UBINODES

THE THREAT MODEL.

A research scenario on data protection, surveillance, and operational security.

Data Protection · Surveillance · Operational Security

ubinodes.org/research-2026

ABOUT THIS DOCUMENT.

This is a work of research. The scenario described here is fictitious. The company, the agency, the people and the events are inventions, assembled to illustrate a method. Everything around that fiction — the capabilities of modern surveillance, the conduct of certain state institutions, the documented cases cited in the closing section — is real.

The purpose is practical. There are countless communication tools, operating systems and security applications available, and any open discussion of which to trust quickly dissolves into competing opinions, where the loudest voice tends to win. A concrete worst-case scenario fixes the debate to the ground. It lets us ask, of any tool or habit, a single disciplined question: would this protect the people in the scenario, in their situation, against the adversary they face? Used that way, a threat model becomes the lens through which every later decision about software and procedure can be judged objectively.

It is written for two kinds of reader: those who simply want to understand the mindset that serious data protection requires, and those who will go on to design the procedures that put it into practice. It assumes no technical background. It does assume the willingness to take an uncomfortable premise seriously: that doing nothing wrong is not, by itself, protection.

CONTENTS.

1	Why a Threat Model
2	The Mindset
3	Threat Assumptions
4	The Scenario
5	Protection by Role
6	When the State Is the Adversary
7	Outside Forces
8	Techniques
9	The Discipline
10	Conclusion
11	Sources

1. Why a Threat Model.

A threat model is not a list of fears. It is a structured account of who might want your information, what they are capable of, and what it would cost you if they obtained it. From that account flow every sensible decision about which tools to use, how to use them, and how much effort each situation deserves.

Most security advice fails because it skips this step. It recommends an application, or a setting, or a habit, without first asking what the application is supposed to defend against. The result is a collection of disconnected precautions that feel reassuring and protect very little. Worse, it invites endless argument: every person brings their own experience and their own favoured tools, and with no shared scenario to test them against, the discussion goes nowhere.

The scenario that follows deliberately describes a worst case. This is not pessimism for its own sake. Security is not backward-compatible: once your information is exposed, no later precaution can recall it. A model calibrated to the worst plausible adversary therefore protects you against every lesser one as well, and it builds the habits before they are needed rather than after they have failed. The aim is not to make everyone behave like a hunted dissident. It is to let each person locate their own position honestly, and protect themselves accordingly.

There is a deeper reason to be proactive, and it is architectural. Reactive security — the kind that responds only after something has gone wrong — depends entirely on someone else being positioned to notice the harmful act and shut down the channel it travels through. It outsources your safety to a third party who must be present, capable, willing and trustworthy at the decisive moment. Often they are none of these; sometimes the third party is the adversary. This dependence is not merely a habit that can be corrected. It is built into the shape of certain systems. A network organised around central servers — highly available hubs that relay between one another, with users gathered at the edges — concentrates the very points on which reactive security relies, and so concentrates the points at which it fails. Whoever controls the hub controls the channel: they can watch it, sever it, or turn it against the people it was meant to protect, and nothing done afterwards can recover what passed through it. Security that cannot be switched off by anyone but you is not a matter of better monitoring; it is a matter of a different design — one with no central point to seize, and no single party whose compromise compromises everyone.

2. The Mindset.

Before any tool is chosen, a particular way of thinking has to be in place. The technology is downstream of the attitude.

- **The attack surface is yours to shape.** An adversary who does not know who to watch, where to look, or which channel matters must spread their effort thinly. Every habit that keeps them uncertain — separating roles, varying patterns, revealing nothing unnecessary — widens their problem and raises their cost. Uncertainty, deliberately maintained, is a defence in itself.
- **“Nothing to hide” is not a defence.** The belief that the innocent have nothing to fear is comforting and false. The danger is rarely what you have actually done; it is what can be made of what you have said, once a fragment of communication is lifted out of context and recast. A single sentence, stripped of its surroundings, can be made to suggest almost anything. Poor security does not merely risk exposing wrongdoing — it lets anyone build a story about you from your own data and make it stick. The defence is not to have clean conduct, which you may already have. It is to give the adversary nothing to work with.
- **Protect everyone you touch.** A threat model that protects only the principal is incomplete. Pressure is applied wherever it is easiest: a family member, a supplier, a contractor, a source. The decision to protect information is therefore also a decision to protect the people connected to it — those who never chose to be part of the work and have no defences of their own.
- **Be proactive, never retroactive.** Caution taken before exposure protects you; caution taken afterwards is theatre. The habits in this document are worth building while they still cost nothing, because the moment they are needed is the moment it is too late to begin.

