AI Policy

Petit Legal – use of generative AI and AI tools

Document	Details
Version and date	1.0 – 13 November 2026
Lawyer	Sander Petit
Firm	Petit Legal

1. Purpose and scope

This policy governs how the Lawyer and the employee(s) of Petit Legal, and anyone else working temporarily for the firm, such as trainees, temporary staff, locums and contracted freelancers (together: staff), use AI tools. The policy is deliberately proportionate and practical, designed for a sole practice with one employee. Its aim is for the firm to benefit from the opportunities of AI without eroding the core values of the legal profession (independence, partisanship, confidentiality, integrity and expertise) or its legal obligations.

The policy applies to all AI tools that process or generate text, speech or other data, including where they are built into another program (such as a summary or translation feature), and regardless of whether the tool runs on a firm device or a private device.

This policy follows the revised Recommendations on AI in the Legal Profession issued by the Netherlands Bar (NOvA) (November 2026). Those recommendations are guidance and do not in themselves create obligations; where they reflect an existing obligation (such as the duty of confidentiality, the GDPR or the AI Act), their binding nature derives from that obligation. The policy applies to all forms of AI, including generative AI (text, image, audio and video), AI-assisted search and analysis services, and agentic AI systems that perform tasks autonomously.

This is the public version of the AI policy of Petit Legal, to which the quotation refers (section 6). The internal supplier assessment and working documents do not form part of it.

Status of this policy. This policy describes how the firm uses AI and which safeguards it applies in doing so. It is an internal framework and does not create any obligations or guarantees towards clients or third parties beyond what the law, the professional rules and the engagement require. The legal relationship with the client is governed by the quotation or engagement letter and the firm's general terms and conditions, including the limitation of liability contained therein. Time limits and working methods in this policy are starting points: the Lawyer may deviate from them in a specific case, giving reasons, provided that the statutory and professional obligations are complied with. The firm may amend this policy; the published version states the version number and date.

2. Principles

- **Final responsibility.** AI supports, but does not steer. The Lawyer remains ultimately responsible for every piece of advice, every document and every act, including where AI has been used.
- **Confidentiality first.** Professional secrecy and legal professional privilege take precedence over convenience. Client data is only entered into a tool approved for that purpose (sections 4 and 5).
- **Human review.** AI output is never adopted without review (section 8).

- **Transparency.** The firm is clear, internally and towards clients, about how AI is used (section 6).
- **Careful selection and documentation.** A documented assessment is made in advance for each tool (section 7).
- **Risk-based.** Risk arises at the intersection of the tool and its specific use: the same tool may be harmless in one application and not in another. The register (Annex A) therefore also records, for each tool, the uses for which it may be used, may be used only with additional safeguards, or may not be used.
- **Data minimisation.** No more information is shared than is reasonably necessary for the purpose (section 5).
- **Independence and partisanship.** AI is never decisive in representing the client's interests (section 8).

3. Roles and responsibilities

Role	Responsibility
Lawyer	Adopts the policy, decides on the admission of tools, carries out the supplier assessments, assesses whether a DPIA is required, maintains the register (Annex A), decides per matter whether AI is used, obtains the client's consent and reviews the output. The Lawyer is ultimately responsible for compliance, including where the employee has used AI. The Lawyer is also the person responsible for AI within the meaning of the NOvA recommendations: approves the policy and tools, negotiates data processing agreements, carries out the annual review and is the escalation point for incidents.
Employee	Follows this policy and the training, uses only tools approved by the Lawyer, refers any doubt about a category or tool to the Lawyer and reports incidents immediately.
Locum	When the Lawyer is absent, the Lawyer arranges a locum. The locum is bound by this policy and by professional secrecy, and is given access only to the tools and data needed for the locum work.

4. Permitted tools

Only tools on the list in Annex A may be used for firm work. Each tool has one of three statuses. A tool that is not on the list may not be used for client data until the Lawyer has assessed it (section 7).

