



AI for legal, risk, and compliance

Using AI to accelerate review and discovery while retaining accountable professional judgment.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 12 / EMBER & SKY

EXECUTIVE WORKING EDITION - 08 OCT 2026

12

AT A GLANCE

Using AI to accelerate review and discovery while retaining accountable professional judgment.



Legal, risk and compliance teams handle large volumes of text under tight deadlines, which makes them natural candidates for AI support and also places the strictest demands on accuracy and confidentiality. AI can help find clauses, summarize long documents, compare versions, draft first versions, triage intake and monitor changes in obligations. It cannot take responsibility for a legal conclusion, and it can produce fluent errors, including citations to sources that do not exist. This playbook sets out where AI supports these teams, how to protect privilege and confidential data, how to design review so counsel stays accountable, how to measure quality, and how to implement in stages. It supports professionals and does not replace them or give legal advice.

WHO IT IS FOR

General counsel, heads of risk and compliance, legal operations and privacy leads, plus the technology and security partners who support them.

TIME TO USE

70 minutes to read; two workshops to apply (task inventory and review design)



What you will have at the end

- A ranked list of legal, risk and compliance tasks by value, risk and suitability for AI support
- A review standard that keeps a named professional accountable for every output used
- Confidentiality and privilege controls agreed with security and outside counsel where relevant
- A test set and scoring method for the first use case, including checks for fabricated sources
- A 90-day plan from a low-risk pilot to a decision to scale



Contents

01	What AI means for legal, risk, and compliance work	03
02	Where AI fits and where it does not	04
03	How it works, in plain language	06
04	Operating model, roles and accountability	08
05	Risks, controls and human review	09
06	Measuring quality and value	11
07	Implementation guide: start where checking is easy	13
08	Common pitfalls and an illustrative scenario	15
09	Questions to ask vendors, and build or buy	16
10	A first quarter that produces evidence	18
11	Terms used in this publication	19



EXECUTIVE BRIEF

What AI means for legal, risk, and compliance work

01



AI is most useful here as a fast, tireless first reader and drafter whose work a qualified person always checks.

Much of the work in these functions is reading and comparing text: contracts, policies, regulations, regulator letters, case files, evidence, audit findings and incident reports. Language models are good at summarizing, extracting, classifying, comparing and drafting. They are not good at knowing what is true, what is current, or what a court or regulator would decide. They generate likely text, and likely text can be wrong in a way that reads as authoritative.

That combination shapes the right posture. Use AI to reduce the time to a first view, to widen the first pass of a review so fewer items are missed, and to remove drafting drudgery. Keep judgment, interpretation, advice to the business and any statement to a regulator, court or counterparty with accountable professionals. Professional duties, including competence and confidentiality, continue to apply when tools are used, and rules differ by jurisdiction and professional body. Check current guidance from your regulator or professional body.

- AI supports the team; it does not become the author of legal positions or the owner of risk decisions
- Every output used in work product is reviewed by a qualified person who is accountable for it
- Confidentiality and privilege are design constraints from the first day, not later add-ons
- Fabricated citations and misquoted text are known failure modes and need a specific check
- The best early use cases are bounded, repetitive and easy to verify against the source document

PLAIN TEST

If a competent junior could do the task from the documents in front of them, and a senior could check it in a fraction of the time, AI assistance probably fits. If the task is the judgment itself, it does not.



WORKING QUESTION

Which recurring task in your team is mostly reading and comparing, and where would a faster first pass change how you spend senior time?

WHERE IT FITS

Where AI fits and where it does not

02

Match the task to the risk. Locating and organizing information is a stronger fit than concluding, advising or representing.

A practical way to sort tasks is by what happens if the output is wrong and how easily it can be checked. Extracting the renewal date from a contract is easy to verify against the page. Summarizing a regulator's consultation paper for a briefing is checkable but needs care. Advising on whether a practice breaches a rule depends on facts, jurisdiction and interpretation, and belongs to a professional. Drafting for external use, such as a filing, needs the highest scrutiny because errors become public and consequential.

