SLTA | SWISS LEGALTECH ASSOCIATION

LAW&TECH INSIGHTS MAGAZINE

2026-9

SLTA TECH & LAW INSIGHTS

AI BROWSERS AND DATA PROTECTION

A SWISS AND EUROPEAN PERSPECTIVE FOR THE LEGAL PROFESSION

NICOLAS TORRENT, SYLWIA POMIESZCZ

Editorial Team :

Nicolas Torrent, Rucsandra Bratu,
Sylwia Pomieszcz, NomungerelJamsranjav

INTRODUCTION

A new generation of consumer browsers now embeds a large language model directly into the browsing session, with OpenAI's Atlas and Perplexity's Comet as the most visible examples. Whereas a traditional browser renders pages when the user requests them, an AI browser observes the entire session, retains a memory of it, and, in agent mode, acts within it: filling forms, sending messages, and executing transactions on behalf of the user.

This shift matters for lawyers. The tool is no longer confined to processing what the user chooses to submit; it now processes what the user does throughout the session. Browsing history reveals categories of data that Swiss and European law treat with particular caution, and the professional secrecy obligations that govern the legal profession do not bend to accommodate consumer product design.

This article sets out the data protection issues raised by AI browsers under the GDPR and the Swiss Federal Act on Data Protection (FADP), together with the professional conduct layer that applies specifically to Swiss lawyers. It also flags two recent US developments that illustrate where the confidentiality question is heading in practice.

1. DATA COLLECTION, PROFILING AND SENSITIVE DATA

AI browsers process the entire browsing session, and not only the page on which the user is asking a question. Atlas's "browser memories" feature stores descriptions of every site visited, and the setting that allows the content of visited sites to be used to improve the model is enabled by default. Opting out restricts what the provider may do with the data, but does not restrict what is collected in the first place.

The legal consequence is easily underestimated. Browsing history inevitably reveals information about health, political opinions, religious beliefs or sexual orientation, all of which fall under the definition of sensitive data (Art. 9(1) GDPR; Art. 5(c) FADP). The Court of Justice of the European Union has confirmed in *Meta Platforms v Bundeskartellamt* that activity data revealing such information falls within the sensitive-data regime regardless of whether the data subject explicitly revealed such data or it is inferred from activity data.[1]

2. LAWFUL BASIS, MINIMISATION AND GDPR PRINCIPLES UNDER STRAIN

The lawful basis for session-wide processing and for model training is genuinely unclear at this stage. The current reference point is Opinion 28/2024 of the European Data Protection Board of 17 December 2024, which indicates that legitimate interest for the development and deployment of AI models, as well as the question of model anonymity, are to be assessed on a case-by-case basis.[2]

Memory and personalisation features sit uneasily with the principles of purpose limitation and data minimisation (Art. 5(1)(b) and (c) GDPR; Art. 6 FADP). Persistent agent memory can prevent the effective exercise of the rights to access, rectify, erase or export personal data, since the user cannot readily identify what has been retained, where, or in what form. Autonomous agent actions that produce legal or similarly significant effects on the user or on third parties, such as an unintended purchase, a message sent in the user's name, or a contract concluded on the user's behalf, fall within the rules on automated individual decisions (Art. 22 GDPR; Art. 21 FADP).

The allocation of responsibility for such actions is a real question. Depending on the facts, responsibility may rest with the user who granted the agent access, with the deployer that operates the tool at the organisation level, or with the vendor whose product acted outside the boundaries the user could reasonably have expected. For a law firm or in-house team, the safe working assumption is that liability for what the agent does will start with the entity that deployed it, and that recourse against the vendor is likely to be a subsequent, contractual matter.

The closest guidance to date comes from two regulatory interventions on agentic AI published in February 2026. Neither targets AI browsers specifically, but both apply directly to the agent mode that these browsers now offer. The Spanish data protection authority published an 81-page guidance on agentic AI on 18 February 2026,[3] and the Dutch data protection authority issued a formal warning on 12 February 2026 concerning autonomous agents with broad system access.[4] Both interventions converge on a single message: autonomy does not dilute accountability. The agent is a technical means for which the deployer remains responsible, and neither open-source nor local deployment discharges data protection obligations.

A further layer applies specifically to the way the browser stores its memories. Under the ePrivacy framework (Art. 5(3) of Directive 2002/58/EC), the storage of information on, or the reading of information from, a user's device requires the user's prior informed consent, unless the storage is strictly necessary for a service the user has requested. The "browser memories" feature does not obviously qualify as strictly necessary, which means that a distinct consent layer applies in addition to the GDPR analysis.

3. ROLES, DPIA AND ORGANISATION-LEVEL GOVERNANCE

The qualification of the browser vendor as controller or processor remains unresolved in practice, and most consumer AI browsers do not offer an Article 28 GDPR data processing agreement at all. This absence leaves the deployer without a proper contractual framework for handling the data flows that the tool generates.