3. Threat Assumptions.

A scenario is only as useful as the assumptions behind it. The following are the operating assumptions of this model. They are stated plainly, and they are not rhetorical: each is grounded in documented events, several of which are cited in the Sources section.

- **Institutions can and do act unlawfully.** It is a comfortable fiction that the police, the prosecution service, customs, the military and the judiciary operate within the law and correct themselves when they do not. In practice, across many countries, these institutions are willing to abuse their authority, and the checks meant to restrain them are frequently absent or captured. This is not confined to distant or obviously authoritarian states. It includes wealthy democracies. France is no exception, and neither are its peers. The legal apparatus of the European Union and its member states has repeatedly been turned to serve concentrated private interests — financial, pharmaceutical, military-industrial — rather than the public it claims to protect. Treating these institutions as automatically trustworthy is the single most expensive assumption a person can make.
- **The law itself can be the weapon.** The most effective attacks are not illegal; they are conducted under colour of law. A sealed indictment, an extraterritorial statute, a broadly drawn warrant, an “investigation” — each can be deployed against a target who has broken no law, to extract a settlement, remove a competitor, or simply to intimidate. The documented use of one country’s anti-corruption law to pressure a foreign rival into selling itself is not a theory; it happened, and a senior executive spent years in prison as the instrument of that pressure.
- **Surveillance capability is total and commercially available.** Assume that any device can be unlocked, that any phone can be turned into a microphone, that metadata and location are collected even when message content is not, and that the tools to do all of this are sold openly to states and to private actors. Commercial forensic tools routinely extract the full contents of supposedly secure phones. Commercial spyware has been found on the devices of journalists, lawyers and activists across the world. End-to-end encryption protects the content of

a message; it does not hide that the message was sent, by whom, to whom, or from where.

- **Reach is international.** Intelligence-sharing arrangements mean that crossing a border is not crossing to safety. The fourteen states most deeply enmeshed in such sharing — the United States, the United Kingdom, Canada, Australia, New Zealand, Denmark, France, the Netherlands, Norway, Germany, Belgium, Italy, Sweden and Spain — should be treated as a single surveillance space, within which information about a citizen of one can be lawfully obtained through another.
- **Consequences can be severe and unanswerable.** The assumptions above lead to a hard conclusion: a person may be detained without charge, have their property seized, see collaborators abroad harassed or harmed, and find no legal remedy available, because the institutions that would provide the remedy are the same ones applying the pressure. The model is built for this case not because it is the most likely on any given day, but because it is the one against which all lesser threats are also covered.

4. The Scenario.

Consider a company. Call it OilCo: a large, profitable energy firm based in the United States. Its books are accurate. Its conduct, on the matter at hand, is lawful. None of what follows depends on OilCo having done anything wrong, and in this scenario it has not.

OilCo is informed that it is being pursued by a government agency — for the sake of the scenario, the national tax authority, the IRS — for tax fraud. The agency claims to hold proof. It states that the company faces fines exceeding a million dollars and that named executives face imprisonment. OilCo reviews its records again. The numbers are correct. The documents the agency claims to hold do not match anything the company actually produced.

The question, therefore, is not the one the agency wants asked — “are you guilty?” — because the company is not. The real questions are sharper. How does an agency come to hold documents that contradict the company’s own accurate records? Were those records stolen and altered? Were they fabricated outright? Has information been twisted, out of context, into the appearance of a crime? And, most importantly: how should the company have protected itself so that no such material could exist to be used against it?