TASK	FIT	REASON AND CONDITION
Contract clause location and extraction	Strong	Verifiable against source; use clause references and page pointers
First-pass contract review against a playbook	Strong with review	Flags deviations; lawyer decides acceptability
Document review and eDiscovery triage	Strong with sampling	Prioritization and clustering; validated by sampling and established protocols
Summaries of policies, regulations and long files	Moderate	Useful orientation; check against the source and its date

TASK	FIT	REASON AND CONDITION
Regulatory change monitoring	Moderate	Surfaces items for human review; may miss or mislabel changes
Drafting routine agreements and notices from approved templates	Moderate	Template and fallback language controlled by the team
Policy question answering for employees	Moderate	Grounded in approved policies; clear route to a person
Legal research and case law	Weak without verification	Fabricated or misdescribed authorities are a known risk
Legal opinions, regulatory submissions, litigation positions	Not suitable for unsupervised use	Accountable professional judgment required

OPTION	BEST WHEN	WATCH FOR
General assistant on approved terms	Drafting, restructuring and summarizing non-privileged text	Data handling; no source grounding
Assistant grounded in your clause library and playbooks	Consistent review and drafting against your positions	Library quality and currency
Specialist legal tools	Research, review platforms and contract lifecycle features	Vendor claims; integration and data terms
Rules and workflow without AI	Intake forms, routing and deadlines	Often enough; cheaper and more predictable

COMMON MISTAKE

Starting with legal research or advice because it looks impressive. Begin with tasks whose outputs you can check against a document in minutes. Build trust and measurement habits there first.



 WORKING QUESTION

For your three busiest tasks, can a reviewer check the output against a source in a few minutes, or does checking take as long as doing the work?

 HOW IT WORKS

How it works, in plain language

03

Models predict useful text from patterns. They become dependable in legal settings only when they are given the right documents, narrow instructions and a way to check what they produce.



A language model has no database of law and no memory of your matters. It has learned patterns in text and produces a likely continuation of what it is given. When it is asked a question from memory, it may blend real and invented material. When it is shown the relevant document and asked to work only from it, quality rises sharply and checking becomes simple. That is why the reliable patterns in this field are all variations on the same idea: supply the source, constrain the task, require pointers and verify.

01 Supply the source

Give the tool the contract, policy, regulation text or evidence set it should work from, in a protected environment approved for that data.

02 Constrain the task

State the question, the format and the rules, for example: answer only from the document, quote the clause, and say if it is not present.



03 Require pointers

Ask for section numbers, page references or quoted text so a reviewer can find the basis in seconds.

04 Verify

A person confirms quotes against the original, checks any cited authority exists and says what it says, and notes the date of the source.

05 Record

Keep the version used, the instructions and the reviewer decision so the work can be explained later.

PATTERN	WHAT IT DOES	TYPICAL FAILURE
Extraction	Pulls terms, dates, parties, obligations into a table	Missed clause in a long or oddly formatted document
Comparison	Shows differences between versions or against a standard	Misses a subtle change in meaning where wording looks similar
Classification and triage	Sorts documents or requests by type and urgency	Rare but critical items placed low
Drafting from templates	Assembles first drafts from approved language	Unapproved wording introduced; wrong party or defined term
Question answering over documents	Answers from supplied policies or contracts	Answer from outdated or conflicting version
Open question from model memory	Produces general commentary	Plausible but invented or outdated statements

KEY DISTINCTION

There is a difference between asking a model what the law says and asking it what a supplied document says. The second can be verified in minutes. The first needs independent research by a professional.

WORKING QUESTION

For your first use case, what exact source documents will the tool be shown, and how will a reviewer confirm each output against them?

OPERATING MODEL

Operating model, roles and accountability

04

Define who may use which tools on which matters, who reviews outputs, and who owns the library of approved language that drives consistency.

Legal and compliance work is already structured by matters, clients, privilege and authority levels. An AI operating model should follow those lines. Decide which categories of work may use which tools, which require outside counsel approval or client consent, and which are excluded. Make the decision visible in a short, plain policy that lawyers and staff can apply without guessing. Where the team serves internal clients, tell them when AI has contributed to work product, according to your own policy and professional norms.

