan feature of each ecosystem's official app, or a generic Bluetooth Low Energy scan with an open-source tool. Scan **away from your own devices**, and repeat in at least three different places while moving — a device genuinely following you reappears across locations.

Physical search (one to several hours; no tools): when scanning finds nothing but suspicion remains, a physical search is often the most effective step. **Search in a fixed order, not by intuition.**

- **Vehicle** (the most common placement) — outside: inside bumpers front and rear, wheel arches, underbody rails, behind the licence-plate frame, near the tow hook; inside: under and between seats, spare-wheel well, boot side pockets, deep in the glovebox, roof-lining edges. Check every magnetic-friendly metal surface — magnetic cases are a common accessory.

- **Carried items**: bag compartments and under base padding, coat linings and seams that feel odd, wallet slots, suitcase linings and wheel housings, and any regularly carried item with unusual weight or thickness.
- **Home and gifts**: plush toys, ornaments and electronic accessories given by the person concerned; children's schoolbags and toys (especially amid custody disputes). This territory may also involve audio and video recording devices, which are handled differently — professional help is advisable.

Safety note while searching: if the abuser may observe you searching (cohabitation, for instance), the search itself can escalate risk. Discuss timing with professional support first.

Layer four: behavioural correlation (on the clean device; the most underrated)

When the technical checks come back negative, this layer is often the only way forward. List, one by one, the things the other person knows, and work backwards to **the possible source of each item**:

- Do they know your **location**, or the **content of conversations**? The former may come from vector A or B; the latter points to vector C or to notification access.
- Do they know about **messages you sent**, or **drafts you never sent**? The range of possible sources for the latter is very narrow.
- Is what they know **immediate**, or **delayed by some hours**? The latter fits periodic sync or backup.
- Do they know something that appeared on **only one particular device**? That narrows the scope to that device directly.

Record in writing the date of each incident, exactly what the other person knew, and where you were at the time — this record is also the evidentiary basis for section 7.2.

Layer five: network-side observation (requires professional help)

The first four layers all take place inside the device under examination or its accounts; when those checks all come back negative, and the principle in section 2.4 reminds us that the device itself cannot be trusted, **the only vantage point not yet contaminated is the network boundary**.

Monitoring software has to send data out, and that traffic passes through some router. The content is encrypted, but **the shape of the traffic is not**: when it is sent, how much, to where, and the ratio of upstream to downstream are all visible.

These clues do not carry equal weight. Our own measurements show that **what carries the discriminating power is "upstream markedly greater than downstream", not "the reporting is very regular"** — the latter can be evaded by jitter at zero cost, and it misclassifies benign heartbeat traffic along with everything else. Chapter 6 is that measurement, together with the boundaries within which it applies.

This layer needs equipment and expert interpretation, and is not within self-help range.

Layer six: professional forensics (when to stop doing this yourself)

Stop checking on your own and seek professional help if any of the following applies:

- you have found a device but cannot determine its type or whether it records audio or video;
- you have found unidentified software with high privileges on the device, and the situation involves immediate risk to your safety;
- the surveillance appears to come from a technically capable abuser, or involves a work device and enterprise management privileges;
- the case may enter legal proceedings and needs forensically sound reporting;
- you have cleaned up repeatedly but the surveillance keeps returning — this usually means a persistence channel remains undiscovered.

Acting yourself and preserving forensic evidence are difficult to reconcile: any clean-up destroys the scene. **If the legal route matters to you, preservation takes priority over removal.**

---

6. Network-side observation: periodic reporting in encrypted traffic

6.1 Why an observation point outside the device is needed

The first four layers of chapter 5 share one weakness: they all take place inside the device under examination or its accounts, and that is exactly what may already be under control. The fifth layer moves the observation point off the device.

Its basis is a structural fact: **for surveillance to be worth anything, the data has to leave the device.** Encryption protects content, but it cannot hide that something is being sent out at regular intervals. The question therefore becomes: from a position that sees only traffic metadata, how easily can periodic reporting be recognised?

The question has practical consequences, because it determines whether a service can honestly be offered: if the criteria only work under ideal conditions, then claiming to detect surveillance from the network side is selling a feeling of safety. So we measured its limits.

6.2 The experimental environment

The measurements were run in **BCN-349**, a four-node controlled environment built in EVE-NG with all addresses drawn from the RFC 5737 documentation ranges. The full topology, the provisioning scripts, the raw data and the analysis programs are published with this white paper (see Appendix D).

Node	Role
PHONE	The monitored end. Generates the periodic reporting flow and ordinary background application traffic at the same time
GW	**The observation point** (in the role of a home router). Captures traffic throughout
C2	The monitoring server. Receives uploads over TLS and replies with a short ack
WEB	A benign server, used by the background traffic

The subject under test is the behavioural pattern of periodic reporting, not any real monitoring product. The experiment involves no monitoring software sample and evaluates no product.