On the question of whether a data protection impact assessment is required (Art. 35 GDPR; Art. 22 FADP), our position is that session-wide agent access is high-risk processing and calls for a DPIA before deployment. Three factors support this qualification: the systematic monitoring of an individual's activity across all services accessed via the browser, the routine involvement of sensitive data as identified in section 1, and the autonomous execution of actions that can produce legal or financial effects. Any one of these would already point to a DPIA obligation, and the three combined leave little room for argument against.

A related concern arises at the organisational level. An employee who installs an AI browser on a work device effectively causes the organisation to become the controller of whatever the tool ingests, and, depending on how the data is used downstream by the vendor, potentially a joint controller with the vendor as well. The pragmatic response is to treat AI browsers as managed software rather than as personal preferences: block them by default at endpoint level, allow them only after the DPIA referred to above has been completed, and document the internal authorisation. In the Swiss employment context, session-wide monitoring on work devices additionally engages the restrictions on systems that monitor employee behaviour at the workplace (Art. 26 of Ordinance 3 to the Employment Act).

4. CROSS-BORDER TRANSFERS

Session data routinely flows to US-based providers, which brings the rules on international transfers into play (Chapter V GDPR; Art. 16 FADP). Compliance requires either reliance on an adequacy mechanism, namely the EU-US or Swiss-US Data Privacy Framework where the provider is certified, or the use of standard contractual clauses together with a transfer risk assessment along the lines set by *Schrems II*.[5]

For Swiss lawyers, a distinct additional concern arises: the exposure of client data to foreign lawful-access regimes, notably under the US CLOUD Act. This concern is one of the main reasons why Swiss firms tend to avoid non-Swiss and non-EU cloud AI tools, and it is well documented in Swiss practice.[6]

5. THE SWISS LAYER: PROFESSIONAL SECRECY AT THE CORE

The FADP is technology-neutral and applies fully to AI-supported processing. The Federal Data Protection and Information Commissioner (FDPIC) has been explicit about transparency expectations, requiring clarity on the purpose, functionality and data sources of AI-based processing, and requires a DPIA in high-risk cases.

For practitioners, the most consequential layer is professional secrecy. Article 321 of the Swiss Criminal Code and Article 13 of the Federal Act on the Free Movement of Lawyers (BGFA), reflected at the deontological level in Article 4 of the Swiss Code of Professional Conduct, protect confidences entrusted to lawyers. An AI browser that observes a lawyer's entire session, including client portals, e-mail and document platforms, exposes those confidences to a third-party vendor and therefore raises a secrecy issue, and not merely a privacy issue. Article 62 FADP adds a broader statutory duty of confidentiality that extends beyond the Article 321 remit.

The Swiss Bar Association has issued dedicated Guidelines on the Use of AI, which serve as the anchor for this section.[7] These Guidelines set out deployment options for compliant use, and identify locally operated solutions, in which data does not leave the firm's infrastructure, as the most appropriate. Consumer AI browsers cannot meet that criterion in their current form: their operating model routes session content and derived memories to a cloud vendor outside the firm's infrastructure, and this routing is not an incidental feature but the mechanism through which the product functions. The Guidelines' preferred deployment model is therefore not one that these products can implement.

Two further conduct rules bear directly on AI browsers. The first is the duty of diligence under Article 12(a) BGFA, read in current commentary as encompassing an extended duty of technological competence. The second is Article 35 of the Swiss Code of Professional Conduct, under which unsecured digital communication requires the client's consent.

The sanction regime deserves particular emphasis for a lawyer audience. Under the FADP, sanctions target responsible individuals, with fines reaching up to CHF 250,000 (Arts. 60–63 FADP). This is a materially different conversation than GDPR corporate fines, since the exposure is personal.

The broader regulatory context should be kept in mind. Switzerland signed the Council of Europe AI Convention on 27 March 2025, and the EU AI Act (Regulation (EU) 2024/1689) reaches general application in August 2026, a date that matters for Swiss firms serving EU clients.

6. SECURITY AS A STRUCTURAL LEGAL RISK

Indirect prompt injection affects the entire product class, and security researchers report that it cannot be fully patched. Hidden content on a page can instruct the agent, which then acts with the user's full credentials and sessions, effectively bypassing same-origin protections. From a legal perspective, this undermines any argument based on "appropriate technical and organisational measures" (Art. 32 GDPR; Art. 8 FADP).

A striking illustration comes from a phishing benchmark conducted by a security vendor in September 2025 against the 100 most recent OpenPhish and PhishTank attacks: Edge blocked 54 % of the malicious pages, Chrome 47 % and Dia 46 %, while Comet blocked only 7 %. Later testing put Atlas at approximately 5.8 %.[8] These figures come from a security vendor, predate later patches, and should therefore be treated as indicative rather than definitive. The order-of-magnitude gap between traditional browsers and AI browsers remains the relevant point.

7. CONFIDENTIALITY, PRIVILEGE AND LITIGATION TO WATCH

Two recent US developments illustrate where the confidentiality question is heading in practice.