ROLE	OWNS	TYPICAL TASKS
General counsel or head of function	Policy on use and risk appetite	Approves use cases and exclusions; reports to the board or committee
Legal operations lead	Process, tools and measurement	Runs pilots, vendor reviews, metrics and training
Reviewing lawyer or compliance officer	Accuracy and judgment on each output	Verifies quotes, authority and conclusions; signs off
Knowledge lawyer	Clause library, playbooks and templates	Maintains approved positions, fallbacks and precedent
Privacy and security lead	Data handling and access	Assesses vendors; sets rules for personal and privileged data
Business or matter owner	Use of outputs in the business	Provides context; escalates unclear items
Outside counsel	Specialist advice and jurisdictional questions	Advises on rules, privilege and use in disputes

What a short use policy should say

- Which tools are approved, for which categories of data, and which are prohibited

- What must never be entered, such as privileged material in unapproved tools or special-category personal data without approval
- The review standard: who checks, what they check and how sign-off is recorded
- How to treat citations and authorities: every one verified before use
- When to disclose AI assistance to a client, counterparty or regulator, according to applicable rules and engagement terms
- How to report an error or incident, and who decides next steps



Invest in the clause and playbook library. AI that reviews against your fallback positions is far more useful than AI that reviews against generic ideas of what is standard. Knowledge lawyers who keep the library current become a quiet center of value, because every review, draft and answer draws on their work.

WORKING QUESTION

If an AI-assisted draft reached a client with a fabricated citation in it, who in your structure would have been accountable for catching it, and would they have known it was their job?



RISKS AND CONTROLS

Risks, controls and human review

05



Key risks are fabricated or outdated content, breach of confidentiality or privilege, bias and unfairness in decisions, and over-trust. Controls are specific and mostly procedural.

Frameworks such as the NIST AI Risk Management Framework and ISO/IEC 42001 provide structure for governing AI systems. Data protection law governs personal data in documents, and the EU AI Act and sector rules may apply depending on the use, with obligations that vary by role and risk level. Professional conduct rules, court practice directions and client engagement terms may also address the use of AI. This playbook does not give legal advice. Confirm current requirements with qualified counsel in each relevant jurisdiction, and convert



them into written controls.

RISK	EXAMPLE	CONTROL
Fabricated or misdescribed authority	A cited case does not exist or does not say what is claimed	Verify every authority against a primary source before use; keep a record
Outdated or wrong-jurisdiction content	A summary reflects a superseded rule	Source date shown; jurisdiction stated in the task; expert review
Loss of confidentiality or privilege	Client material entered into an unapproved tool	Approved tools only; data terms reviewed; no training on inputs; access controls
Personal data misuse	Employee investigation data processed without a basis	Privacy assessment; minimization; purpose limits; retention rules
Biased or unfair outcomes	Triage systematically ranks some groups lower	Test across categories; human decision; periodic audits
Missed items in review	Key document marked low priority	Sampling of low-ranked items; recall targets; defensible protocol
Over-trust	Reviewers skim and approve	Required quote checks; seeded errors in audits; time for review
Prompt injection in documents	Text in a counterparty file instructs the tool	Treat documents as data; limit tool actions; review ingested files

Controls to confirm before use on live work

- Tool and data terms reviewed by privacy, security and, where relevant, outside counsel
- Matter categories permitted, restricted and prohibited are written down
- Reviewers are trained to check quotes, authorities, dates and jurisdiction
- Outputs carry pointers to source text, and reviewers use them
- Sign-off by a named qualified person is recorded for work product
- Logs and retention follow policy and any legal hold requirements
- Clients and counterparties are informed where rules or contracts require it
- An incident route exists for errors, leaks and misuse

EXAMPLE (ILLUSTRATIVE)

A compliance team uses an assistant to compare a new internal policy with an obligations register. The tool lists possible gaps with quoted text from both. A compliance officer confirms each gap, decides which matter, and signs the review note. The tool never closes a finding on its own.



WORKING QUESTION

Where in your current workflow would a plausible but wrong output slip through, and what single check would stop it?



MEASURING VALUE

Measuring quality and value

06

Judge AI assistance by accuracy on your own documents, reviewer time, completeness and the quality of decisions, not by speed alone.



Set a baseline before any pilot. For the chosen task, record how long a competent reviewer takes, what they miss and how consistent two reviewers are with each other. That last figure is often an uncomfortable surprise and a fair benchmark: AI does not need to be perfect, but it needs to be measured against how well the current process actually performs. Then build a test set from past, non-sensitive or properly authorized matters, with answers agreed by senior reviewers.