This is the inversion at the heart of the model. The honest party is the one under threat. The state agency, willing to advance a case on manufactured or manipulated evidence, is the adversary. The company’s guilt or innocence is, from a security standpoint, beside the point — because truth alone does not protect anyone once an institution has decided to act against them. What protects them is having given that institution nothing to seize, nothing to leak, and nothing to distort.

5. Protection by Role.

Security is not uniform. It is matched to exposure. The more sensitive the information a person holds — and the more useful they are as a target — the more rigorous their protection must be. The scenario above lets us set out a tiered model, from the most exposed role to the least, with the principle that the same disciplines apply at every level, only with diminishing intensity.

Two ideas run through every tier. The first is least privilege: each person holds only the access their own work genuinely requires, so that compromising any one of them exposes as little as possible. The second is need-to-know: information is divided and held in compartments, so that no single account, device or person is a master key to everything. Where access is granted, it is also made to expire, so that yesterday's necessity does not become tomorrow's liability.

- **Leadership (ultra-high).** The most senior figures hold, or can reach, almost everything: finances, contracts, plans, personnel records, relationships. This makes them the single most valuable target and the one most worth impersonating. They warrant the strongest measures available: communication conducted only through encrypted channels; the most sensitive work done on devices kept separate from everyday use; strong multi-factor authentication; the ability to decrypt the most sensitive material logged and, where possible, requiring more than one person's approval; and clean devices carrying no internal data when travelling.
- **Technical staff (ultra-high).** Those who administer the systems hold the keys to everything, and so face the same level of scrutiny as leadership, often higher. No single administrator should be able to reach every system. Their access is segmented, their credentials held in secured stores, their actions logged, and significant changes made only with more than one person's sign-off.
- **Legal advisers (high).** Counsel handling a live case receive, by necessity, highly sensitive material. That material is confined to what the case requires, held in encrypted storage, granted on access that expires, and audited and destroyed or securely archived once the matter concludes.

- **Finance and accounting (high).** Because the scenario turns on financial records, those who keep them are the most probable origin of any leaked or altered document. Their software is isolated, their actions logged in a tamper-evident way, and any external transfer of files tracked.
- **General staff (medium).** Most employees do not handle the most sensitive information, but they are the most exposed to everyday attacks — deception, malicious attachments, the careless click. They warrant encrypted communication, access limited to what their role needs, and steady, practical training in recognising manipulation.
- **Interns and temporary members (low).** Those at the edge of the organisation should not touch sensitive material at all. They work with non-sensitive or synthetic data, on tightly limited access that is removed entirely when they leave.
- **Friends and family (very low).** The outermost layer is social, and the hardest to govern, because it cannot be managed by technical means. It rests on a single understood habit: discretion about what is seen and heard. If every inner tier holds, there is nothing here to leak by accident.

The lesson of the tiers is cumulative. Security kept well at every level multiplies: it becomes disproportionately harder for any adversary to assemble, from scattered fragments, a coherent picture they can act on.

6. When the State Is the Adversary.

It is worth turning the scenario around to see the threat clearly, because the most dangerous adversary in this model is not a hacker or a competitor. It is an institution acting under the authority of the state, against a target who has done nothing wrong.

Such an adversary does not need to break the law, because it writes and enforces it. Its methods are procedural. Evidence obtained improperly can be re-derived through a parallel route built to look legitimate. A warrant can be granted on thin or misleading grounds, and used to seize every device and record a person holds — not because they are evidence, but because the seizure itself is disruptive and because, once data is taken, the owner loses all control over how it is read, used or leaked. Detention can be prolonged, in some jurisdictions for months, without a charge ever being brought, on the broad pretexts of financial crime or national security. And throughout, the institutions that should provide a remedy — the courts, the oversight bodies — may be the very ones applying the pressure, or unwilling to act against them.