6.3 Criteria and method

All criteria are computed solely from traffic metadata visible at GW; **nothing is decrypted**:

Criterion	Definition	Intuition
T: timing criterion	Coefficient of variation of the intervals between consecutive upstream events, `CV < 0.15`	Human activity is irregular; a program's reporting is regular
A: direction criterion	`upstream bytes ÷ downstream bytes > 3`	Ordinary use is mostly download; exfiltration is mostly upload
T∧A: combined criterion	Both hold at once	—

Variables: three reporting periods (2, 5 and 15 seconds) × three jitter ratios (0%, 20% and 50%), nine configurations in all. Jitter is defined as the actual interval falling uniformly at random within plus or minus that proportion of the period. The nine run in parallel within the same observation window, each on a different destination port so the flows can be separated; benign background traffic runs at the same time and serves as the denominator for false positives. **Thirty windows are observed per configuration**; below thirty, no statistic is reported for that cell, only the raw observations.

> Both thresholds (`CV < 0.15` and a ratio > 3) are fixed values chosen in advance. The choice of threshold governs the conclusion, and its cost is set out in section 6.5 — **the direction criterion's threshold was never put to any test in this experiment**.

Instrument verification: before each run we confirm that the capture really sees every flow under test and that the monitored end has no bypass around the observation point; and the zero-jitter configuration serves as a positive control — if it is not reliably caught by the timing criterion, the capture or the analysis has failed, and that run is void rather than recorded as a negative. This maps directly onto the third of the *three sources of a negative result* in chapter 5.

6.4 Results

All thirty observation windows **passed instrument verification**: before and after each run we confirmed that all nine flows under test were present (30/30), that the zero-jitter configuration (the positive control) was indeed caught by the timing criterion (30/30), and that all 315 connections arriving at the monitoring server came from the observation point's external address, none of them bypassing it (315/315). Void windows: 0. Every statistic below rests on all thirty windows.

Detection (30 windows per cell)

Period	Jitter	Events per window	Median CV (IQR)	T	A	T∧A
2 s	0%	150	0.0010 (0.0002–0.0011)	30/30	30/30	30/30
2 s	20%	151	0.1158 (0.1132–0.1181)	30/30	30/30	30/30
2 s	50%	150	0.2904 (0.2797–0.2977)	**0/30**	30/30	0/30
5 s	0%	60	0.0002 (0.0002–0.0005)	30/30	30/30	30/30
5 s	20%	60	0.1167 (0.1116–0.1205)	30/30	30/30	30/30
5 s	50%	61	0.2892 (0.2772–0.3030)	**0/30**	30/30	0/30
15 s	0%	20	0.0002 (0.0001–0.0002)	30/30	30/30	30/30
15 s	20%	20	0.1127 (0.1064–0.1179)	30/30	30/30	30/30
15 s	50%	21	0.2926 (0.2641–0.3131)	**1/30**	30/30	1/30

> The single hit in the 15 s / 50% cell fell in the 25th window, which observed only 17 events and gave a CV of 0.1455 — just under the threshold. **That is sampling variation in the CV estimate, not evidence that longer periods are easier to detect** — the same mechanism can equally produce a spurious miss. Dispersion shows it most clearly: the 15 s / 50% configuration, with the fewest events, spans a CV range of 0.1455–0.3652, while the 2 s / 50% configuration, with the most events, spans only 0.2692–0.3200. **Sample size affects the stability of the estimate, not the central tendency of the criterion** — the median CV at all three periods is almost identical (0.290 / 0.289 / 0.293), because CV is dimensionless. With n = 1 we report the raw observation only and claim no rate.

False positives (577 benign flows across the thirty windows; median 19 per window, range 15–23)

Benign flow type	Flows	T false positives	A false positives	T∧A false positives
Page-load flows	487	16 (3.3%)	0	0
Sync flows	30	2 (6.7%)	0	0
keepalive flows	60	**60 (100%)**	0	0
Total	577	78 (13.5%)	**0**	**0**

Benign flows misclassified per window: median 2 for T (range 2–4), and 0 for both A and T∧A.