METRIC	QUESTION IT ANSWERS	HOW TO MEASURE
Extraction accuracy	Are the captured terms right?	Compare with a verified set, per field and clause type
Recall of key issues	Did it find what matters?	Planted and known issues in test contracts or files



METRIC	QUESTION IT ANSWERS	HOW TO MEASURE
Precision of flags	How many flags are noise?	Share of flags a reviewer accepts
Quote and citation accuracy	Do the pointers lead to the stated text?	Check every quote and authority in a sample
Reviewer time	Is review faster with the tool?	Timed comparison, including checking time
Consistency	Do outputs vary between runs or users?	Repeat tests; compare reviewers
Turnaround	Are requests resolved sooner?	Cycle time from intake to response
Escapes	What got through that should not have?	Errors found later by senior review or audit

S.

01 Choose one task

For example, extracting key commercial terms from supplier agreements.

02 Build the test set

Use a spread of documents including poor scans, unusual drafting and missing clauses.

03 Agree the answers

Two senior reviewers mark the correct results and resolve disagreements.

04 Run and score

Score accuracy, recall and pointer quality. Study every miss for its cause.

05 Set a pass mark

Set it by task risk before the pilot, and retest after any change of tool, model or instruction.

ILLUSTRATIVE FIGURES ONLY

If a first-pass review of a standard agreement takes 90 minutes and the assisted process takes 50 minutes including verification, the time saved is 40 minutes per agreement. Count the checking time. If checking is skipped to hit the number, the saving is not real.

WORKING QUESTION

What miss rate on key issues would your team accept for this task, and who decided that figure?

**IMPLEMENTATION GUIDE**

Implementation guide: start where checking is easy

07



Begin with a bounded task, approved tools and strong review. Expand only when measurement and confidence justify it.

01 **Inventory tasks**

List recurring tasks with volume, time, risk and how easily outputs can be verified. Select one or two for the pilot.

02 **Set guardrails**

Agree the use policy, approved tools, data rules and review standard with security, privacy and the function head.



03 **Prepare the library**

Collect approved templates, clause positions, policies and the obligations register that the tool will draw on.

04 **Build the test set**

Assemble documents and agreed answers, and decide pass marks.

05 **Configure and test**

Set instructions, require pointers, run the tests, study misses and adjust.

06 **Pilot with real work**

Run with a small group of reviewers, record time, corrections and issues, and meet weekly.



07 Decide and train

Decide to scale, change or stop. Train the wider team on the review standard, not only on the tool.



Readiness for a first pilot

- A sponsor at head-of-function level and a named pilot lead
- A task whose output can be checked against a document
- Approved tool with reviewed data terms and access controls
- Clause library or playbook that is current and owned
- Reviewers with time set aside to check and give feedback
- Agreed rules on what data may and may not be used
- A baseline of current time and quality for the task



Build a habit of recording corrections. Each time a reviewer changes an output, note why. Patterns show where instructions need tightening, where the library has gaps and where the tool simply should not be trusted. Over a few months this record becomes the most practical guide to what AI can and cannot do in your team.

Training matters as much as tooling. Teach lawyers and analysts how to frame a task, how to supply the source, how to ask for pointers and how to spot the signs of invented content: unusually tidy citations, absent page references and confident answers to questions the document does not address. Make the verification step a visible part of the workflow rather than a personal habit.



WORKING QUESTION

Which task could you pilot in the next month where the output is a table of facts that a reviewer can tick off against the source?



PITFALLS AND SCENARIO

Common pitfalls and an illustrative scenario

08



The failures that matter in this field are quiet: unchecked output, unclear accountability and data in the wrong place.

- Using a public consumer tool for client or privileged material because it was convenient
- Accepting citations without opening the source
- Letting speed targets pressure reviewers into skipping checks
- Training the team on prompts but not on the review standard
- Applying one tool and one policy to every matter type regardless of sensitivity
- Treating regulatory monitoring output as complete rather than as a prompt for review
- Failing to update the clause library, so the tool reviews against outdated positions



ILLUSTRATIVE SCENARIO (HYPOTHETICAL)

A corporate legal team pilots contract review on supplier agreements. The tool flags a liability cap as acceptable because it matches a standard phrase, but a separate clause carves out most losses. A reviewer reading only the flag would have missed it. The team adds a rule requiring the tool to list all clauses affecting liability and to quote them, and reviewers now read the quoted set. Recall on this issue improves in the next test round.