The status applies per tool; admission applies per use. Even a green tool is used only for the uses listed as permitted in Annex A, and the additional safeguards listed there apply to enhanced-safeguard uses. Whether a tool is free or paid is not a criterion: what matters is how the actual processing is arranged contractually and technically.

Status	Meaning	Conditions
Green	May be used with client data.	In principle only with a signed data processing agreement (exceptions with accepted residual risk: Claude Pro and Pocket, see below), a completed assessment, a DPIA where required under Article 35 GDPR, and with the client's consent for the matter (section 6). Only through the firm's business subscription.
Yellow	Only without client data and without confidential firm data.	Suitable for general research and use with public information or case descriptions that pass the recognisability test in section 5.

Status	Meaning	Conditions
Red	Not permitted for firm work.	Applies to tools without a data processing agreement, with unclear data processing or insufficient safeguards, and to consumer subscriptions.

If a client consents to the use of an AI tool, that consent takes precedence over the above status. For a tool without a data processing agreement with the firm, this applies only to Claude Pro, under the conditions below; for Pocket, the arrangements in section 9 apply. The client's consent does not cover the personal data of opposing parties, witnesses and other third parties: such data is omitted or replaced as far as possible (section 5).

Outcome of the initial assessment

Claude (Team or Enterprise, or API with a data processing agreement): green subject to conditions. The contract provides a DPA with SCCs, no training on customer data and an IP indemnity outside the liability cap, but processing takes place in the US and the outstanding actions from the internal supplier assessment must first be completed. Processing outside the EEA requires a documented impact assessment in advance (section 7). Even where processing takes place in the EU via Amazon Bedrock or Google Vertex, the risk of access by US authorities (CLOUD Act) remains, because Anthropic and those cloud providers are US companies. Zero data retention is available only on request for the API and Claude Code for Enterprise, not by default for Team or Enterprise (chat). The contract documents were requested from Anthropic in November 2026.

Perplexity: yellow. There is a DPA, but the processing location and security measures have not been established, there is no right to object to third-party model providers and there is no fixed notification period for data breaches. Perplexity is a US company; the impact assessment referred to in section 7 is required before the status can change to green.

Pocket: approved for recording, transcribing and summarising meetings, under the conditions of section 9 and with accepted residual risk (below). The contract documents have been requested from Pocket. Other AI note-takers: red.

Consumer versions (such as Claude Free and Perplexity Free): red. Claude Pro: green under the conditions below. Other paid consumer versions (such as Claude Max and Perplexity Pro): yellow, with the training setting switched off.

Tools currently used by the firm

The firm currently uses Claude Pro and Perplexity Pro. These are consumer versions: there is no data processing agreement, the supplier acts as an independent controller, and the safeguards of the business terms (such as the contractual prohibition on training) do not apply.

Claude Pro may be used for matter work involving client data (status green), under the following conditions:

- the client has given consent through the quotation (section 6) and has not objected;
- the setting that allows conversations to be used to train models is switched off; the Lawyer checks this every six months and whenever Anthropic's terms change, and records the check;
- no more is entered than the question requires, and names and other identifying details of opposing parties and third parties are omitted or replaced where the work permits (section 5);

- not for special categories of personal data or criminal-offence data (Articles 9 and 10 GDPR), identification numbers, or recordings and transcripts of conversations, and not in the cases where section 6 requires specific consent, unless that consent has been given.

Accepted residual risk. On 13 November 2026 the firm decided to admit Claude Pro under these conditions. The residual risk accepted in doing so is: there is no data processing agreement with Anthropic; under its consumer terms Anthropic itself determines the retention and use of data, and may change those terms; data is processed in the United States, with the risk of access by US authorities; and the prohibition on training rests on a setting, not on a contractual commitment. This decision is reviewed at the annual evaluation and whenever the terms change (NOvA recommendation A.2). Switching to Claude Team or the API with a data processing agreement remains the longer-term objective.