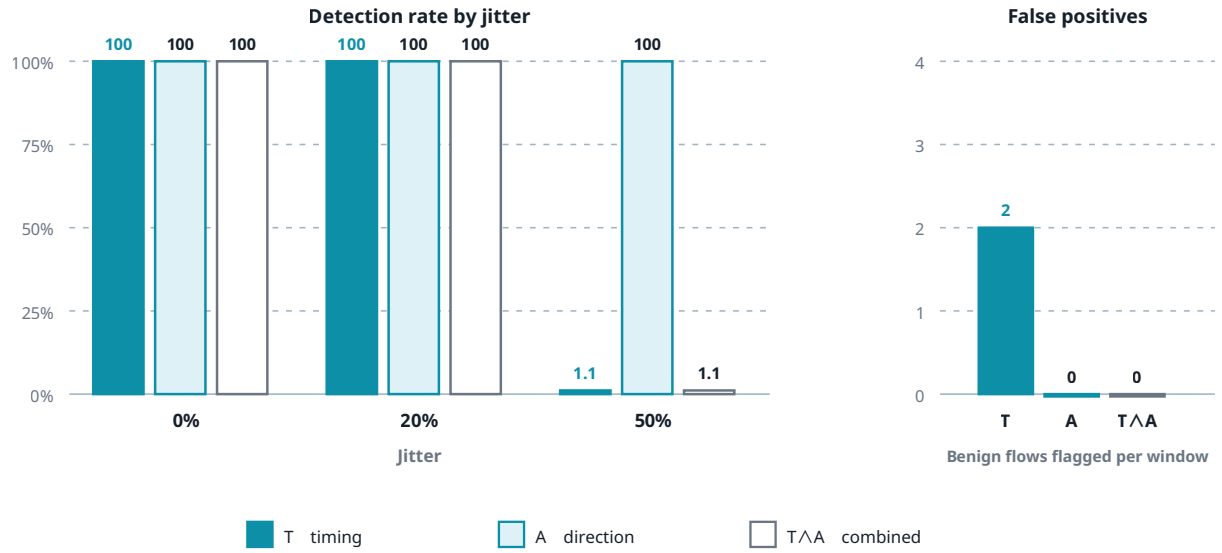

Figure 3: BCN-349 measurement. Left: detection rate, each bar pooling three reporting periods over 90 observation windows (the full 3×3 result is in the table in 6.4) — past the breakpoint, jitter defeats the timing criterion while the direction criterion is unaffected. Right: median number of benign flows flagged per window (about 19 benign flows per window) — the timing criterion flags benign keepalives consistently, which is why it cannot be used on its own.

Three findings

1. **The direction criterion is entirely reliable in this environment; the timing criterion is not.** A hit in all 270 observations (nine configurations × thirty windows), with zero false positives across 577 benign flows; T missed almost everything at 50% jitter.
2. **The timing criterion's failure point can be calculated in advance, and can therefore be evaded at zero cost.** Under uniform jitter the coefficient of variation of the intervals is `CV = j / √3`, independent of the length of the period. A threshold of 0.15 therefore corresponds to a critical jitter of `j* ≈ 26%`. Measurement agrees with theory to within three parts in a thousand: at 20% jitter the measured median was 0.1158 against a theoretical 0.1155; at 50% jitter, 0.2904 against 0.2887. **Evading the timing criterion requires no skill at all — set the jitter above a quarter, and jitter costs nothing to implement.**
3. **The timing criterion misclassifies benign periodic traffic consistently.** Two keepalive flows were judged positive by T in all thirty windows: 60/60, 100%. This is not an occasional false positive but

an inevitability: a keepalive is by definition a small packet sent at regular intervals. What distinguishes it from exfiltration is not timing but direction.

What this converges on

In this environment, **the combined criterion T∧A is no better than the direction criterion A used alone**: A already has zero false positives, and adding T only makes it miss, at high jitter, targets it would otherwise have caught. The timing criterion is worth something only where the adversary does not know that anyone is watching the timing — and that is precisely the assumption Gap two in chapter 4 identifies as one that may not be taken as a premise: **a protection mechanism's effectiveness cannot rest on the attacker's cooperation.**

6.5 What this data does and does not support

Supports:

- From a position that sees only traffic metadata, connections with periodic and upstream-asymmetric characteristics can be separated from the background traffic established in this experiment.
- The two criteria differ markedly in their sensitivity to jitter, and the timing criterion's failure point matches the theoretical prediction (see finding 2), so that failure point can be extrapolated to jitter ratios this experiment did not test.
- **The timing criterion does not stand on its own**: it can be evaded at zero cost, and it consistently misclassifies common benign periodic traffic.

Does not support:

- **It does not support the claim that the direction criterion is also free of false positives on a real network.** This is the most important limitation in this chapter. The highest upstream-to-downstream ratio among all benign flows was 1.44, and the lowest among all reporting flows under test was 9.85 — between them lies an empty band several times wide. **That means the threshold of 3 was never put to any test**: this experiment cannot distinguish "the threshold was well chosen" from "the background traffic was too easy". The reason is that **the background traffic contains no upstream-dominant benign traffic** — cloud backup, photo-library sync, outbound video and file uploads are exactly the class that would challenge A. A's zero false positives are a product of this background model, not a property of A. To claim that A is usable on a real network, one would have to re-measure in an environment that includes backup and sync traffic. **This white paper makes no such claim.**
- **It does not support** the claim that network-side detection can replace device inspection. The network side cannot see the content of the surveillance, nor identify which application is responsible; it can only indicate that one connection has the characteristics of exfiltration.
- **It does not support** any inference about the reporting behaviour of real monitoring products. The subject under test was controlled reporting behaviour; the periods, jitter and transport of real products were not measured.

- **It does not support** extending these results to interpretation outside encrypted traffic, or to network conditions this experiment did not cover.

