A Practical Guide for **Understanding and Using AI in Your Court**



Second Edition

How to Prepare Your Court for the Future

equivant
COURT

Table of Contents

Introduction..... 2

Security, Trust, and Humans: Why Oversight Still Comes First.....3

Short Term AI Planning: Where Your Court Starts.....7

Long Term AI Planning: Building on What You Learn.....10

Responsibly Exploring AI: Giving Staff a Safe Place to Use It.....12

Reputable Resources for Understanding and Using AI.....15

Final Thoughts.....17

Introduction

The question facing courts has changed: it is no longer whether to use AI, but how to govern the AI that is already in the building.

When we published the first edition of this guide in December 2024, most courts we spoke with were still deciding whether AI belonged in their operations at all. Two years later, that decision has largely been made for them. AI now shows up in filings submitted to your court, in the software your staff already use, and in the browsers on their desktops, whether your court has approved it.

That shift changes the job for court leaders. The challenge is no longer getting comfortable with AI. It is deciding which tools are approved, what data can go into them, who reviews the output, and what happens when something goes wrong.

This second edition is written for the court administrators and IT leaders who make those decisions. It keeps the practical focus of the original and adds what two years of real-world use have taught the justice community:

- Real examples of major risks, and how courts and rule-makers have responded
- A short-term plan for evaluating and piloting AI, followed by long-term planning once you have a pilot behind you
- Guidance on giving staff a safe place to use AI, including how to tell consumer tools from enterprise-grade ones
- Checklists and worksheets to use with your team
- An updated list of reputable resources

Security, Trust, and Humans: Why Oversight Still Comes First

AI can support court work, but it cannot carry accountability. Every use of AI in your court needs someone to review the output and own the result.

Courts adopt new technology carefully for good reason. They hold protected case data and personally identifiable information (PII), and their decisions affect people's liberty, property, and families. That caution has served courts well with AI. The courts moving ahead are the ones putting governance in place first, then expanding use from there.

equivant Court takes the same approach with its own AI work: governance and data protection first, with a person in the loop on every workflow that touches case or client data. You can read how we put that into practice on our **AI Governance at equivant Court** page.

Where AI is showing up in courts

The uses we described in 2024 have moved from early experiments toward everyday operations:

- **Public-facing help.** Chatbots and self-help portals that point court users to the right forms, information, and next steps.
- **Document work.** Reviewing filings for completeness, checking redactions, and processing e-filed documents.
- **Data and operations.** Analyzing caseload data to spot patterns, bottlenecks, and backlogs.

Five risks to manage, and what they look like in practice

However your court uses AI, these risks should be managed under your local, state, and federal data security policies.

Risk	What Can Go Wrong	What it Looks Like in Practice
Accuracy and Reliability	AI produces confident, plausible output that is simply wrong.	In 2025, a California appellate court sanctioned an attorney $10,000 after finding that 21 of 23 case quotations in his brief were fabricated by AI tools. The court published the opinion as a warning to the bar. Someone at the court still had to catch and verify every citation. (McGuireWoods summary)
Data privacy and security	Confidential or nonpublic information entered into a public AI tool leaves your control.	California now requires court AI policies to prohibit the entry of confidential, personally identifying, or nonpublic information into a "public" AI system, meaning one where outsiders can access the data or use it for training. (Rule 10.430) See the callout in the Responsibly Exploring AI section.
Transparency and accountability	People affected by AI-assisted work cannot tell whether AI was involved or how to get a human review.	California requires disclosure when public-facing written, visual, or audio work consists entirely of AI output. New York takes a different route: no disclosure, but the filer is fully responsible for verifying accuracy. Courts need to decide which model fits their policy.

Risk	What Can Go Wrong	What it Looks Like in Practice
Dependence on technology	Review becomes a rubber stamp, and human expertise erodes.	In June 2026, a federal judge in Mississippi removed all four attorneys from a case over AI-fabricated citations. That included local counsel who had signed filings without checking them, a role the judge described as serving as a "rubberstamp." (Mississippi Today)
Legal and ethical considerations	AI use runs ahead of the rules or conflicts with them.	Rules are now being written state by state. See the next section.

Legal and ethical considerations: a state-by-state landscape

There is no single, comprehensive federal framework for AI use in courts. Rules are being written state by state, and often court by court, which means your court's obligations depend on where you are.

