

AI Product Evaluation Checklist

For Law Firms & Judicial Offices

Use this checklist during vendor review, pilot approval, and contract negotiation for any AI-enabled product. As artificial intelligence tools become increasingly prevalent in legal and judicial settings, rigorous evaluation is essential to protect client confidentiality, uphold professional obligations, and ensure responsible deployment.

This checklist covers twelve critical domains — from data inputs and security controls to exit rights and institutional fit — providing a structured framework for evaluating any AI vendor before, during, and after procurement.

12 Evaluation Domains

Comprehensive coverage from data lineage to contract controls

Vendor Review

Structured questions for every stage of procurement

Legal & Judicial Focus

Tailored for law firms and judicial offices specifically

Data Inputs, Lineage & Storage

DOMAINS 1 & 2

Before deploying any AI tool, your organization must have complete clarity on what data enters the system, how it moves, and where it ultimately resides. You should be able to identify exactly what data will enter the tool and confirm that the vendor documents its origin, movement through the system, and final destination. Critically, you must understand how the product helps you control who can access the data at each stage, and whether the tool's access can be limited to only the minimum data necessary. This is especially critical if the data enters via API or integration with other tools in your tech portfolio.

Special attention is required when the tool may process client, case, personnel, sealed, confidential, or otherwise sensitive information. These categories carry heightened professional and legal obligations that must be reflected in vendor agreements.

Storage & Jurisdiction

Confirm where data is stored and in what jurisdictions.
Understand how long data is retained — including prompts, outputs, logs, and metadata.

Model Training Risk

Determine whether the vendor uses your data to train, fine-tune, improve, or benchmark its models. This position must be stated clearly in the contract — not just in marketing materials — and retention settings should be adjustable or disableable where appropriate.

Subprocessors, Model Providers & Security

DOMAINS 3 & 4

Modern AI products rarely operate in isolation. Vendors typically rely on a network of subprocessors, model providers, cloud hosts, and infrastructure partners — each of whom may receive or process your data. You must confirm that the vendor has identified all key third parties and can specify which ones may receive your data. Equally important, the vendor should commit to notifying you before adding or changing material subprocessors, and all downstream providers must be subject to equivalent confidentiality, security, and data-use restrictions.



Technical Safeguards

Confirm what technical and organizational security measures protect data in the system, including encryption in transit and at rest for uploaded documents, prompts, and outputs.



Access Controls

The product should support role-based access controls, SSO, MFA, and audit logging. Access should be limitable by matter, team, office, or user role.



Incident Response

The vendor must have a documented incident response process for security events involving your data, with clear notification timelines and remediation obligations.

Testing, Validation & Known Limitations

DOMAIN 5

AI tools deployed in legal or judicial settings must be rigorously tested before adoption. Vendors should be able to demonstrate that the product has been evaluated for accuracy, reliability, bias, and security. Ask for documentation of red-teaming exercises, evaluation methodologies, and known failure modes. Ideally, the product will have been tested on legal or adjudicative workflows similar to your own.

Transparency about limitations is non-negotiable. Vendors must disclose hallucination risk, confidence limits, and the specific situations where outputs may be unreliable. Users should be clearly warned not to treat AI-generated outputs as authoritative without independent review — a safeguard that is especially critical in high-stakes legal and judicial contexts.

Red-Teaming Documentation

Has the vendor provided documentation of adversarial testing, evaluation methods, and known failure modes?

Legal Workflow Testing

Has the product been tested on legal or adjudicative workflows similar to yours?

Hallucination Disclosure

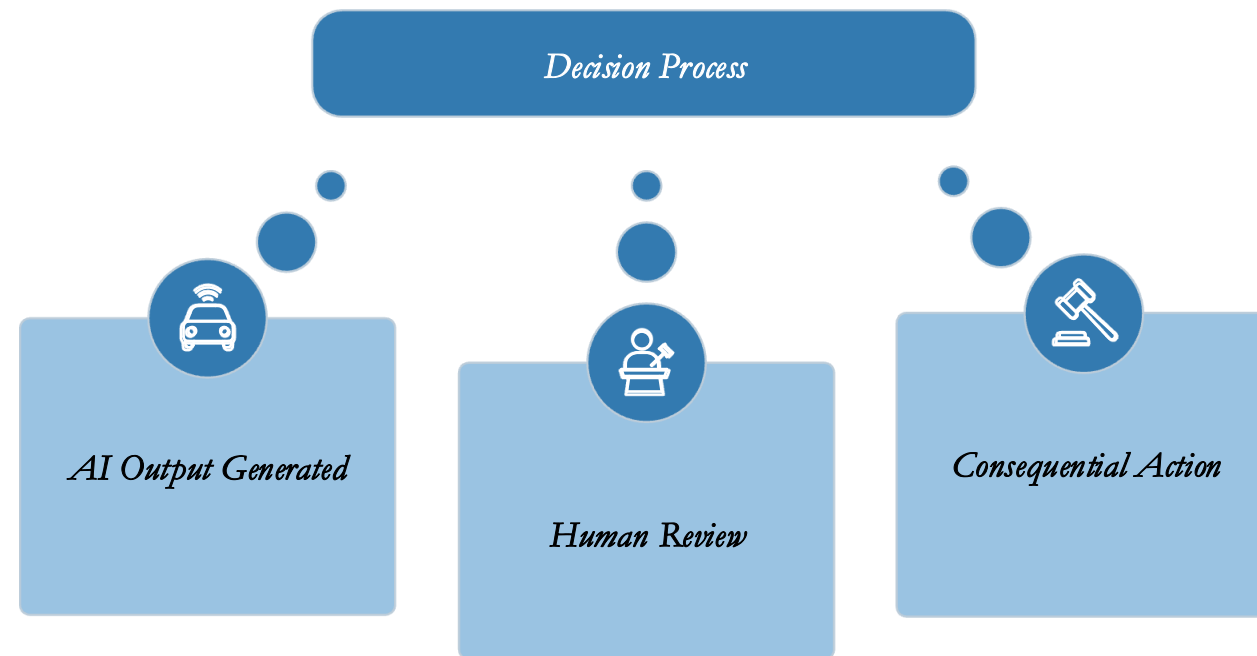
Does the vendor disclose hallucination risk, confidence limits, and situations where outputs are unreliable?

