

Insights: Publications

5 Key Takeaways | Bring Us Your Hardest IP Problems: An AI-Powered Partnership for Building What Comes Next

August 20, 2026

Written by **Alton L. Absher, III**, **Charles W. Gray**, **Tyler McAllister** and **Patrick M. Njeim, MSc.**

Kilpatrick's Alton Absher, Charles Gray, Tyler McAllister, Patrick Njeim, and Tra Unger, joined by Spencer Thevenin

of Dish, presented “Bring Us Your Hardest IP Problems: An AI-Powered Partnership for Building What Comes Next”

at the 21st annual KTIPS (Kilpatrick Townsend Intellectual Property Seminar). The session opened with a live audience survey repeating questions first asked a year earlier, turned to the professional responsibility framework

now governing a lawyer's use of Generative AI, and then demonstrated six agentic tools the firm built in-house, running from invention intake and quality assurance through portfolio analysis and the assertion lifecycle, from developing a demand letter to preparing for oral argument.

The panel provides these key takeaways:

1. The lawyer, not the tool, remains accountable. Competence under Model Rule 1.1 and its Comment 8 now reaches the benefits and risks of relevant technology, and the USPTO Rules of Professional Conduct mirror that duty. Ethics authorities treat Generative AI as a nonlawyer assistant subject to the supervisory obligations in Model Rules 5.1 and 5.3, and guidance issued in 2026 scales that obligation to the autonomy of the system. In practice, what courts have sanctioned is not the use of AI. It is the failure to verify, and the failure of candor once an error surfaces.

2. Practitioner use has moved from casual to committed, and from convenience to substance. Our 2025 and 2026 survey results reveal an evolution in how legal professionals leverage Generative AI in patent prosecution. We are witnessing a clear transition from casual experimentation to daily, deeply integrated reliance. Practitioner usage has matured beyond basic entry-level tasks (such as search and summarization) toward substantive workflows, including novelty evaluations, quality checking, and document drafting. Furthermore, the motivation behind adoption has shifted from a fear of being left behind to top-down organizational mandates focused on turnaround speed and enhanced work

product quality. Ultimately, Generative AI is evolving from an ad-hoc efficiency hack into an integrated component of modern patent practice.

3. For work that is specific to a practice, building beats buying. The firm's AI development group is creating tools that span the IP lifecycle, all of them developed internally, to close the gap between what practitioners need day to day and what the available products actually do. The reason to build is narrow and worth stating plainly: the tooling has to understand prosecution practice, and general-purpose products do not. Each tool starts from a workflow the firm already runs and is built around it. The transferable question for an in-house team is not build versus buy in the abstract, but which parts of your workflow are specific enough that a generic tool cannot know them.

4. The step change is agentic workflows, not chat. What was demonstrated were multi-step systems that chain tools together and return a work product, such as an assertion workup, a portfolio analysis, or a ranked set of predicted questions from the bench, rather than a chat window returning prose. What makes them usable is where they stop. Each workflow pauses at defined points and requires a lawyer to approve before anything consequential happens. When evaluating any tool, ask what it will not do on its own, not only what it can produce.

5. Put the guardrails in the build system, not only in a policy document. Every tool goes through a written development lifecycle approved by information security: no engineer releases their own work, matter content runs only where nothing is retained and nothing trains an outside model, the system may read but cannot send or file without a person, and each tool is scoped, inventoried, and monitored. The tools also report their own limits, labeling a conclusion low confidence where the underlying record is thin. A policy depends on everyone following it every time. A build system does not.

For more information, please contact:

Alton Absher, aabsher@ktslaw.com

Charles Gray, cgray@ktslaw.com

Tyler McAllister, tmcallister@ktslaw.com

Patrick Njeim, pnjeim@ktslaw.com

Tra Unger, tunger@ktslaw.com