A note on the measurement proxy: this experiment uses periods of 2, 5 and 15 seconds as a proxy, whereas real monitoring software typically reports at minute to hour scale. The proxy remains meaningful because the behaviour of the timing criterion depends on **the ratio of jitter to period** and on **the number of samples observed**, not on the absolute length of the period — the theoretical relationship in finding 2 is independent of the period, and the measurements at all three periods were indeed consistent. But it **cannot** be used to infer how long one would have to observe in a real situation before enough events accumulate — that depends on the absolute period and must be measured separately.

---

7. What to do after finding surveillance

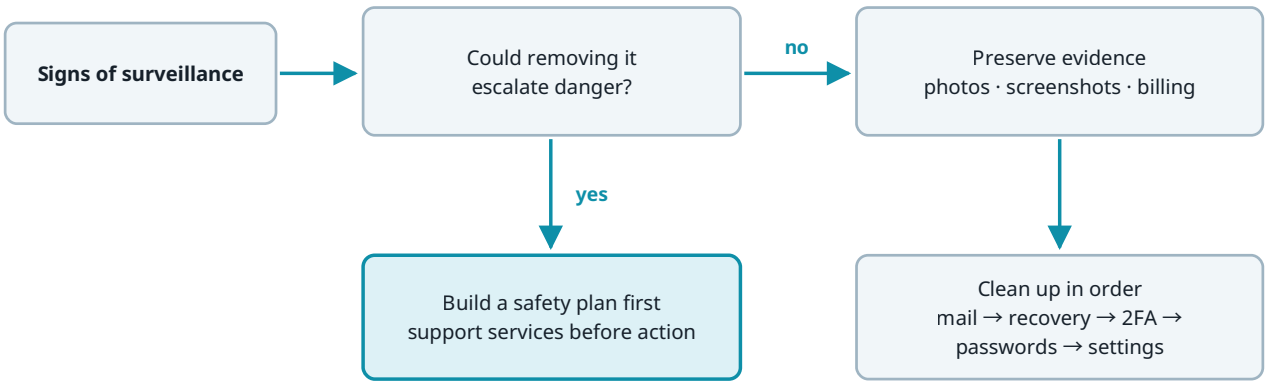

Figure 4: What to do after finding surveillance. Removal is never step one.

Most technical guides end at detection. For the person under surveillance, **the moment of discovery is when the real problem begins**.

7.1 First priority: personal safety

Removal is an act the other person will notice immediately. Once a tracker is removed or shielded, its location reports stop or freeze in place; once software is deleted or an account signed out, reporting ends at once. To an abuser these are unambiguous signals that the victim has become aware and has begun to act.

In research on intimate-partner abuse, **the stage at which a victim resists or prepares to leave is the stage of greatest danger.**

So, before acting, establish:

- do you have a workable safety plan (somewhere safe to go, someone to contact, the documents and money you would need)?

- if the other person arrives or makes contact immediately after removal, how will you respond?
- has a professional social worker or the police already become involved?

If the answer is no, the reasonable choice is to leave things as they are for now, build support first, and deal with the technical layer afterwards. Surveillance is hard to endure, but a technical clean-up cannot protect your safety; the order cannot be reversed.

If you are in immediate danger, call 110 (in Taiwan); for protective services and counselling, call 113.

7.2 Evidence preservation

Complete the following record before removing anything. Store all of it **somewhere other than the device in question**:

For a physical tracker:

1. **Photograph it in place first**, wide shots and close-ups, so the photographs show where it was hidden;
2. Record the date and time of discovery and the exact position (which part of the car, which bag compartment);
3. Record its appearance, serial number and any visible identifiers; screenshot anything your phone can read from it;
4. **Do not reset it, do not disassemble it, do not remove the battery** — a reset severs the link to the owner's account, which is precisely the connection that points to the abuser;
5. If its function must be interrupted without damaging it, place it in an opaque container or wrapping.

For surveillance on a phone or an account:

1. Screenshots of the suspicious software or settings, with full names and authorisation status;
2. Screenshots of account sign-in records, with device names, times and locations;
3. **Subscription and billing records** — commercial monitoring services must be paid for, and the money trail is among the strongest evidence available;
4. The behavioural correlation record built in layer four (dates, what the other person knew, where you actually were);
5. If possible, have someone **film you operating the device in question** with a second device, to corroborate that the screenshots were not fabricated afterwards.

In both cases: do not delete, do not reset, and do not replace the phone until evidence preservation is complete.