Pocket is approved for meetings (section 9). Here too the residual risk has been knowingly accepted: there is as yet no data processing agreement, recordings are processed in the United States, and security has not yet been assessed on the basis of an audit report. The contract documents have been requested; once received, the Lawyer will assess whether additional measures are needed.

Perplexity Pro has status yellow: only without information from which the client or the matter can be identified, and with the training setting switched off. The Lawyer also checks this setting every six months.

AI features in existing software

Office software (such as Google Workspace, Microsoft 365, PDF software and the practice management system) may acquire new AI features through updates that are switched on by default. Such a feature counts as a new tool. Therefore:

- At major updates and at least every six months, together with the check of the training settings (below), the Lawyer checks the settings of the software in use for new or extended AI features.
- A new AI feature is switched off, or at any rate not used with matter data, until it has been assessed in accordance with section 7 and included in Annex A.
- Particular attention is paid to features that send data to a (model) provider other than the software supplier itself, or that process data outside the EEA.

5. Which data in which tool

The table below sets out, for each type of data, the type of tool (section 4) into which it may be entered. Tools with status red are not permitted for any category and are therefore not listed in the table. The restriction applies to everything the firm enters, uploads or has transcribed, including attachments, images and screenshots with their associated metadata.

Type of data	Green	Yellow	Explanation
Public information and general legal questions	Yes	Yes	Without reference to a specific client or matter.
Internal know-how without confidential content (templates, precedents, general work instructions)	Yes	Yes	Check that no client names or matter details are included.
Case descriptions that pass the recognisability test	Yes	Yes, after test	See the recognisability test below. In case of doubt, the stricter category applies.

Type of data	Green	Yellow	Explanation
Client files, documents and correspondence, including data of opposing parties, witnesses and experts	Yes, with consent	No	Consent in accordance with section 6. Do not enter more data than the question requires.
Special categories of personal data and criminal-offence data (Articles 9 and 10 GDPR)	Only after a separate assessment	No	The Lawyer decides in advance and records the decision, assessing at the same time whether a DPIA is required.
Confidential firm data without personal data (strategy, rates, financial data, security information)	Yes	No	Only in a tool that contractually does not train on the data.
Staff and applicant data	Only after a separate assessment	No	Inform the data subject in advance (Articles 13 and 14 GDPR) and use only a green tool.
Recording or transcribing conversations with clients or third parties	Only with an approved tool after notifying all participants	No	See section 9. Pocket is approved for meetings.
Passwords, access keys, citizen service numbers (BSN) and other identification numbers	No	No	Never enter, not even in a green tool.

If the client has consented to the processing of the above data, or has itself entered such data into a particular AI tool, the Lawyer may also use that data or AI tool in the same or an equivalent manner and with the same data.

Recognisability test

A case description passes the test if the client or the matter cannot reasonably be identified from it by an outsider. Removing names is not enough: a combination of facts (sector, place, dates, amounts, particular circumstances, an artist name or the title of a work) can also make a matter recognisable. The test is stricter than anonymisation within the meaning of the GDPR, because professional secrecy also protects information that cannot be traced to a person. In case of doubt, the stricter category applies.

Data minimisation

Share only the information that is reasonably necessary for the specific question. For a question about one clause, that clause with the necessary context is sufficient, not the whole file. This does not rule out large-scale processing (for example a due diligence review), but requires a conscious choice about what is shared. In addition to the content, the supplier also processes metadata and usage data, such as file names, author details, timestamps and IP addresses; these too can reveal a client relationship.

According to the NOvA (C.1), anonymisation or pseudonymisation is not required in all cases if the tool offers adequate safeguards, but it is advisable for high-risk use. This policy remains stricter for yellow tools: only information that passes the recognisability test may be entered into them.