A few examples show the range of approaches:

- **California** adopted Rule of Court 10.430, effective September 1, 2025. Any court that did not prohibit the use of generative AI by staff or judicial officers had to adopt a use policy by December 15, 2025. The policy covers staff use for any purpose and judicial officers' use outside their adjudicative role.
- **New York** adopted Part 161, effective June 1, 2026. It permits the use of AI in preparing court papers and does not require disclosure. Anyone who uses AI must "carefully review the paper and independently ensure" it contains no fabricated material. This is useful model language for courts drafting their own policies.

- **Maryland's Appellate Court** referred an attorney to the Attorney Grievance Commission in late 2025 over fabricated citations and urged rule-makers to act (opinion, PDF). In March 2026, the state's Rules Committee voted unanimously to make clear that an attorney's signature certifies every cited authority has been verified.
- **Federal courts** face the same problem. The Mississippi case above shows that AI misuse is not just a state court issue (Mississippi Today).

The common thread across all of these is human accountability. Whoever signs, files, or publishes AI-assisted work owns its accuracy. For court leaders, that raises a practical question: when your staff find a suspected fabricated citation in a filing, what happens next?

For a closer look at the rules and proposals taking shape, and questions to ask any AI vendor about your court's data, read our blog post, AI Rules for Justice Are Already Being Written. Is Your Court in the Room?

Worksheet: Does your AI policy cover these risks?

□ Lists which AI tools are approved, and who approves new ones

□ Prohibits entering confidential, PII, or nonpublic information into public AI tools

□ Covers AI features built into software you already use, not just standalone tools

□ Requires a person to review AI output before it is used, filed, or shared

□ Names who is accountable when AI output is wrong

□ States when AI use must be disclosed, and to whom

□ Requires review of AI output for bias or harmful content

□ Sets a process for staff who find suspected fabricated citations in filings

□ Aligns with your state's court rules and your vendor contracts

□ Includes staff training and a schedule for reviewing the policy

Short-Term AI Planning: Where Your Court Starts

The best first step with AI is a small, well-governed pilot, not a big launch. The steps below can be completed within a few months and provide a foundation for long-term planning.

Many courts feel behind on AI and are tempted to skip ahead to buying a tool. The courts getting real value are doing the groundwork first. Work through these steps roughly in order:

1. **Reevaluate your workflows.** What's working well? What needs improvement? Identify processes that are slow, error-prone, or repetitive. Do this before you evaluate any new technology, so the problem drives the tool rather than the other way around.
2. **Know your rules.** Check the state court rules, local rules, and legislation that apply to your court, and keep them at the center of any internal policy. As the previous section shows, these are changing quickly.
3. **Find the AI already in use.** Ask staff which AI tools they use for work today, approved or not. Then check which software your court already licenses that has added AI features. You can't govern what you haven't found.
4. **Identify and rank your risks.** Use the five risks in the previous section as a starting point, and plan how you'll mitigate each one.
5. **Write or update your AI policy.** Use the worksheet in the previous section to check for gaps. Discuss the policy with your team regularly and update it as technology and rules change.
6. **Hold your vendors to the same standard.** Make sure every vendor and in-house application aligns with your AI policy. If one doesn't, plan to address it right away.

7. Strengthen your security basics. Use multi-factor authentication (MFA), encryption, and role-based access controls. Set a clear schedule for security updates in your policy and keep to it.

8. Check your data and infrastructure. AI is only as good as the data behind it. Make sure your data is accurate, well-organized, and stored in systems that securely support AI. Depending on your tech stack, this may become a long-term project, but start assessing it now.

9. Run one small pilot. Pick a low-risk, high-volume task that stays internal, such as summarizing internal documents or drafting routine correspondence for staff review. Set a clear success measure and an end date, then review what you learned before expanding.

> **“You’ll only be able to move forward with AI if you have the policies and standards in place that allow you to do so.**
>
> -Leigh Sheldon on Scaling Justice

Questions to ask any AI vendor

□ Is our court’s data used to train models for anyone else?

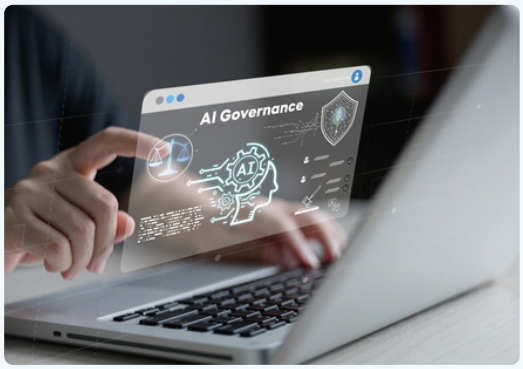
□ Where is our data stored and processed, and is it kept separate from other customers’ data?

□ Which security and compliance standards does the AI functionality meet (for example, CJIS or SOC 2)?