PITFALL	EARLY WARNING SIGN	RESPONSE
Unchecked authority	Citations without links or page numbers	Mandatory verification step; reject outputs without pointers
Shadow use	Staff use personal accounts for client work	Approved alternative; clear policy; monitoring and training
Accountability gap	Nobody can say who signed off	Named reviewer and recorded decision for each use



PITFALL	EARLY WARNING SIGN	RESPONSE
Stale library	Tool recommends superseded positions	Owner and review dates for each template and playbook
Over-scoping	Pilot spreads across many matter types	Limit to one task until results are stable

WORKING QUESTION

Where could someone on your team use AI today without anyone in the function knowing, and what would you want them to do instead?

BUYING QUESTIONS

Questions to ask vendors, and build or buy

09

Ask for proof on your documents, clear data terms and the ability to see where every output came from.

VENDOR QUESTION	WHAT TO LOOK FOR
Is our data used to train or improve models, and can this be disabled in writing?	Written terms, enterprise controls and deletion rights
Where is data stored and processed, and who can access it?	Region options, subprocessor list, access logging
Can the tool show the source of every output?	Quotes and page pointers that open the original passage for every output?
How are privilege and confidentiality protected?	Tenant isolation, role-based access, matter-level permissions
Can we test on our own documents with agreed scoring?	A structured trial and your own test set, not a scripted demo

VENDOR QUESTION	WHAT TO LOOK FOR
How are model or feature changes communicated?	Release notes, notice periods, ability to retest before adoption
How do you handle legal authorities or research?	Links to primary sources; clear statement of limits and verification steps
What audit trail is available?	Logs of inputs, outputs, users and review decisions



Build or buy decision

- A packaged legal tool covers your main task and integrates with your document system
- You need your own playbooks, templates and controls built in
- The team can maintain the library and run evaluations
- Costs are modeled for pilot and full use, including review time
- Exit and data return terms are documented



Remember that vendor assurances do not transfer your obligations. Your team remains responsible for the quality of work product and for compliance with data and professional rules. Ask your own advisers to review the contract terms, and revisit them when the vendor changes products or ownership.

WORKING QUESTION

If a regulator or client asked how an AI-assisted document was produced and checked, could you answer from your records?



30 / 60 / 90

A first quarter that produces evidence

10

Start with a verifiable task, protect confidentiality from day one, and keep named professionals accountable.

First 30 days: ground it

- Inventory tasks and rank by value, risk and ease of verification
- Draft a short use policy with security, privacy and the function head
- Select one pilot task and name a pilot lead and reviewers
- Record the baseline for time, quality and reviewer consistency
- Confirm approved tools and data terms with advisers

Days 31 to 60: test it

- Build the clause library or playbook the tool will draw on
- Create a test set with agreed answers and pass marks
- Configure instructions that require quoted pointers
- Run tests, study misses and tune
- Train reviewers on the verification standard

Days 61 to 90: decide

- Pilot on live work with sign-off recorded for each use
- Compare time and quality with the baseline
- Review corrections and incidents weekly
- Report results to the head of function
- Decide to scale, change or stop; plan the next task

GLOSSARY

Terms used in this publication

11

TERM	MEANING
Hallucination	A fluent output not supported by any source, including invented cases or quotes.
Privilege	Legal protection of certain confidential communications from disclosure; rules vary by jurisdiction.
eDiscovery	Identifying, collecting and reviewing electronic information for legal proceedings or investigations.
Technology-assisted review	Using software to prioritize or classify documents, validated by sampling.
Clause library	A maintained set of approved contract language and fallback positions.
Playbook	A documented set of preferred positions and escalation rules used when reviewing contracts.
Obligations register	A list of legal and regulatory duties, owners and controls.
Recall	The share of relevant items that a review process actually finds.
Precision	The share of flagged items that turn out to be relevant.
Grounding	Basing an answer on supplied documents rather than model memory.
Prompt injection	Text in a document or message that tries to override a tool's instructions.
Legal hold	A requirement to preserve records relevant to actual or expected proceedings.

A note on use

This publication is a public orientation resource, not legal, technical, financial or professional advice. Adapt it to your people, information, workflows and obligations, and check current tool terms, regulations and risks before acting.



TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