6. Clients: consent and information

1. The Lawyer handling the matter obtains the client's consent in advance for the use of AI in the matter, in line with the NOvA recommendations. This is arranged in the quotation or engagement letter and, for ongoing matters, before the use. An exception applies where the client indicates that it has itself used AI tools in providing information, or where this is evident from the information provided. The client gives this consent by accepting the quotation. That consent covers the tools and uses stated in the quotation (Claude, Perplexity and Pocket). New consent is required for any other tool or any other use involving client data.
2. In addition to the general consent in the quotation or engagement letter, the Lawyer obtains separate and specific consent, and records it in the file, where: (a) the Lawyer does not personally review the final result in full by hand, as with a sample check in bulk work (section 8); (b) professional secrecy cannot be fully safeguarded with the tool; (c) the client has previously objected to the use of AI; (d) special categories of personal data are processed; (e) the Lawyer knows that the client has a confidentiality agreement with a third party that prohibits processing with AI or sharing with third parties; (f) matter data could be used to train a model; or (g) this policy requires it elsewhere.
3. At intake, the Lawyer asks whether the client has its own guidelines on the use of AI by suppliers (such as an AI ban or a consent requirement) or confidentiality arrangements with third parties that restrict AI processing. Any restriction is recorded in the file and passed on to the employee.
4. This policy (or a summary) is available to clients, on request and via the Lawyer's website (www.petitlegal.nl). The firm informs clients which categories of AI tools are used and who processes data in that context, including processing outside the EEA. The quotation refers to the PDF version of this policy on the website. That public version contains the policy without the internal supplier assessment and working documents, states the version number and date, and is republished after every adopted amendment. What the quotation says about the tools used must correspond with the published version.
5. If a client refuses consent, no AI is used with client data in that matter. This is recorded in the file. The client may at any time, without giving reasons, state in writing that no AI may be used; from that moment this also applies to an ongoing matter. AI support is not a condition for accepting an engagement. Working without AI may take more time; that time is charged in accordance with the fee arrangement, but the client's preference prevails.
6. The firm takes into account AI-generated documents that clients provide themselves: these are a starting point, never a final product, and are reviewed.
7. **Information to the client.** The quotation states which tools the firm uses and for what purpose. If this changes materially, for example through a new supplier or processing in another country, the quotation text is amended and existing clients are informed.
8. In the quotation, the Lawyer points out to the client that the client may compromise the confidentiality of the communication by entering correspondence with or documents from the Lawyer into a public AI tool, or by recording a conversation with the Lawyer with a public transcription app, and that the information may thereby lose its status as privileged information.
9. **Billing.** When billing by the hour, the Lawyer charges only the time actually spent, even if AI has performed part of the work in less time. This includes the time spent drafting, reviewing and refining AI-assisted output. Licence costs of AI tools are not passed on to clients. Fixed-price arrangements that factor in AI are permitted.
10. **Disclosure to the court.** The Lawyer discloses to the court the use of AI in preparing a court document where the nature or extent of that use is reasonably relevant to the

assessment of the document, for example where AI-generated or AI-edited image or sound material is submitted as an exhibit. There is no general duty to disclose every use of AI.

11. **Identification.** AI makes it easy to falsify images, sound and documents. The identity of a new client is therefore verified through an independent second channel, and not solely on the basis of video contact or digital documents. Extra vigilance is required for relationships established entirely remotely, unusual urgency, deviating payment instructions or unclear documentation.