> **The most common mistake**: discarding the device or deleting the software immediately, in shock or anger. That destroys the only physical evidence linking the conduct to a specific person, and later proof becomes very difficult. **A reset is the most thorough removal, and the most thorough destruction of evidence.**

7.3 Legal routes in Taiwan

- **Stalking and Harassment Prevention Act**: the conduct it covers includes persistent monitoring, following or ascertaining the whereabouts of a specific person by electronic communications, the internet or other equipment. A victim may report to the police; following investigation a written warning may be issued to the perpetrator, and a further offence within two years of that warning may support an application for a protection order.
- **Criminal Code, offences against computer security**: unauthorised intrusion into another's computer or related equipment, unauthorised acquisition or alteration of another's electromagnetic records, and producing programs specifically for committing these offences, all carry penalties.
- **Criminal Code, offences against privacy**: unauthorised recording — by audio, photography, video or electromagnetic record — of another's non-public activities, conversations, speech or intimate body parts carries penalties.
- **Personal Data Protection Act**: unlawful collection, processing or use of personal data causing harm to others carries corresponding liability.
- **Domestic Violence Prevention Act**: where the parties are family members or intimate partners, a protection order may be sought, which may include prohibiting harassment and stalking of the victim.

When reporting, we suggest going to a **Women and Children Protection Unit**, or stating clearly that the case involves stalking and domestic violence, so that it is handled by officers familiar with such cases. Bring the evidence preserved under section 7.2, and keep the original device. **Legal assessment of a specific case belongs to a lawyer**; this white paper is not legal advice.

Situation	Channel
Immediate danger	110
Domestic violence, sexual assault, child protection	113 protection hotline
Filing a stalking case	Local police Women and Children Protection Units
Legal consultation	Local bar associations; the Legal Aid Foundation

7.4 The order for rebuilding a clean environment

Getting the order wrong is the most common reason a clean-up fails. **If you change a password while the account's recovery options are still controlled by the other person, the new password may be reset within minutes.**

1. **On the clean device, deal with the email account first**: it is the recovery exit for every other account and must be recovered first.
2. **Clean up recovery options**: remove any recovery email address or phone number that is not yours.

3. **Reset two-factor authentication**: bind it to a device only you hold, and regenerate the backup codes.
4. **Change passwords and revoke all sessions**: revoking sessions is the critical step — with a password change alone, existing sign-ins may remain valid on some services.
5. **Work through the other accounts**: cloud, social, financial, carrier.
6. **Clean up the settings layer**: forwarding rules, call diversion, location sharing, family groups, paired devices, configuration profiles.
7. **Only then deal with the device in question**: after evidence preservation is complete, restore it to factory settings — and **do not restore from an old backup**, which may carry the surveillance settings with it. Set it up again by hand.
8. **Review every other device**: tablets, computers, in-car systems, smart-home devices.

Afterwards, repeat the layer-one account audit within a few weeks. **A return of the surveillance usually means one channel on the list was never found.**

---

8. Recommendations for vendors and policymakers

For tracker vendors and mobile platforms:

1. **Detection should not stop at ecosystem borders**: abusers face no ecosystem constraint when choosing a device, so cross-platform coverage should be completed quickly.
2. **Confront the shared-space miss**: judgment logic centred on abnormal co-travel is least discriminating in the principal abuse scenario. What is needed may not be sharper parameters but a different judgment approach.
3. **Provide a path for people without smartphones**: for example low-cost standalone detectors, or testing tools that third parties — social workers, police units — can operate on a victim's behalf.
4. **Bring the account layer into a single safety-check entry point**: a complete audit currently requires crossing several settings screens. One entry point should cover signed-in devices, location sharing, forwarding rules, recovery options and third-party authorisations, with one-tap revocation.
5. **Sharing relationships should expire and require periodic reconfirmation**: location sharing and family groups should not remain valid indefinitely with no review prompt.
6. **Persistent disclosure for dual-use software**: software capable of continuously reporting location or content should, whatever its store category, maintain a persistent indicator on the monitored device that cannot be switched off. **Legitimate supervision is not harmed by the supervised party knowing.**
7. **Redesign the "what happens after discovery" flow**: alerts and safety prompts should not uniformly advise immediate removal. There should be a branch for "this may be intimate-partner surveillance" leading to evidence preservation and support resources.

8. **The recipient of notices should be independently configurable**: security notices should not go only to the bill payer or the family-group organiser.
9. **Verifiable device handover state**: users should have a way to establish which configuration profiles and high-privilege applications have been installed on a device since it left the factory or was last reset.

For policymakers and practitioners:

1. **Bring technical detection into protective services**: victims currently must overcome the technical barrier alone; basic detection capability and referral mechanisms should exist at the level of Women and Children Protection Units and domestic-violence prevention centres.
2. **Publish accessible evidence-preservation guidance**: the content of section 7.2 should be a standard explanation every receiving officer can give on the spot.
3. **Front-line staff need basic literacy in digital surveillance**, and the channels that help victims gather evidence should extend to digital evidence.

---

9. Limitations and further work

9.1 The external-validity boundary

The conclusions of chapters 2 to 5 are of the form "these gaps exist at the level of mechanism principle", not "this is how a given platform or product currently performs".