There is a second figure who matters here: the person who, from conscience, brings wrongdoing to light. Whether an insider reporting genuine fraud or a source assisting an investigation, this person is the most fragile point in any case against the powerful, and the protection of their identity is paramount. It is protected by reducing contact to the necessary minimum, by stripping every trace of metadata and origin from what they provide, by referring to them only through a code rather than a name, and by holding the link between identity and information so tightly that not even most of those working on the case could reveal it under pressure. The principle is operational deniability: arrange things so that there is nothing to give up, even if someone is made to try.

7. Outside Forces.

Beyond the direct contest, three further forces shape the environment, and each requires its own posture.

- **International reach.** When an adversary can act across borders, the protection of national law disappears, and the laws of the adversary's jurisdiction may be turned against you. The pattern is recognisable: collaborators abroad — a journalist in one country, a researcher in another — suffer a convenient office fire, a burglary in which only the devices and files vanish, an investigation that is closed with unusual speed. These are not coincidences; they are signals of an adversary with surveillance power, insider access and global mobility. The practical conclusions are to treat the major intelligence-sharing states as one space, to assume metadata and hardware interception are in play regardless of encryption, and, where the stakes justify it, to keep human and digital assets in more neutral jurisdictions.
- **The media.** The press is both shield and hazard. It can expose injustice and rally support; it can also, through carelessness or appetite for a story, endanger the very sources it relies on, or be used as cover for an approach that is really an attack. The posture is controlled disclosure: deciding in advance who may speak to the press, releasing curated and accurate information through trusted journalists rather than reacting to leaks, and getting ahead of a damaging narrative instead of chasing it.
- **Public opinion.** The largest and least predictable force is the public, whose view is shaped by emotion, imagery and speed far more than by fact. Support can curdle into condemnation on a single viral headline, and coordinated networks can amplify a falsehood faster than any correction. The defence is not censorship, which fails and inflames; it is pre-emptive transparency — publishing clear, accurate timelines and rebuttals before distortions take hold, and keeping the human reality of the situation visible, because people relate to people, not to documents.

A final point ties this section to the whole: when an adversary may try to prevent you from speaking at all — by seizing your work, or detaining those who carry it — the ability to publish independently becomes a security

measure in its own right. Information that is easy to find, hard to remove, and not dependent on any single platform is far harder to suppress.

8. Techniques.

The mindset and the scenario point to a set of technique categories. These are described as capabilities, not products. Specific tools change constantly, fall in and out of trust, and vary by country and situation; what endures is the function each technique performs. The choice of any particular application is exactly the decision this threat model exists to inform, and is properly the subject of a separate, living review.

- **Encrypt everything at rest.** Full-disk encryption with a strong passphrase, so that a seized or stolen device yields nothing. For the most sensitive material, consider storage whose very existence can be plausibly denied.
- **Encrypt everything in transit.** End-to-end encrypted messaging and email as the default for any communication that matters, with the understanding that this protects content, not the fact of contact.
- **Separate identities and devices.** Keep the most sensitive work off everyday machines. Where it is warranted, use dedicated communication channels and accounts that exist only for a purpose and are not linked to personal life.
- **Disposable channels for the most fragile contact.** For the protection of sources, single-use accounts and contact methods that retain no metadata and are discarded after use.
- **Strong, layered authentication.** Long, unique credentials held in an encrypted store, combined with a second factor — preferably a physical one — so that a stolen password alone opens nothing.
- **Reduce the device's reachability.** Disable radios when not in use; understand that a powered-on device is a reachable device, and power down what handles sensitive material when the work is done. For the highest sensitivity, physical isolation from any network.
- **Plan for seizure.** Decoy accounts and innocuous data on devices likely to be inspected; the assumption that anything carried across a border may be copied.
- **Legal pre-arrangement.** Counsel retained before trouble arrives, not after; a trusted third party briefed in advance to alert a lawyer, the press or a consulate if a member is detained.

- **Independent, resilient publishing.** The means to put information into the open in a way that is easy to find and difficult to remove, so that suppression of the messenger does not suppress the message.

The point of the list is not to adopt all of it indiscriminately. It is to match each technique to the role and exposure set out in Section 5, so that effort is spent where it changes the outcome.

