

Why your invention disclosure process is losing ideas before they reach you

The friction in most invention disclosure processes isn't obvious from inside the IP team — the forms seem reasonable, the timelines fair, the inventors engaged. And yet the pipeline stays thin, the same researchers keep showing up, and the IP team suspects it's missing ideas that never made it through the door.

The friction is visible from the inventor's side. An engineer generating two or three promising ideas a quarter is also writing code, meeting on product specs, and managing a workload that has nothing to do with IP. The disclosure form arrives as one more obligation — structured around legal requirements rather than engineering language, with no clear signal of what happens next. Most inventors defer it. Many don't come back.

This guide covers what it takes to change that dynamic — moving from a process where inventors come to you to one where you find the signal before it's ever formally submitted. It draws on a panel conversation between Tom Wolfe, Director of Legal for IP and Product at Kodiak AI; Alyssa Sandrowitz, Associate General Counsel for IP at Woodward; Ann Yang, Senior Director of IP at RefleXion Medical; and moderator Alec Sorensen, CEO of Tradespace.

You can read this guide alongside the webinar replay or independently.

[Watch the webinar replay →](#)

What the disclosure process is competing with

Ann Yang's first attempt at a disclosure process at RefleXion came from her law firm background: a standard form, a structured submission, a presentation to the VP of R&D and CTO. It was logical. Nobody used it. Engineers didn't want to present to their manager's manager. The form felt like extra work. Submissions came in slowly, and the ones that did arrive were often half-baked.

So she modified it. She'd meet with inventors first, take notes, build the summary herself, and present it to leadership on their behalf. That worked better — until the underlying problem reasserted itself. People still didn't have time to put together even a short description, even knowing someone else would take it from there.

Alyssa Sandrowitz inherited a similar dynamic at Woodward. The disclosure form she found when she arrived was seven pages long. Annual submissions hovered around 40 for a \$3.5 billion aerospace company with significant R&D investment. She describes the engineers' response plainly: why do we need to do this? It takes too long.

Both of those experiences point to the same underlying problem. A disclosure process designed around legal requirements asks engineers to translate their work into a different language at a moment when they have no particular reason to do so. The process is competing with everything else on their plate — and it usually loses.

Moving from passive intake to active capture

The clearest version of the shift from friction to flow is the difference between waiting for disclosures and going to find them.

Tom Wolfe at Kodiak runs his disclosure pipeline on a simple rule: either a draft is in progress, or it's waiting on the inventor. There's no third state. He's built that discipline by integrating into the development environment itself — monitoring changes to the GitHub repo, new safety documentation, new rationale documents. He's looking in the places where inventors are working instinctually, in the systems they use every day, rather than waiting for them to surface something through a form.

When he finds a promising lead, he doesn't open with a patent conversation. He goes back to the inventor and asks about the work — what they're doing, what problem they're solving. Engineers are easy to get talking about what they're working on, he notes. The IP frame comes after the conversation, not before it.

Ann at RefleXion built a version of this through the company's quality and regulatory systems. Medical device development requires design control documentation — product specifications, user manuals, materials that have to be released into the quality system regardless of what the IP team needs. By getting involved in that review process, she could see what was going into the product without asking engineers to create additional documentation or attend additional meetings. The signal was already there. She just needed a seat at the table where it lived.

Alyssa at Woodward embedded herself into the product lifecycle process. Woodward runs through defined design review phases — preliminary, critical, detailed — that are required by customers and partners. Rather than scheduling separate invention harvesting meetings, her team now sits in on those reviews, surfaces innovations in context, and takes the existing artifacts rather than asking engineers to generate new ones.

The common thread is that all three moved primary engagement with inventors upstream — before the disclosure process starts, not at the moment it needs to happen.

Three things most IP teams **get wrong about invention disclosure**

Getting more disclosures requires reducing friction, but not the friction most IP teams focus on. This piece covers the inventor-side dynamics that drive submission rates — including why inventors routinely undervalue their own ideas and what the IP team can do about it.

[Read the article →](#)

What AI changes **about the capture process**

The practical constraint on active capture has always been bandwidth. Going to find inventions — sitting in design reviews, monitoring development workflows, building individual relationships across engineering teams — takes time that lean IP teams rarely have in excess.