- This white paper tested **no** tracker product, **no** phone's detection feature and **no** platform protection. It **must not** be cited as an evaluation or comparison of any product or platform.
- The third-party measurements cited here apply only to the subjects, versions and dates those studies state. Vendor behaviour changes, and is known to change.
- The vector taxonomy in chapter 3 and the gap analysis in chapter 4 rest on the mechanisms described in public documentation. Where a vendor's actual implementation diverges from its documentation, that reasoning does not apply to that vendor.
- **The methodology in chapter 5 is unvalidated**: we have not measured the detection rate or false-negative rate of any layer, nor verified that the clean-up order in section 7.4 actually stops surveillance. These are reasonable procedures designed from mechanism principles, not procedures known to work.
- The limits of chapter 6 are stated in section 6.5; its conclusions must not be extended to conditions that experiment did not cover.
- This white paper does not name specific products, nor record how monitoring software is obtained or installed; this is a deliberate choice (see section 3.3). Its cost is that readers cannot independently verify that portion of the text — a cost we judge acceptable against the reduction in usefulness to abusers.
- The legal discussion covers current Taiwanese law and does not constitute legal advice.

9.2 Pre-registered experimental design (v2.0)

The following design is published before any measurement begins, so that others can scrutinise the design itself before seeing any results. **Publishing the design before generating the data means the criteria cannot afterwards be bent to fit the outcome.**

Common settings

- **Subjects**: one tracker from each of the three major ecosystems, paired with one test device from each of the two major platforms; the surveillance vectors are established separately on our own test devices. Every device is restored to factory settings before each run.
- **Control design**: each scenario establishes exactly one vector — with two established at once, a detection could not be attributed to either. An unmonitored device of the same model serves as control.
- **Repetitions**: 20 runs per combination per scenario for the tracker cases; 10 runs per combination for vectors B to D. Below those counts, no statistic is reported for that cell; only the raw observations.
- **Instrument verification**: before each run, confirm with known-present surveillance that the layer of checking can in fact see it (a positive control). **A failed positive control aborts that run**, excluded from the statistics.
- **Statistical presentation**: time-based figures are reported as medians with interquartile ranges rather than means (the distributions are right-skewed; means get dragged by outliers).
- **Disclosure**: all raw records and analysis scripts ship with v2.0, in the format of Appendix D of this white paper. **The monitoring samples themselves will not be distributed.**

Experiment matrix

ID	Scenario	Variable under test	Primary measurements
E1	Cross-ecosystem detection matrix	Tracker ecosystem × phone platform	Whether an alert ever appeared; latency of the first alert
E2	Alert-latency distribution	Time of day (daytime / night)	Distribution of separation-to-alert time
E3	Shared space	Whether the owner remains co-located with the victim	Whether alerts trigger; difference against E2 under otherwise identical conditions
E4	Per-layer detection rate	Surveillance vector × layer of checking	Whether that layer finds that vector; false-negative rate
E5	Clean-up completeness	Clean-up order (per section 7.4 / arbitrary)	Whether surveillance actually stops; whether it recurs weeks later
E6	Whether changing phones works	Whether the account is handled at the same time	Whether surveillance resumes once the new device signs in

Each scenario changes exactly one thing.

Pre-declared falsification conditions

If any of the following occurs, the corresponding claim must be treated as refuted, and v2.0 will record it explicitly:

- In E1, all ecosystem × platform combinations trigger alerts within comparable times → **Gap two** does not hold.
- E3's alert behaviour does not differ appreciably from E2's → the shared-space reasoning within **Gap three** does not hold.
- In E4, device-layer checking reliably finds vector B or vector D surveillance → the comparison table in section 3.5 needs revision.
- In E5, arbitrary ordering is no less effective than the order in section 7.4 → that section's claim that "order matters" does not hold.
- In E6, changing the device alone lifts the surveillance → the closing claim of section 3.5 is wrong.

Ethical and legal premise: per section 4.1 of our Research Ethics Code, all test subjects are the researchers themselves or volunteers with written informed consent. No one is tested without consent, no real victim's device or data is used, and trackers remain on the researchers' own belongings throughout.

9.3 Feedback we hope to receive

- Practitioners' views (social workers, police officers, lawyers) on the feasibility of chapter 7;
- Additions or counter-evidence to the gap analysis in chapter 4;
- Persistence channels the layered checks do not cover;
- Whether Principle 0 is too demanding for a victim with no access to a second device;
- Local case patterns in Taiwan (**please do not include identifying information**).

Write to `contact@odysec.org`; PGP-encrypt sensitive material to `security@odysec.org`.

Appendix A: Detection checklist

> Printable as a standalone page.

Before you start - ☐ Judge: could acting on this escalate the abuser's behaviour? If yes, contact 113 or the police first - ☐ A clean device unconnected to my accounts has been obtained - ☐ No existing account has been signed in on that clean device - ☐ No search for help has been made on the phone in question

Preparation: confirm the check itself works

> This section is not preparatory work; it is **the precondition for a negative result being believable**. If any item fails, then "nothing found" only means the checking mechanism was never active (see chapter 5, *The three sources of a negative result*).