7. Supplier assessment

Before a tool becomes green, the Lawyer carries out a supplier assessment and documents the considerations. The minimum requirements are:

- A signed data processing agreement that complies with Article 28(3) GDPR, with a valid transfer mechanism (Articles 44 to 46 GDPR) where data is processed outside the EEA.
- Insight into storage and processing locations and an up-to-date list of sub-processors, including the procedure for changes and objections.
- A contractual commitment that data will not be used for training, and clear retention and deletion periods.
- Demonstrable security (Article 32 GDPR): certifications or audit reports that the firm has actually seen, not merely stated on a website.
- A data breach notification period that allows the firm to meet its own notification obligation (Article 33 GDPR, 72 hours).
- Exit arrangements: export or return of data, deletion after termination and a reasonable notice or transition period.
- An assessment of limitations of liability, unilateral amendment clauses, governing law and dispute resolution.
- A data protection impact assessment (DPIA, Article 35 GDPR) where the processing is likely to result in a high risk. The Lawyer assesses for each tool whether this is the case.
- Inclusion in the record of processing activities (Article 30 GDPR) and in the firm's AI register.
- Insight into the full chain: whether the service relies on an underlying model provider or an intermediate middleware layer, where that party is established and which safeguards apply there.
- Where data is processed or stored outside the EEA, or where the supplier has a non-European parent company: a documented impact assessment in advance, including with regard to access by foreign authorities (such as under the US CLOUD Act). Storage within the EEA does not exclude that risk as long as the supplier is part of a group with a US company. The residual risk and the decision to accept it are recorded and reviewed annually. The supplier may not fall back on a non-approved location without consent, not even in the event of an outage.
- Zero data retention where possible: prompts, context and output are processed only temporarily in working memory, without storage in logs or databases, without caching across sessions and without use for training, fine-tuning or RLHF, including at underlying providers. If chat history is needed, it must be capable of automated and irreversible deletion on request. The label "ZDR" alone is not sufficient; scope and exceptions are checked.
- Commitments as to who at the supplier and its subcontractors has access to stored data and under what circumstances. The supplier's staff do not view unencrypted matter data without the firm's consent. Preferably: encryption at rest and in transit, customer-managed keys and separation of data from that of other customers.

- A written protocol for requests by third parties (litigants, police and judicial authorities, intelligence services, regulators) for the firm's data: no direct disclosure without the Lawyer's consent unless the supplier is under a mandatory legal obligation, immediate prior notification, resistance using all reasonable legal remedies and, if necessary, minimal disclosure as a last resort. Arrangements on the costs of resistance are made in advance. The Lawyer informs the supplier in writing in advance that it will hold privileged information.
- Rights to data: the firm retains all rights to prompts, context, uploaded files and output, including as against the supplier's subcontractors.
- An adequate liability arrangement without unreasonable limitation in the event of a data breach or other serious incident, and confirmation that the supplier is insured for it.
- Measurable service levels (availability, response and resolution times for outages and incidents), maintenance preferably outside office hours and no unilateral change to the terms at short notice without consultation.
- A right to terminate if the agreement is transferred to a party that cannot continue the service at the same quality or that is established outside the EEA.
- A commitment that the supplier keeps its staff AI-literate, provides the documentation the firm needs for its own AI literacy, and does not use prohibited AI practices.

The assessment includes a check of the current terms. Terms and sub-processor lists change regularly; the Lawyer reassesses green tools at least annually and upon every notification of a change.

Four-step approach

A new tool goes through four steps: (1) orientation, (2) demo, (3) pilot with test data (see the test set in section 10), and (4) deployment. The next step is only taken if the previous one gives no reason to abandon the tool. The Lawyer tests the supplier's key promises personally, more extensively where more sensitive data will be processed and more lightly for one-off use or use with public information only. Where reasonably possible, during the pilot the Lawyer checks the separation between users, the actual deletion of data and resilience against prompt injection (for example with a test document containing hidden instructions).

Specialist assistance with the product and contract assessment, or joint procurement with other small firms, is considered for higher-risk tools.

Certificates and audit reports

Certificates (such as ISO/IEC 27001, ISO/IEC 42001 and SOC 2 Type II) are useful, but never decisive. The Lawyer requests the full report and checks (i) which services and systems were tested, (ii) the audit period and validity, and (iii) whether the report also covers the underlying model provider and intermediate layers.