□ Where does a person review AI output before it affects a case or a court user?

□ Can we turn AI features on or off, and control who has access to them?

- Our blog post AI Rules for Justice Are Already Being Written. Is Your Court in the Room? goes deeper on vendor questions.



Worksheet: Plan your first pilot

Question	Your Answer
What task will the pilot cover?	
Why this task (volume, time saved, low risk)?	
What data will the AI tool touch? Any confidential or PII?	
Which approved tool will you use?	
Who reviews the AI output, and how?	
How will you measure success?	
When does the pilot end, and who decides the next steps?	

Long Term AI Planning: Building on What You Learn

Once your court has a pilot behind it, long-term planning becomes concrete. You are no longer guessing what AI might do. You are building on what you have seen it do in your own operations.



In the first edition, we suggested picturing your court 10 to 15 years out. AI now changes too quickly for that horizon to be useful on its own. A three-to-five-year vision, reviewed every year against what your pilots teach you, gives you direction without locking in assumptions that may not hold.

Start with the problems, not the technology

Ask your team: what would we fix first if we had more time and fewer manual steps? The answers usually point to a handful of areas where AI is already proving useful in courts:

- **Staff assistants for routine work.** Helping staff draft routine correspondence, summarize documents, and search internal procedures and manuals, with a person reviewing every output.
- **Smarter intake and e-filing.** Classifying incoming documents, flagging incomplete filings, and pulling key data into the case record so clerks spend less time on data entry.
- **Better help for the public.** Self-help tools that answer common questions in plain language and in multiple languages, with a clear path to a person when the question goes beyond general information.

- **Data-driven operations**. Spotting backlogs, forecasting caseloads, and flagging overdue reports or deadlines before they become problems.
- **Governance as a standing function.** An AI committee, a current inventory of AI tools in use, and regular reviews, so oversight grows alongside adoption rather than lagging behind it.

Turn the vision into a plan

For each goal, lay out the steps that will get you there and the conditions that need to be true before you move to the next stage, such as a successful pilot, an updated policy, or staff training. The more thorough your answers, the better.

Questions for your long-term plan

□ Why should we use AI for this goal, and what problem does it solve?

□ What does success look like, and how will we measure it?

□ What data does it require, and is that data ready?

□ How will we keep it secure and compliant with our rules?

□ Where does a person stay in the loop?

□ What risks does it create, and how will we mitigate them?

□ What would make us pause or stop?

□ Who owns it, and when will we review it?

Responsibly Exploring AI: Giving Staff a Safe Place to Use It

Your staff are probably already using AI at work. The safest response is not a ban on its own. It is an approved, secure tool, paired with clear rules on where data can go.

In 2024, we suggested getting comfortable with AI by trying it on personal tasks. That advice has aged. AI is now built into the email, search, and office software your staff use every day, and many people turn to consumer AI tools to save time at work without realizing where their data goes.

This has happened in some of the most security-conscious organizations. In 2023, Samsung restricted employee use of generative AI tools after staff uploaded sensitive internal code to ChatGPT (CNBC). Samsung's longer-term answer was not just a ban. It was developing an internal AI tool its staff could use safely (The Korea Herald via Asia News Network). That same year, several major banks, including JPMorgan Chase, restricted staff use of ChatGPT over concerns about sharing sensitive financial information (Forbes Australia).

Consumer AI tools my keep what you enter, and some use it to train future models unless you opt out. OpenAI, for example, says conversations in its individual ChatGPT plans may be used for training unless the user opts out, while its business plans are excluded by default (OpenAI).

Quick Tips for Safe AI Use

- Check with your court which AI tools are approved before using any of them for work.
- Never enter confidential, sealed, personally identifiable, or other nonpublic information into a tool your court has not approved.
- Sign in with your work account, not a personal one, when using an approved tool.
- If you're not sure, ask IT first.

Not all AI tools handle your data the same way

The same brand name can sit on both sides of this table. What matters is the account you sign in with and the agreement behind it.

	Consumer/ Personal Accounts	Exterprise/ contracted tools	Internal/self hosted tools
Examples	Free or paid personal ChatGPT accounts; consumer Copilot	ChatGPT Business or Enterprise; Microsoft Copilot signed in with an organization (Entra ID) account	An AI model that your court or vendor runs in its own controlled environment
Used to train models?	Maybe, unless you opt out (OpenAI)	Not by default for ChatGPT business plans; not used to train foundation models under Microsoft's enterprise data protection (Microsoft)	Set by your court
Who controls your data?	The provider, under consumer terms	Your organization, under a contract	Your court
Appropriate for court data?	No	Only if approved under your AI policy	Only if approved under your AI policy

A note on terms: "open-source" describes AI models whose code and weights are publicly released. It is not the same as "unprotected." A court or vendor can run an open-source model within its own secure environment, placing it in the right-hand column.