- ☐ Bluetooth on
- ☐ Location permission granted
- ☐ OS up to date
- ☐ Tracking alerts not disabled
- ☐ The other ecosystem's detection app installed

Layer one: account audit (clean device) - ☐ Signed-in device list reviewed entry by entry - ☐ Location sharing and family groups reviewed entry by entry - ☐ Email forwarding and filter rules reviewed entry by entry - ☐ Recovery email address and phone number confirmed as mine - ☐ Two-factor authentication bound to a device I hold - ☐ Third-party application authorisations reviewed entry by entry - ☐ Backup destination account confirmed as mine

Layer two: device inspection (the phone in question; highest risk) - ☐ Accessibility services list - ☐ Notification access list - ☐ Device administrator or device owner privileges - ☐ Configuration profiles and work profiles - ☐ Sources permitted to install other applications - ☐ Battery and mobile-data usage rankings - ☐ Paired Bluetooth devices and known wireless networks

Layer three: physical search and Bluetooth scanning - ☐ Official app's manual scan run away from my own Bluetooth devices - ☐ Scan repeated in at least three **different places**, recording unknown devices that recur - ☐ Car exterior: bumper interiors, wheel arches, underbody, behind plate frame - ☐ Car interior: under/between seats, spare-wheel well, boot side pockets, glovebox, roof lining - ☐ Bags: compartments, under base padding - ☐ Clothing: coat linings, odd seams - ☐ Wallet, suitcase lining and wheels - ☐ Gifts from the person concerned - ☐ Children's schoolbags and toys

Layer four: behavioural correlation (clean device) - ☐ What the other person knows has been listed - ☐ Location-type and content-type information distinguished - ☐ Immediate versus delayed knowledge assessed - ☐ Dates, contents and my own movements recorded

If any sign of surveillance is found - ☐ **Do not remove it immediately; do not reset; do not discard** - ☐ Whether removal would escalate danger has been assessed - ☐ Tracker: photographed in place (wide + close), time, place and identifiers recorded - ☐ Phone and accounts: screenshots and sign-in records preserved (stored off the device in question) - ☐ Subscription and billing records preserved - ☐ My own movement timeline recorded - ☐ Evidence taken to a Women and Children Protection Unit - ☐ Clean-up follows the order: mail → recovery options → 2FA → passwords and sessions → other accounts → settings layer → device - ☐ Not restored from an old backup - ☐ Layer-one audit repeated a few weeks later

Appendix B: Terminology

中文	English	日本語
藍牙追蹤器	Bluetooth tracker	Bluetooth トラッカー
群眾外包定位網路	crowd-sourced location network	クラウドソース型位置ネットワーク
意外追蹤警示	unwanted tracking alert	不要な追跡の通知
輪替識別碼	rotating identifier	ローテーション識別子
跟蹤軟體	stalkerware	ストーカーウェア
雙用途程式	dual-use application	デュアルユース•アプリ
帳號層監控	account-level surveillance	アカウント層の監視
無障礙服務	accessibility service	ユーザー補助サービス
通知存取	notification access	通知へのアクセス
裝置管理員	device administrator	デバイス管理者
管理描述檔	configuration profile	構成プロファイル
工作設定檔	work profile	仕事用プロファイル
工作階段	session	セッション
救援管道	account recovery options	アカウント回復手段
雙因素驗證	two-factor authentication	二要素認証
乾淨裝置	clean device	クリーンな端末
證據保全	evidence preservation	証拠保全
持久化	persistence	永続化
判定準則	decision criterion	判定基準
偽陰性	false negative	偽陰性
變異係數	coefficient of variation	変動係数
流量中繼資料	traffic metadata	トラフィックのメタデータ
外部效度	external validity	外的妥当性
預先登記	pre-registration	事前登録

Appendix C: References

The basis for every technical claim in this white paper is listed here. **Where a third-party measurement is cited, please carry its subject and its date along with it.**