Dependency and data portability

The risk of dependency on a supplier (vendor lock-in) is taken into account, weighing more heavily the larger the share of the firm's core work a tool supports. For each green tool, a simple exit plan is in place before deployment. The firm ensures that the following data can be exported and reused with another supplier: uploaded files in their original format; prompts, prompt history and associated output and metadata; prompt libraries and standing instructions (such as the firm's own skills); and, where applicable, firm-specific adapter files or model weights. Structured data is supplied in json, csv or xml as far as possible. After termination, data, including backups and embeddings, is deleted within a fixed period.

Interim reassessment

In addition to the annual reassessment, a tool is reassessed in the interim upon a change at the supplier (terms, sub-processors, owner or a major model update), upon a change in legislation or in the NOvA recommendations, and after an incident.

8. Use and review of output

1. AI can make mistakes and generate fictitious information, including non-existent case law, statutory provisions and quotations. Every source, quotation and statutory reference in AI output is checked against the primary source before it is used in a document, advice or letter.
2. Case law is checked via the ECLI number and the original reasoning in an official source (not just the AI summary): does the judgment exist, is it relevant and has it not been superseded? Statutory provisions are checked in the official publication, with attention to entry into force, transitional law and article numbering. Factual assertions are checked against the file. Knowingly or negligently adopting errors may breach the duty of truthfulness (Article 21 Dutch Code of Civil Procedure).
3. The intensity of the review is proportionate to the consequences of an error. Departing from full manual review (for example a documented sample check in bulk work) is only permitted after the client has been informed of the risks and has given consent (section 6), and after the output has proved reliable on a representative subset. The departure is recorded in the file.
4. Be alert to tools that tend to confirm the prompt, even when it is wrong. Ask critical and open questions and explicitly ask for counter-arguments. Deliberately use opposing prompts to test a conclusion and consider standing instructions that enforce objectivity and citation of sources. This weighs more heavily where AI is used to support a position already taken.
5. AI is never decisive in assessing a client's position, choosing a litigation strategy or weighing interests.
6. The review required depends on the type of task. From light to heavy: (a) a procedural sub-task, such as an exhibit index or formatting; (b) improving one's own work, such as spelling and readability; (c) identifying patterns or anomalies that the Lawyer then assesses one by one; (d) a preparatory substantive analysis that the Lawyer then fully reassesses. The closer the AI contribution comes to the substantive assessment, and certainly where it is shared with litigants or the court, the more intensive the review. The Lawyer does not allow the client to pressure him into relying blindly on an AI analysis for reasons of time or cost.
7. The Lawyer critically assesses whether the output takes account of the client's individual circumstances, particularly in advice on prospects of success, and is alert to strikingly one-sided answers as a possible indication of bias. In cross-border matters, it is checked whether the output is based on the applicable law: a question asked in Dutch is often answered under Dutch law, even where another law applies. Extra vigilance is required after major model updates.
8. AI is not used to make court documents unnecessarily voluminous or to overwhelm litigants or the court with information that is not relevant to the dispute.
9. Documents intended for the court, an opposing party or a client are reviewed and signed or released by the responsible lawyer personally.
10. **Agentic AI.** AI agents (tools that independently perform multiple steps, edit files, call systems or send messages) are only used if it has been mapped in advance which tools, integrations and services the agent can call and which data goes to which party in doing so. The agent is given no more permissions than the task requires. Every step with possible

legal consequences, such as sending correspondence, filing documents or amending client documents, requires prior confirmation by the Lawyer. The agent's actions are logged where the tool allows, and the result is reviewed before it goes to the client or the court. The Lawyer is alert to prompt injection via documents, emails or websites that the agent reads, and to an overly broad interpretation of the instruction.

