Analysis · May 26, 2026

← Previous·Next →

Real vs apparent privacy: the questions worth asking

An operational synthesis of cycle 2: the questions that distinguish a service with architectural privacy from one with declarative privacy. A questionnaire for the European professional before adopting any digital tool for sensitive data.

In plain terms: Two services with the same legal notice can behave very differently. One protects by technical design. The other protects by contractual promise. The difference is not read in the notice — it is discovered by asking the concrete questions. The quality of the answers says as much about the product as their own content.

The difference between architectural privacy and declarative privacy

Across the seven previous articles of this cycle we have moved through different layers of the same matter. The law of international transfers with Schrems II. The mathematical idea of the cryptographic hash that seals each Cuaderno. The architectural choice of the kill switch and the institutional capture that almost always accompanies it. The mechanism of end-to-end encryption and the operational question of where the keys reside. The alignment of incentives according to the business model. Self-sovereign cryptographic identity. Self-hosting as a proportional strategy. Each article dealt with one angle. This one, the last of the cycle, brings them together into a questionnaire.

The distinction worth retaining is simple: there are services whose privacy is *architectural* and there are services whose privacy is *declarative*. The first is embedded in the technical design: certain violations of the privacy commitment are technically difficult or impossible because the architecture does not allow them. The second is deposited in the text of the legal notice: certain violations would be contractually sanctionable if they occur, but nothing technically prevents them. Both models can comply with the GDPR; but one protects by construction and the other protects by promise, and the difference is operationally enormous.

The questions that follow are designed to distinguish one case from the other. They are not advanced technical questions. They are the questions that any honest provider can answer in its public documentation. The quality and precision of the answer says as much about the product as the answer itself. The questions are grouped into six layers; it is worth asking them all before adopting the service for sensitive data, not only the ones the first instinct identifies.

Layer 1: architecture

Let us pin down a term before going on. By *operator* we mean the company that provides the service: the entity that controls the servers and the software, not a specific person. With that clarified, the root architectural question is: what does the operator do with the content between sender and recipient? There are three possible answers, and it is worth learning to tell them apart, because all three are sometimes advertised with similar vocabulary.

- The first: the content passes through an operator server in the clear, where the operator can read it even if it promises not to.
- The second: the content passes through an operator server encrypted, where the operator cannot read it if the keys reside exclusively on the users' devices.

- The third: the content passes through no operator server, because no operator server exists in that specific flow.

The difference between these three is not one of degree: it is one of kind.

The complementary question —already formulated in the Cuaderno on encryption— is: who holds the cryptographic keys that allow the content to be read? If the user and only the user holds them, the encryption is real. If the operator also holds them in any form —even under the name «account recovery» or «device synchronisation»—, the encryption is nominal. The question admits no honest intermediate answer.

Layer 2: business model

The question about the business model matters as much as the architectural question, and for the same substantive reason: incentives produce, over time, systematically different products even with identical declared purposes. How does the operator make money today? A single source, two, a mix? If funding includes advertising or data monetisation, what data is monetised and on what GDPR legal basis is it done? Does the purpose declared in the legal notice cover the third-party data that the professional intends to entrust to the service?

And the second-order question, not always formulated: what is the operator's financial situation at a three- to five-year horizon? A company in a venture-capital phase operates under different pressures from a company in stable profitability. The change of funding model is, repeatedly, the moment at which the implicit contract with users is rewritten without negotiation.

Layer 3: jurisdiction

For the European professional, the question of jurisdiction is not rhetorical. In which jurisdiction is the operator incorporated? In which country are the servers that process the data physically located? Is the answer to the two previous questions the same or different, and if it differs, which law applies? A European region operated by a US company is not, for the purposes of Schrems II, a European answer: the company is subject to FISA 702 regardless of where the servers are.

The complementary operational question is: if an intelligence order valid in the operator's jurisdiction arrived tomorrow demanding the handover of my data or that of my clients, what would happen? If the honest answer begins with «the company would be obliged to hand them over», the service does not protect against that order no matter how much the advertising suggests otherwise. If the honest answer begins with «the company could not hand them over because it does not hold them in the clear», the service does protect; and the difference depends almost entirely on the first two layers, not on the quality of the privacy policy.

Layer 4: operator and kill switch

What technical capacity does the operator retain to remotely suspend, block, delete, or degrade the service? The question is not paranoid: it is operational. Digital platforms have exercised that capacity repeatedly in recent years, sometimes on their own initiative, sometimes under government order, sometimes after changes of ownership or policy. If the capacity exists, it is worth knowing under what contractually declared scenarios it is exercised, and reserving a margin for the undeclared scenarios that the practice of recent years has shown to be just as relevant: an unexpected court order, an international sanction, a change of corporate governance, an acquisition by an entity with a different policy.

The sister question is that of the continuity plan: if the operator were to exercise the capacity against the professional —for whatever reason, fair or not—, how much uptime would remain available, what data-export

procedure exists, and to which alternative provider could one migrate? If the answer begins with «it shouldn't happen», it is not an operational answer; it is a promise.

Layer 5: identity and access

Who controls the credentials for accessing the service? If the operator can reset the user's access without the user's participation —a procedure typically called «account recovery»—, the operator is, technically, the custodian of the account and can also hand it over to whoever requests it through the appropriate procedure. If the operator cannot reset the access because the identity resides cryptographically on the user's device, the operator also cannot hand it over, not even under order. Both modalities are legitimate depending on the context; but, once again, they are different, and it is worth knowing which one is being adopted.

What happens to the professional's data if the professional loses access? Are there recovery mechanisms —of account, of file, of session— that depend on the operator? Are those mechanisms compatible with the sector's professional ethics if the operator is coerced into using them?

