






1

[illegible]

2

[illegible]

3

The AI Toolbox: What Counts as AI?

Predictive & Analytical AI
Finds patterns and makes predictions: e-discovery review, litigation analytics, spam filters. Lawyers have used it for years.

Generative AI
Creates new text, images, or code from a prompt: ChatGPT, Claude, Gemini, Copilot. Powerful, but it can “hallucinate.”

Agentic AI
Plans and carries out multi-step tasks on your behalf, such as browsing, filling in forms, or sending email. It needs the closest supervision.

Embedded AI
AI features built into tools you already own: Word, Outlook, Zoom, Westlaw, Lexis, and practice management software.

4

Myth vs. Reality

Myth: “AI looks up the answer.”
Reality: Generative AI predicts likely words. It can invent cases and quotations with total confidence (Mata v. Avianca, S.D.N.Y. 2023).

Myth: “Legal-specific tools don’t hallucinate.”
Reality: A 2024 Stanford study found leading legal research AI tools still gave incorrect or unsupported answers roughly 17% to 33% of the time.

Myth: “AI will replace lawyers.”
Reality: AI replaces tasks, not judgment. The duties of competence and supervision stay with you.

5

Myth: “If we don’t use AI, we have no AI risk.”
Reality: AI is already embedded in everyday software, and staff may be using personal accounts on their own (“shadow AI”).

Myth: “My chats with AI are private.”
Reality: Consumer tools may retain and train on inputs. A court held a defendant’s chats with a public AI tool were not privileged (U.S. v. Heppner, S.D.N.Y. 2026).


Myth: “One tool can do it all.”
Reality: Different jobs call for different tools, and different guardrails. That is the point of today’s session.

More Myths, Debunked

6



7



8



9

Risks of bias and discrimination in AI-generated content

Risk of Bias in AI Outputs
AI may reflect biases present in training data, leading to incomplete or unfair results.

User Responsibility (Input Controls)
Prompts must be carefully drafted to avoid generating inappropriate or misleading content.

Prohibited Content
Do not generate or rely on output that is offensive, discriminatory, or exclusionary.

10

Why Your Firm Needs an AI Policy



11

Why a Policy, and Why Now?

AI Is Already in Your Office
Staff may be using free tools on personal devices whether or not you have approved them. A policy brings "shadow AI" into the open.

The Rules Require Reasonable Measures
Rules 5.1 and 5.3 require managers to ensure that lawyers, staff, and the tools they use comply with professional obligations.

Clients, Courts & Carriers Are Asking
Outside counsel guidelines, judicial standing orders, and malpractice and cyber insurers increasingly ask about your AI rules.

12

What Your AI Policy Should Include

13

Core Policy Components

Scope & Purpose

Who is covered (attorneys, staff, contractors, vendors), which tools, and why the policy exists.

Approved Tools & Approval Process

A published list of vetted tools and a clear process for requesting and approving new ones.

Prohibited Inputs & Data Protection

What may never go into an AI tool, especially consumer tools without firm-level protections.

Human Review & Verification

No AI output leaves the firm without review by a responsible attorney or staff member.

14

Sample Policy Language: Prohibited Inputs

Strict Prohibition of Sensitive Data Input

Employees must not enter confidential or legally protected information into AI platforms unless using approved tools.

Approved Exceptions (firm tools only)

Only firm-approved generative AI tools that keep prompts within the firm's technical boundaries may be used with protected data.

Examples of Protected Data

Personal data, trade secrets, intellectual property, financial information, and other sensitive client, firm, employee, or third-party information.

15


Disclosure & Client Communication
When to tell clients and courts about AI use, and how the firm tracks court rules and standing orders.

Billing for AI-Assisted Work
Bill for time actually spent, don't charge for learning general-purpose tools, and disclose any AI costs passed through.

Training, Reporting & Enforcement
Required training, a clear path for reporting incidents, consequences for violations, and a schedule for review.

Core Policy Components (cont.)

16



Deploying Your AI Policy

17

Deploying Your Policy: A Roadmap

1. Form an AI governance team: leadership, IT and security, risk management, and staff
2. Inventory how AI is already being used and where the pain points are
3. Define the jobs to be done before shopping for tools
4. Vet vendors and pilot with non-client or anonymized data
5. Launch with training and a published approved-tools list
6. Monitor use, audit outputs, and revisit the policy at least annually

18



Finding the Right Tool: Match the Tool to the Task

Define the Job First
Start with the problem (research, drafting, intake, billing), not the product. The best tool is the one built for the job.

Vet the Vendor
Confirm how data is stored, used, and protected before anyone uploads a single document.

Pilot Before You Commit
Test with a small group on realistic, non-client scenarios. Measure time saved and error rates before a firm-wide rollout.

19

Vendor Due Diligence: Questions to Ask

Data Use & Retention
Is our data used to train the model? How long are prompts and outputs kept, and can we delete them?