11. **Shared memory and knowledge bases.** Memory features, projects, knowledge bases or standing instructions containing information about multiple matters may cause information from one client to carry over into work for another. Matter data is therefore kept separate per matter (for example one project per matter); no matter data is placed in a firm-wide memory or firm-wide knowledge base. The conflict check also verifies whether AI memory or knowledge bases contain information about a party now opposing a (former) client.
12. **Internal transparency.** For every draft, the employee tells the Lawyer whether and how AI was used, so that the review can be tailored accordingly. Using AI is not taboo; experiences and mistakes are shared so that lessons can be learned.

9. AI note-takers and recordings

1. Recording or transcribing conversations with clients, opposing parties or other third parties with an AI tool is only permitted with a green tool approved for that purpose, or with the client's consent. Pocket is approved for this purpose.
2. All other participants are informed in advance that the conversation is being recorded or transcribed, for what purpose and by whom the data is processed.
3. Recordings and transcripts are linked to the matter file, not to a personal account or private device, and are deleted in accordance with the retention period of the file.
4. For conversations in which legal professional privilege or professional secrecy is particularly sensitive, such as litigation preparation and negotiations on confidentiality, the Lawyer personally decides whether a recording is appropriate. The employee never records such conversations independently.
5. The consent the client gives through the quotation covers meetings with the client. For meetings with opposing parties or other third parties, they are informed in advance (point 2), and no recording is made if any of them objects.
6. Transcripts and summaries from Pocket are a working aid and not an approved record. If the Lawyer shares a summary with the client, he reviews it first; the client may request corrections. Recordings and transcripts are not kept in Pocket for longer than necessary for the file.

10. Knowledge and training

The firm ensures an appropriate level of AI literacy among everyone working with AI, as required by Article 4 of the AI Act. The Lawyer and the employee attend training at least annually on the workings and limitations of generative AI, hallucinations and error detection, confidentiality and the GDPR, prompting, cybersecurity and this policy.

In addition to theory, the Lawyer and the employee gain practical experience with the approved tools, within the framework of this policy. Training attended is recorded as part of the continuing professional education record.

Core skills. AI does not replace learning core skills oneself. An employee or trainee doing legal work also carries out legal research, analysis and writing personally and by hand, so that work can continue if AI is unavailable and errors in AI output are recognised.

Test set. The Lawyer compiles a small, anonymised test set of the firm's own manually drafted and reviewed documents that are representative of the practice (for example a contract review, a letter of demand and an opinion on copyright or neighbouring rights). New tools and major updates of existing tools are tested against it. The test set is periodically updated with new case law and legislative changes.

11. Oversight, register and incidents

1. The Lawyer maintains a register of approved tools, their status, the date of assessment and reassessment, the contract, the supplier contact point and, for each tool, the permitted, enhanced-safeguard and excluded uses with the associated measures (Annex A).
2. At least once a year, the Lawyer checks compliance with the policy on the basis of a sample of matters in which AI was used, and briefly records the findings.
3. Incidents (inadvertently entering client data into a non-approved tool, a suspected data breach, incorrect AI output in an outgoing document) are reported to the Lawyer immediately. In the event of a data breach, the firm assesses whether notification to the Dutch Data Protection Authority and to data subjects is required, and whether the client must be informed. The notification period under Article 33 GDPR is 72 hours after becoming aware. An incident also includes a request by a third party to a supplier for the firm's data, or an AI feature activated unnoticed through an update. In the event of a request by a third party, the Lawyer immediately assesses whether legal professional privilege must be invoked.
4. In the event of professional conduct questions or a possible breach of professional secrecy, the Dean of the local Bar is consulted where necessary.
5. **Risk assessment.** The Lawyer assesses the risk of AI use on the basis of three questions: (1) what is the risk of the tool itself (sections 4 and 7); (2) which data is processed (section 5); and (3) what type of task is performed and what are the consequences of an error (section 8). The higher the risk on any of the three, the stronger the safeguards. This approach is aligned with the structure of the NIST AI Risk Management Framework, without seeking formal certification.