9. The Discipline.

Tools and techniques rest on a handful of older habits that no software can supply. They are the discipline beneath the method.

- **Live a believable cover, and keep work and private life apart.** Let no part of your conduct invite a second look, and never let the boundary between the two blur.
- **Need to know.** Hold only what your own task requires, and do not seek out what does not concern you. What you do not know cannot be taken from you.
- **Behave naturally and blend in.** Draw no attention. The aim is to be unremarkable.
- **No loose talk.** Assume that channels, calls and public spaces may be listened to. Conceal sensitive details, and say as little as the work demands.
- **Punctuality is a signal.** Be on time, and read lateness as information: an unexplained absence may mean something has gone wrong.
- **Prepare beforehand, and vary your patterns.** Make arrangements in advance, and avoid the predictability that makes a person easy to follow.
- **Leave no traces.** Be tidy. Do not leave sensitive material lying about, and do not keep it where you live; commit what you can to memory rather than writing it down.
- **Protect those beside you.** Under pressure or questioning, reveal nothing of the people you work with. If one member is compromised, assume their access may be turned against the rest, and act accordingly.

The spirit of all of it is captured in a single line worth remembering: amateurs practise until they get it right; professionals practise until they cannot get it wrong.

10. Conclusion.

This model illustrates how complex, asymmetric and borderless the protection of information has become. The honest party may be entirely in the right and still be destroyed, because truth alone no longer guarantees protection. What endures a crisis is not innocence but preparation — layered, role-matched, and practised before it is needed.

A threat model is therefore not a checklist to be completed once and filed. It is a living framework, recalibrated as the adversary and the tools evolve. Its value is in the question it teaches everyone to ask, of every tool and every habit: against this adversary, in this situation, does this actually protect the people who depend on it. Read the scenario, find your own position within it, and raise your defences to match. The cost of doing so is small. The cost of assuming you will never need to is the one no later precaution can undo.

11. Sources.

The scenario in this document is fictitious. The capabilities and conduct it assumes are documented in, among others, the following.

01. Frédéric Pierucci, The American Trap (with Matthieu Aron) — first-hand account of the use of an extraterritorial anti-corruption statute to pressure a foreign company. https://en.wikipedia.org/wiki/Fr%C3%A9d%C3%A9ric_Pierucci

02. "Jailed French executive who felt force of US bribery law", BBC News, 2019. https://www.bbc.com/news/world-europe-47765974

03. "Zones d'ombre sur la vente de la branche énergie d'Alstom à GE", Le Monde, 2019. https://www.lemonde.fr/economie/article/2019/01/15/zones-d-ombre-sur-la-vente-de-la-branche-energie-d-alstom-a-ge-temoignage-d-un-ancien-cadre-emprisonne-deux-ans-aux-etats-unis_5409271_3234.html

04. "The Feds Can Now (Probably) Unlock Every iPhone Model In Existence", Forbes, 2018 — commercial phone-unlocking capability. https://www.forbes.com/sites/thomasbrewster/2018/02/26/government-can-access-any-apple-iphone-cellebrite/

05. Cellebrite iPhone cracking capability by model. https://9to5mac.com/2022/04/29/cellebrite-iphone-cracking/

06. Pegasus spyware — use against journalists, lawyers and activists. https://en.wikipedia.org/wiki/Pegasus_(spyware)

07. "Journalists Targeted with Pegasus Spyware", Amnesty International, 2025. https://www.amnesty.org/en/wp-content/uploads/2025/03/EUR70918622025ENGLISH.pdf

08. Spyware litigation tracker — legal challenges relating to mercenary spyware. https://citizenlab.ca/spyware-litigation-tracker-legal-challenges-and-formal-complaints-related-to-mercenary-spyware/

09. Edward Snowden, Permanent Record — first-hand account of mass surveillance.

10. "Cybersecurity for the People", The Intercept, 2017 — practical communications security.

https://theintercept.com/2017/05/01/cybersecurity-for-the-people-how-to-keep-your-chats-truly-private-with-signal/

This document is research, published for discussion. The scenario is fictitious; the references are real.