Academic research

1. Heinrich, A., Stute, M., Kornhuber, T., & Hollick, M. (2021). *Who Can Find My Devices? Security and Privacy of Apple's Crowd-Sourced Bluetooth Location Tracking System.* Proceedings on Privacy Enhancing Technologies, 2021(3), 227–245. https://petsymposium.org/popets/2021/popets-2021-0045.pdf — Mechanism analysis of crowd-sourced location networks; the basis for the technical background in section 3.1.
2. Heinrich, A., Bittner, N., & Hollick, M. (2022). *AirGuard — Protecting Android Users from Stalking Attacks by Apple Find My Devices.* ACM WiSec 2022 (Best Paper Award). https://arxiv.org/abs/2202.11813 — Design and in-the-wild observation of a cross-ecosystem detection tool; the basis for Gap two.
3. Shafqat, N., Gerzon, N., Van Nortwick, M., Sun, V., Mislove, A., & Ranganathan, A. (2023). *Track You: A Deep Dive into Safety Alerts for Apple AirTags.* Proceedings on Privacy Enhancing Technologies, 2023(4), 132–148. https://petsymposium.org/popets/2023/popets-2023-0102.pdf — Measured alert latency (including the observation of daytime delays beyond eight hours) and a demonstrated cloned tracker evading alerts; the basis for section 4.1 and Gap two. **That measurement targeted a specific product and was carried out in 2023; later OS versions may behave differently.**
4. Chatterjee, R., Doerfler, P., Orgad, H., Havron, S., Palmer, J., Freed, D., Levy, K., Dell, N., McCoy, D., & Ristenpart, T. (2018). *The Spyware Used in Intimate Partner Violence.* IEEE Symposium on Security and Privacy 2018, 441–458. https://www.ieee-security.org/TC/SP2018/ — The first systematic measurement of the intimate-partner surveillance software ecosystem; establishes dual-use applications as its bulk. Basis for the executive summary and Gap three.
5. Freed, D., Palmer, J., Minchala, D., Levy, K., Ristenpart, T., & Dell, N. (2018). *"A Stalker's Paradise": How Intimate Partner Abusers Exploit Technology.* ACM CHI 2018 (Best Paper Award). https://dl.acm.org/doi/10.1145/3173574.3174241 — Qualitative study with 89 participants; basis for chapter 2's assumption of zero technical capability alongside abundant access.
6. Roundy, K. A., Mendelberg, P. B., Dell, N., McCoy, D., Nissani, D., Ristenpart, T., & Tamersoy, A. (2020). *The Many Kinds of Creepware Used for Interpersonal Attacks.* IEEE Symposium on Security and Privacy 2020. https://acartamersoy.github.io/papers/roundy_sp20.pdf — Shows that software used for interpersonal attacks extends well beyond surveillance; supporting basis for Gap three.
7. Havron, S., Freed, D., Chatterjee, R., McCoy, D., Dell, N., & Ristenpart, T. (2019). *Clinical Computer Security for Victims of Intimate Partner Violence.* USENIX Security Symposium 2019. https://www.usenix.org/conference/usenixsecurity19/presentation/havron — A practical model for professionally assisted technical examination; basis for layers five and six and the policy recommendations in chapter 8.

8. Freed, D., Havron, S., Tseng, E., Gallardo, A., Chatterjee, R., Ristenpart, T., & Dell, N. (2019). *"Is My Phone Hacked?" Analyzing Clinical Computer Security Interventions with Survivors of Intimate Partner Violence.* Proceedings of the ACM on Human-Computer Interaction, CSCW 2019. https://dl.acm.org/doi/10.1145/3359304 — Analysis of actual intervention cases; basis for the "removal is immediately noticed" risk assessment in section 7.1.

Standards documents

1. IETF Detecting Unwanted Location Trackers (DULT) working group. *DULT Threat Model*, draft-ietf-dult-threat-model; *Detecting Unwanted Location Trackers Accessory Protocol*, draft-ietf-dult-accessory-protocol. https://datatracker.ietf.org/wg/dult/about/ — The "cross-platform detection standard" referred to in chapter 4 means these two drafts. **Draft status changes; cite the current version.**

Organisational material

1. Coalition Against Stalkerware. Definitions, practical guidance, and participating-organisation material. https://stopstalkerware.org/
2. IPV Tech Research (a joint research programme including Cornell Tech). Literature and resources on intimate partner violence and technology abuse. https://www.ipvtechresearch.org/

Vendor and platform documentation

1. The official support documentation and privacy descriptions of the tracker vendors and both major mobile platforms. Because these are revised frequently and vary by region, this white paper pins no specific version; **for concrete behaviour, rely on the vendor documentation current at the time you read it, and on your own device.**

Legislation

1. Taiwan's Stalking and Harassment Prevention Act, the Offences Against Computer Security and Offences Against Privacy chapters of the Criminal Code, the Personal Data Protection Act, and the Domestic Violence Prevention Act.

Our own reports

1. Odysec (2026). *Beyond the Tunnel: Measuring the Protection Boundary of a VPN.* Odysec Research Report WP-003. — Source of the measurement discipline and of the principle of verifying the instrument before reading the result.

Appendix D: Experimental environment and data

The **BCN-349** environment used in chapter 6, its provisioning scripts, raw records and analysis programs are published with this white paper and can be downloaded from the research publication page.

The environment uses only freely available Debian 13 and **deliberately uses no image that cannot lawfully be redistributed**, so readers can rebuild it in full. Addresses are drawn from the RFC 5737

documentation ranges, and the keys and certificates are fixed lab-only values — **they are public and must never be used for anything real.**

Revision history

Version	Date	Change
1.0	2026-08-30	Initial release.

Citation: Odysec (2026). *The Eyes Beside You: Detecting and Responding to Close-Range Surveillance.* Odysec Technical White Paper WP-001, version 1.0.

Disclaimer: This white paper serves technical research and education; it is not legal advice. Consult a practising lawyer for any specific case. If you are in immediate danger, call 110 (in Taiwan); for protective services and counselling, call 113.