12. Compliance and evaluation

Non-compliance with this policy by staff may lead to employment law measures. The policy is evaluated at least annually, taking into account changes in the AI Act, GDPR enforcement, the NOvA recommendations and suppliers' terms. The policy is also revised in the interim upon a change at a supplier, a change in legislation or the NOvA recommendations, or after an incident. At the evaluation, the firm also considers whether AI can contribute to more affordable legal assistance, and whether joint procurement of AI tools with other small firms is worthwhile.

Questions and complaints. Clients may always put questions about the use of AI to the Lawyer. A complaint about the use of AI is handled in accordance with the firm's complaints procedure.

Annex A – Register of tools

Classification based on publicly available terms as at 14 October 2026. Based on the completed internal supplier assessment. Supplemented on 13 November 2026 with uses per tool (NOvA recommendation A.1).

Tool	Status	Permitted for	Conditions and points of attention
Claude (Anthropic)	Green (conditional)	Business plans (Team/Enterprise) or API with DPA: matter work with consent.	Processing in the US, SCCs required. EU processing only via Bedrock or Vertex. Consumer versions prohibited. Assess standard cap and right to amend terms. Contract documents requested (November 2026).
Claude Pro (in use)	Green (conditional)	Matter work with consent through the quotation, under the conditions of section 4.	Consumer version without a data processing agreement; residual risk accepted (section 4). Training setting off, checked every six months. Not for special categories of personal data, identification numbers or recordings.
Perplexity Pro (in use)	Yellow	General research and case descriptions that pass the recognisability test (section 5).	Consumer version without a data processing agreement. Training setting off, checked every six months. No client data.
Perplexity	Yellow	General research with public information or case descriptions that pass the recognisability test (section 5).	Enterprise/API DPA exists, but no right to object to third-party model providers, no fixed data breach notification period and no intended processing of special categories of data.
Pocket (and similar AI recorders)	Green for meetings	Meetings with clients (consent through the quotation) and, after prior notification, with other participants (section 9).	No DPA or sub-processor list publicly available, processing in the US, consumer-oriented terms. Documents requested from Pocket; reconsider after receipt and assessment. Approved for meetings with accepted residual risk (section 4).
AI features built into office software	Red pending assessment	No use with matter data until the feature has been assessed (section 4).	For example AI features in Google Workspace or Microsoft 365. Assess per feature. Example: Claude in Microsoft 365 Copilot currently processes data outside the EU Data Boundary; Microsoft has announced that this will change at the end of 2026.

Uses per tool

Tool	Permitted	Enhanced safeguards	Excluded
Claude (Team, Enterprise or API)	Research and analysis; drafts of letters, opinions and court documents; contract review; summarising matter documents (with consent, section 6).	Special categories of personal data (separate assessment); bulk work with sample check (client consent); agents with access to email, calendar or documents (confirmation of each action with legal effect, section 8); memory and projects (one project per matter).	Consumer versions; passwords and identification numbers; independently sending correspondence or filing documents without prior confirmation.
Claude Pro (in use)	Research and analysis; drafts of letters, opinions and court documents; contract review; summarising and translating matter documents (with consent through the quotation).	Data of opposing parties and third parties: omit or replace where possible; cases for which section 6 requires specific consent; memory and projects (one project per matter).	Special categories of personal data and criminal-offence data; identification numbers; recordings and transcripts; matters of clients who have objected.
Perplexity (Pro)	General research with public information and case descriptions that pass the recognisability test.	Legal sources: always verify against the primary source (ECLI, official publication).	Client data, matter documents and confidential firm data.
Pocket (approved for meetings)	Recording, transcribing and summarising meetings with clients (consent through the quotation) and, after prior notification, with other participants (section 9).	Meetings in which legal professional privilege is particularly sensitive: only after the Lawyer's own decision; review summaries before sharing them.	Recording without a participant's knowledge or after an objection; independent recording of sensitive meetings by the employee.
AI built into office software	None, pending assessment.	—	Use with matter data pending assessment.