The first is United States v. Heppner, a bench ruling of 10 February 2026 followed by a written opinion of 17 February 2026.[9]The court held that communications between a criminal defendant and the consumer version of a generative AI assistant were not protected by either attorney-client privilege or the work product doctrine. The primary ground was that the AI is not an attorney, and the exchanges therefore could not satisfy the client-attorney requirement. The court added, in the alternative, that the consumer privacy policy, which permits retention and training, would in any event have defeated any reasonable expectation of confidentiality, and that the use of AI had not been undertaken at counsel's direction.

Commentators are divided on how far the ruling reaches. Some read it as veering toward a near-categorical exclusion of client AI use, while others read the holding narrowly as fact-specific, turning on consumer-grade use undertaken on the defendant's own initiative rather than at counsel's direction. The narrow reading is the one that matters here. Applied to AI browsers, which observe the entire session, the confidentiality concern generalises, but the outcome remains fact-dependent rather than a blanket rule.

The second development is Amazon v. Perplexity, a preliminary injunction of March 2026 currently on appeal (Ninth Circuit appeal argued on 11 June 2026), and the first federal appellate test of whether an AI browser agent may access a logged-in third-party platform at a user's direction.[10] The dispute turns on unauthorised-access law, namely the US Computer Fraud and Abuse Act and California Penal Code § 502, rather than on data protection, but it bears directly on the agent mode that these browsers now offer.

It is a data point in itself that several compliance advisories now recommend blocking AI browsers outright in any environment handling client or otherwise confidential data.

CONCLUSION

AI browsers concentrate, within a single consumer product, the most difficult questions raised by the current AI moment: session-wide profiling that inevitably touches sensitive data, an unclear lawful basis for model training on personal data, cross-border transfer exposure, unresolved controller allocation, structural security weaknesses that no vendor patch can fully close, and, for Swiss lawyers, a professional-secrecy regime that these tools cannot presently satisfy.

The EU framework supplies the analytical grid: GDPR principles, EDPB Opinion 28/2024, the Spanish and Dutch interventions on agentic AI, and the EU AI Act. The Swiss framework adds the FADP transparency and DPIA expectations and, critically, the professional-conduct layer: Article 321 of the Swiss Criminal Code, Article 13 BGFA, the Swiss Code of Professional Conduct, the Swiss Bar Association's AI Guidelines, and the personal criminal liability of individuals under Articles 60–63 FADP.

For the Swiss legal profession, the practical takeaway is that consumer AI browsers, in their current form, cannot be reconciled with professional secrecy on work devices, or on any device used to access client matters. The compliant path runs through locally operated solutions, organisation-level governance, informed client consent where applicable, and a DPIA that treats session-wide agent access as high-risk by default.

The US developments in Heppner and Amazon v. Perplexity are not directly transposable, but they signal the direction of travel. Confidentiality claims over consumer AI use will not be presumed, and the boundaries of agent-mediated access to third-party services are being drawn now. Swiss firms that adopt AI browsers without a governance layer are, in practical terms, underwriting that uncertainty on personal exposure.

REFERENCES

1. GDPR Arts. 5, 6, 9, eur-lex.europa.eu/eli/reg/2016/679/oj; CJEU C-252/21, *Meta Platforms v Bundeskartellamt,* ECLI:EU:C:2023:537.
2. European Data Protection Board, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models (17 December 2024), edpb.europa.eu.
3. Agencia Española de Protección de Datos (AEPD), guidance on agentic AI (18 February 2026) (original in Spanish). Analysis: Covington (Inside Privacy); Alston & Bird.
4. Autoriteit Persoonsgegevens (AP), warning of 12 February 2026 on autonomous AI agents (OpenClaw and similar), autoriteitpersoonsgegevens.nl.
5. GDPR Chapter V; CJEU C-311/18, *Data Protection Commissioner v Facebook Ireland and Schrems* (Schrems II), ECLI:EU:C:2020:559.
6. Chambers and Partners, *Artificial Intelligence 2026 — Switzerland: Trends and Developments,* practiceguides.chambers.com.
7. Fédération Suisse des Avocats / Swiss Bar Association, *Guidelines on the Use of AI,* digital.sav-fsa.ch/en/ki-guidelines. See also FDPIC, AI and data protection; FADP official text, Fedlex; Swiss Code of Professional Conduct, sav-fsa.ch.
8. LayerX Security, phishing benchmark research (September 2025), layerxsecurity.com. On indirect prompt injection as a structural class problem: AIMultiple.
9. *United States v. Heppner,* No. 1:25-cr-00503-JSR (S.D.N.Y., bench ruling 10 February 2026; written opinion 17 February 2026, Rakoff, J.). Commentary: Harvard Law Review Blog; Gibson Dunn; Debevoise; Akin Gump.
10. *Amazon v. Perplexity,* N.D. Cal. preliminary injunction (March 2026); Ninth Circuit appeal argued 11 June 2026. Commentary: Eric Goldman, Technology & Marketing Law Blog.