Layer 6: future

This last layer tends to be neglected because it demands projection. What would happen if the service were acquired by another company? Almost all acquisitions bring with them a revision of the terms of service in the following months. What would happen if regulatory requirements changed? European law has increased takedown and blocking obligations since 2022, it has not reduced them. What would happen if the operator disappeared? A significant portion of cloud services has no documented exit plan for the scenario of the operator shutting down; the professional discovers the problem when there is no longer time to prepare for it.

There is a formulation worth retaining for this layer: architectures that depend less on the operator are more resilient to changes in the operator. Self-hosting in any of its modalities, self-sovereign cryptographic identity, communications with no server in the middle, all of these reduce the future risk surface through the procedure of reducing the present dependency surface. They do not eliminate it; they reduce it.

The difference between structure and promise

If we had to distil the cycle into a single sentence, it would be this: structural answers hold even if the operator, the administration, or the legislation changes; promise-based answers hold as long as the one who promises can and wants to keep them. Both can be correct at the moment of adoption. Only one of the two holds up independently of the passage of time and the change of circumstances.

This does not mean that every professional must demand structural answers from all the services they adopt. Proportionality remains legitimate: a spreadsheet for internal accounting does not need the same answer as a patient's clinical record. It does mean, yes, that professionalism consists of knowing what kind of answer has been accepted in each case, and of having consciously decided that that kind of answer is proportional to the specific datum.

The questionnaire, in order

Twelve concrete questions that synthesise the cycle, ordered so that the answer to each one informs the next:

1. Does the content pass through an operator server? If it does: in the clear, encrypted with the operator's keys, or encrypted with keys exclusive to the user?
2. If end-to-end encryption is invoked, where do the cryptographic keys reside? Does the operator know or retain any part of them in any form, including «recovery»?
3. What metadata does the service generate and retain? For how long? To whom is it visible?

4. How is the operator funded? If funding includes advertising or data monetisation, does the declared purpose cover third-party data entrusted by the professional?
5. What is the operator's financial situation at a three- to five-year horizon? Are there factors suggesting an imminent change of model (a pending IPO, a financing round running out, a probable acquisition)?
6. In which jurisdiction is the operator incorporated? In which country are the servers physically located? If they differ, which national law applies to the processing?
7. What would happen if an intelligence order valid in the operator's jurisdiction demanded the handover of my data? Could the company technically comply with it?
8. What technical capacity does the operator retain to suspend, block, or delete the service? Under what contractual scenarios? Under what historically documented non-contractual scenarios?
9. What exit plan exists if the operator were to exercise that capacity against me, fairly or unfairly? Is there a documented procedure for exporting data to an alternative provider?
10. Who controls the access credentials? Can the operator reset them without my participation? Does that protect me or expose me?
11. Is there a European, self-hosted, or no-server-in-the-middle alternative for this specific function? What is its real cost, compared with the assessed risk?
12. If today's decision were examined five years from now by an inspector, an auditor, or a client affected by a breach, would the present choice be defensible with the arguments available today, or would it require an apology for not having asked reasonable questions?

The questions do not expect perfect answers. They expect honest answers, which the honest operator knows how to give and the less honest operator avoids formulating with precision. The operational difference between the two classes of operator, we say it without drama, is usually perceived by reading slowly the answers they offer voluntarily, even before having to ask for more.

With this article we close the second cycle of Cuadernos Lacre. We began with the editorial debt inherited from Schrems II and we end with an operational questionnaire. Along the way we have moved through concepts —hash, encryption, identity— and applied analyses —kill switch, business model, self-hosting—. The publication's declared editorial intent was not to overwhelm the reader with the exhaustive list of problems, but to hand them tools so that, faced with any new service, they can tell what kind of answer they are accepting. That distinction —between architecture and promise— is the tool. The rest each professional will put at the service of the data they consider, in their own practice, worthy of the question.

Sources and further reading

- This publication, cycle 2 (May 2026) — *Schrems II, five years on*, *What SHA-256 really is*, *Kill switch and institutional capture*, *End-to-end encryption, truly explained*, *The business model as a signal of trust*, *The 24 words: what a cryptographic identity is*, *Self-hosting as a professional practice*. The seven articles on which this questionnaire rests.
- Regulation (EU) 2016/679 — General Data Protection Regulation. The reference legal framework for all the questions the questionnaire raises, in particular articles 5, 6, 25, 28, 32, 33 and chapter V.
- European Data Protection Board — guidelines and operational opinions on Schrems II, international transfers, impact assessments and proactive accountability (publications 2020-2024).
- Spanish Data Protection Agency — penalties published 2022-2024 against data controllers for inadequate transfer instruments or for formal impact assessments without substantive content.
- noyb.eu — European Center for Digital Rights, led by Maximilian Schrems. A public repository of complaints, appeals and analysis on the real, not apparent, compliance with European data protection rules.

← PreviousSelf-hosting as a professional practiceNext →What a signature cannot fix

Recent readings

- Reflection · June 29, 2026 You Are Not Anonymous

- Reflection · May 27, 2026 What a signature cannot fix
- Analysis · May 25, 2026 Self-hosting as a professional practice

Take this article wherever you need it.

↓ Markdown ↓ Plain text ↓ PDF

The file is downloaded to your device. From there, you can save it, import it into Solo2, or share it as you wish. Cuadernos does not decide the destination for you.

Wax seal · SHA-256 ed96e76396957f3fda1045c6e4e522659b8d150ed991a6b1e11474569f737533

Cuadernos Lacre · A publication of Menzuri Gestión S.L. ·
written by R.Eugenio · edited by the team of Solo2.