If your court wants to stand up its own safe space for staff to experiment, the essentials are:

□ An approved tool under an enterprise agreement, or one hosted in your own environment

□ Sign in with court credentials and MFA, so access follows your existing controls

□ Clear, written rules for what data can and can't be entered

□ A short list of starter tasks that don't touch case data, such as drafting internal communications or summarizing public documents

□ A named contact or a small group of champions that staff can go to with questions

□ A way to collect what staff learn, so it feeds your pilots and your long-term plan

Sidebar: AI and organizational trust

AI risk isn't only about the tools your staff use. It also affects who you hire. The FBI has warned employers about applicants using deepfakes and stolen personal information to interview for remote IT and database jobs, including roles with access to sensitive data (FBI IC3).

In one federal case, companies shipped laptops for a remote IT hire to an address that turned out to be a "laptop farm" run for overseas workers using a stolen identity (U.S. Department of Justice).

Courts hire remote IT staff and contractors with access to sensitive systems too. For those roles, verify identity through trusted channels, confirm that equipment ships to a verified address, and involve your IT security team in hiring.

Reputable Resources for Understanding and Using AI

The resources below are the ones we would point a court leader to first. Much more has been published for courts since our first edition, so we have organized this list by topic. It is not all-inclusive.

Governance and Policy

- **Principles and Practices for Using AI Responsibly and Effectively in Courts** (TRI/NCSC AI Policy Consortium, 2025). Core ethical principles and implementation practices for judges, administrators, and IT leaders, including starting with low-risk tasks.
- **AI Readiness for the State Courts** (NCSC, 2025). A three-level guide for courts just starting out, courts with a first AI project, and courts gathering feedback after one. It includes a readiness assessment, a governance tool, and case studies from state courts.
- **NIST AI Risk Management Framework** and its **Playbook** (National Institute of Standards and Technology). A widely used framework for identifying and managing AI risk, with suggested actions for each step.
- **IEEE GET Program** (IEEE Standards Association). Free access to selected IEEE standards, including standards related to AI ethics and governance.

Tracking the rules

- **AI in State Courts (NCSC).** An interactive map of state court orders, rules, proposed rules, guidelines, and policies on AI. This is the best single place to check what applies in your state.
- **GovTech's AI newsletter**. A free weekly newsletter on AI policy in state and local government.

Getting hands-on

- **Court Chatbots: How to Build a Great Chatbot for Your Court's Website (NCSC).** Practical guidance on designing, launching, and maintaining a public-facing court chatbot.
- **Trends in State Courts (NCSC).** The 2025 edition looks at how courts are using AI for self-represented litigants and guardianship monitoring, and includes 10 steps to launch an AI project.

Training and Education

- **TRI/NCSC AI Policy Consortium for Law & Courts.** Monthly webinars and resources on AI governance, ethics, access to justice, and workforce readiness, including role-specific AI literacy training for court staff.
- **AI and the Rule of Law: Capacity Building for Judicial Systems (UNESCO)**. A free online course on AI and the justice system.
- **Guidelines for the Use of AI Systems in Courts and Tribunals (UNESCO).** International principles and safeguards for keeping AI in courts human-led.
- **Artificial Intelligence in Justice and Public Safety (IJIS Institute, 2022).**

From equivant Court

- **AI Governance at equivant Court.** How we approach data protection, human oversight, and compliance in our own AI work.

- **AI Rules for Justice Are Already Being Written. Is Your Court in the Room?** The rules taking shape across states, and the questions to ask any AI vendor about your court's data.
- **Scaling Justice podcast: AI in Courts: Where Are We Really, and What Should You Do Next?** Guidehouse experts on practical first steps, from governance policy to small pilots



- **Maintaining Court Privacy During the Rise of Artificial Intelligence.** Steps your court can take now to protect its data.

Final Thoughts

We hope this second edition of A Practical Guide for Understanding and Using AI in Your Court helps you cut through the noise and focus on the steps that matter most for your court right now. If you have questions or want to talk with us about preparing your court for AI, contact us.

About equivant Court

equivant delivers customer-centric solutions that simplify justice by boosting efficiency, streamlining processes, and improving communication in courts and justice agencies. equivant's innovative technologies, including an award-winning case management system, have been developed by expert practitioners and have supported more than 350+ justice agencies in serving their communities efficiently, effectively, and equitably.