Human Review, Escalation & Professional Judgment

DOMAIN 6

AI tools in legal and judicial settings must be designed to support — not replace — professional judgment. Any workflow involving consequential decisions must require appropriate human review before action is taken. Clear escalation paths must exist for situations where outputs appear wrong, incomplete, or inappropriate, and users must be able to easily verify the source basis for any AI-generated output.

The tool's design should reinforce the primacy of human expertise. For judicial and adjudicative settings, this obligation is especially acute: there must be explicit safeguards against overreliance on AI-generated analysis, ensuring that independent judicial reasoning remains intact and uncompromised.



This oversight framework ensures that AI tools serve as analytical aids rather than decision-makers, preserving professional accountability at every stage of the workflow.

Confidentiality, Privilege & Output Handling

DOMAINS 7 & 8

Legal confidentiality obligations and privilege protections are foundational requirements — not optional features. The vendor must be able to demonstrate that the product supports these obligations. The tool should be evaluated for risks involving client confidences, work product, sealed filings, chambers material, personnel records, and deliberative material. The contract must explicitly prohibit vendor access or secondary use of this data except as strictly necessary to provide the contracted service, with additional controls for especially sensitive data categories.

Output Traceability

Outputs must be traceable, labelable, and loggable. Citations, source links, or supporting materials should be provided where needed, and users must be able to distinguish generated content from original source material.

Correction & Guardrails

The product must make it easy to correct, reject, or delete flawed outputs. Guardrails should prevent unauthorized filing, dissemination, or reliance on unverified content — critical protections in any legal or judicial context.

Audit, Records & Oversight

DOMAIN 9

Accountability in AI-assisted legal work requires robust audit capabilities. Your organization must be able to review usage logs, administrator logs, and access history at any time. The vendor should be able to support internal review, compliance functions, and supervisory oversight — and logs must be retained long enough to support investigations and accountability reviews.

Comprehensive documentation is essential: you must be able to record who used the system, on what data, and for what purpose. This level of auditability is not merely a best practice — it is a professional and institutional obligation in legal and judicial environments where decisions carry significant consequences for individuals and institutions alike.



Usage Logs

Review usage, administrator, and access history logs on demand



Compliance Support

Vendor supports internal review, compliance, and supervisory oversight



Log Retention

Logs retained long enough to support investigations and accountability



User Documentation

Document who used the system, on what data, and for what purpose

Exit Rights, Data Retrieval & Procurement Controls

DOMAINS 10 & 11

Vendor relationships end — and your organization must be prepared. You must be able to retrieve your data, prompts, outputs, and usage history in a usable format. The contract must clearly specify what happens to your data at contract end or after termination, how quickly the vendor will delete or return data, and whether transition assistance is provided. Avoiding technical or contractual lock-in is a critical procurement objective.

Contract Enforceability

Vendor marketing claims must be matched by enforceable contract terms. The contract must clearly address data use, confidentiality, security, retention, audit rights, and termination procedures — not merely reference them in general terms.

Liability & New Use Cases

Liability, indemnity, and limitation-of-liability provisions must be appropriate for the sensitivity of the use case. A process for approving new use cases after product adoption should also be established to prevent scope creep without proper review.

Vendor marketing materials are not contractual commitments. Every data protection, confidentiality, and security assurance must appear as an enforceable term in the signed agreement.

Fit for Law Firms & Judicial Offices

DOMAIN 12


Not every AI tool is appropriate for every legal or judicial setting. The proposed use must be evaluated against the professional context and risk level before deployment. A tool that is appropriate for administrative tasks may be wholly inappropriate for adjudicative analysis or privileged client communications.



Law Firms

Does the tool align with client obligations, privilege protection, and matter confidentiality? Uses that risk exposing client confidences or work product to vendor systems require heightened scrutiny and explicit contractual protections.

In both settings, organizations should explicitly identify uses that should be **prohibited**, **restricted**, or subject to **heightened review** — and document those determinations before deployment begins.



Judicial Offices

Does the tool avoid undermining impartiality, confidentiality, record integrity, and independent judgment? AI tools in adjudicative settings must never substitute for judicial reasoning or compromise the integrity of deliberative processes.

Summary: A Framework for Responsible AI Adoption

This checklist provides a comprehensive framework for evaluating AI tools in legal and judicial settings across twelve critical domains. From the moment data enters a system to the day a vendor contract ends, every stage of the AI product lifecycle carries professional, ethical, and legal obligations that must be proactively managed.

01

Data & Infrastructure

Evaluate data inputs, lineage, storage, retention, model training practices, and subprocessor chains (Domains 1–3)

03

Professional Safeguards

Confirm human review workflows, confidentiality protections, output handling, and privilege controls (Domains 6–8)

02

Security & Validation

Assess security controls, access management, testing documentation, and known limitations (Domains 4–5)

04

Governance & Contract

Verify audit capabilities, exit rights, procurement controls, and institutional fit (Domains 9–12)

This checklist should be revisited at each stage of the vendor relationship, not only at initial procurement. New use cases, product updates, and changes to subprocessors all warrant fresh evaluation against these criteria.