Security & Access
Encryption, SOC 2 Type II or ISO 27001 reports, single sign-on, access controls, and where data is hosted.

Contract Protections

- Confidentiality terms and prompt breach notification
- A business associate agreement if health information is involved
- Clear ownership of inputs and outputs

20



Choosing the Right Tool: Types of AI Products

21

General-Purpose AI Chatbots

Examples

Consumer versions of ChatGPT, Claude, Gemini, and Copilot, often free or low cost.

Best Used For

Brainstorming, plain-language explanations, outlines, and polishing non-confidential writing.

Guardrails

No client information on consumer plans. Check whether inputs train the model, adjust privacy settings, and verify every fact.

22

Examples

Microsoft 365 Copilot, ChatGPT Enterprise, Claude for Enterprise, and Gemini for Google Workspace.

Best Used For

Summarizing internal documents, drafting emails, and working across the firm's own files.

Guardrails

They can see whatever the user can see, so clean up file permissions first. Confirm the contract bars training on firm data.

Enterprise AI Assistants

23

Legal Research AI

Examples

Lexis+ AI, Westlaw with CoCounsel, vLex Vincent, Bloomberg Law, and specialty tools such as Blue J for tax.


Best Used For

Finding authority, summarizing case law, and first drafts of research memos grounded in real sources.

Guardrails

Not hallucination-proof. Read every case you cite, confirm it says what the tool claims, and check that it is still good law.

24



Drafting & Contract Review AI

Examples
Spellbook, Harvey, and AI features in document assembly and contract management platforms.

Best Used For
First drafts, clause comparison, redlining against a playbook, and issue spotting.

Guardrails
Uploading client documents requires a vetted tool and may require client consent. An attorney must own the final draft.

25

E-Discovery & Document Review AI

Examples
Relativity aiR, Everlaw, DISCO, and other technology-assisted review platforms.

Best Used For
Prioritizing large document sets, privilege screening, and summarizing productions.

Guardrails
Validate results with sampling, follow protective orders on where data may go, and be ready to explain your process.

26

AI Meeting Notetakers

Examples
Zoom AI Companion, Microsoft Teams recap, Otter, and similar meeting assistants.

Best Used For
Internal meeting notes, action items, and dictation.

Guardrails
A third-party bot on a client call may put privilege at risk. Get consent where recording laws require it and control transcript retention.

27

Examples
Laurel, Clio Duo, Smokeball, and AI features in practice management platforms.

Best Used For
Capturing time automatically, drafting billing narratives, and organizing matters and deadlines.

Guardrails
Review every auto-generated time entry before billing (Rule 1.5). Never let a tool calendar a deadline without human confirmation.

Practice Management, Billing & Timekeeping AI

28

Client-Facing Chatbots & Intake Tools

Examples
Website chat assistants, automated intake forms, and AI answering services.

Best Used For
Scheduling, collecting basic intake information, and answering general FAQs.

Guardrails
Don't give legal advice, protect prospective-client information (Rule 1.18), disclose that users are talking to AI, and keep marketing claims truthful (Rule 7.1).

29



One Firm's AI Toolbox

- ▶ **COPILLOT** - Enterprise AI assistant embedded in Microsoft tools for secure drafting, summarization, and workflow support
- ▶ **SPELLBOOK** - Contract-focused AI tool designed to enhance drafting accuracy, identify risks, and improve review efficiency
- ▶ **HARVEY** - Legal AI platform supporting research, analysis, and complex drafting tailored to legal professionals
- ▶ **LAUREL** - AI-driven timekeeping system that automates time capture and improves billing accuracy
- ▶ **BLUE J** - AI-powered tax research tool that analyzes case law and predicts likely outcomes based on fact patterns. Supports faster, more data-informed legal analysis and decision-making

30

Guardrails for Client Confidentiality & Attorney Ethics

31

Universal Guardrails for Every Tool

Know Where the Data Goes

Read the terms. Who stores it, who can see it, and is it used for training?

Get Informed Consent When Needed

Before putting confidential client information into a tool that could disclose it, talk to the client.

Verify Before You Rely

Check every citation, fact, and number. The signature on the filing is yours.

Supervise People and Tools

Train staff, monitor use, and treat AI like a nonlawyer assistant under Rule 5.3.

32

Before You Use Any AI Tool, Ask:

- Is this the right tool for this job?
- Is this tool approved by my firm?
- Am I entering confidential client information, and can I anonymize it?
- Do I need the client's informed consent?
- Does the court or client require disclosure?
- Who will verify the output before it is used?

33

Right Tool, Right Job

- Start with the task, then choose the tool
- Put your AI policy in writing, and train on it
- Confidentiality, verification, and supervision go with every tool
- Questions? cwedwards@pbnlaw.com