AI shifts that constraint in two ways.

The first is signal detection at scale. Tom at Kodiak described building prompts that review changes to the GitHub repo and flag things that merit a look — automating the monitoring work so that his attention goes to the decisions that need human judgment, not to the scanning that precedes them. For software-native companies, this kind of integration into the development environment is increasingly viable. For companies where patient data or regulatory systems are involved, the same principle applies through different access points — design control documentation, FDA submission materials, quality system artifacts.

The second is what happens once a disclosure exists. Ann described engineers who can architect systems but struggle to translate what they've built into written descriptions. AI is useful here — not to replace the IP team's understanding of the invention, but to help inventors create a starting point from the technical materials they've already produced. She noted that she still holds closely to drafting the heart of the invention herself, precisely because that process keeps her intimately familiar with the technology. Her caution about AI-generated output that looks polished is one to take seriously: when a draft arrives fully formed, it can suppress the critical review it needs. The tool is most useful when it compresses the distance between what the engineer knows and what the IP team needs to act, not when it substitutes for either.

Alyssa at Woodward uses AI to do the initial triage work before disclosure committee meetings — prior art search, novelty score, commercial relevance — so that each disclosure can be reviewed in about 15 minutes rather than in a session that runs long and still doesn't get through the queue. The legwork gets done before the meeting, not during it.

What the inventor relationship looks like **when the process works**

The way inventors experience the disclosure process shapes whether they come back.

Ann Yang at RefleXion described the pattern she sees most often: an engineer submits an idea and then hears nothing for months. No update, no decision, no explanation. From the inventor's perspective, the idea went into a black hole. The next time they have something interesting, they may not bother.

Her model runs differently. She described it as bringing a client service mentality into the in-house role — keeping inventors apprised of what's happening with the ideas they've submitted, explaining why something is being deprioritized rather than just deferring it, telling them when something that's not ready now might be ready after another product cycle. Engineers love to know the status of things, Alyssa observed. A portal or dashboard that lets them see where their disclosure is in the process gives them that without requiring the IP team to generate individual updates.

Tom added a related approach: giving ownership of the disclosure queue to engineering managers, making it part of their process rather than an external obligation. At a previous company, the invention disclosure review meeting was engineering-led — managers came with their IP questions, and Tom provided what he needed in return. The review became part of what engineering was already responsible for, not a separate meeting that IP had to convince people to attend.

Alyssa at Woodward tracked the results explicitly: how many rounds of back-and-forth did it take before a disclosure was ready to act on? As that number went down, the quality of what came in went up, and the number of inventors who submitted a second time increased. Those metrics won't appear in any benchmarking survey, but they tell you whether the process is working.



Why Woodward's IP lead chose a platform built around AI — **and what changed**

Alyssa Sandrowitz at Woodward needed a next-generation IPMS that could mirror her company's product lifecycle process and free her team to focus on bigger strategic decisions. This short video covers what she was looking for, what set Tradespace apart, and her advice to IP professionals on embracing AI as a tool.

[Watch the case study →](#)

Five questions that determine whether the process works

Tom Wolfe's approach is simple enough to describe in a sentence. Building it requires answering a set of questions that most IP teams haven't asked explicitly.

Who owns outreach to inventors? If the answer is "the inventors come to us," the pipeline will reflect that. Active capture requires someone accountable for knowing which inventors have something in progress and where each one is.

What does the first interaction look like? If it's a form, inventors will treat it like a form. If it's a conversation — or better, a review of work they've already done — the friction drops and the IP team gets richer information.

How fast does the IP team respond? Ann was explicit about this: inventors shouldn't feel like their ideas went into the abyss. The gap between submission and first response sets expectations for everything that follows. A quick response, even if it's just an acknowledgment of status, changes the relationship.

What happens when the answer is no — or not yet? Ann's approach to this is to explain the reasoning, not just deliver the decision. Something that doesn't fit the current product cycle might fit a later one. Saying that, and following up when the timing changes, keeps inventors engaged.

What gets measured? If the only metric is disclosure volume, the process optimizes for volume. If it also tracks response time, back-and-forth rounds, repeat submitters, and how often disclosures arrive with enough information to act on quickly, the process optimizes for something closer to what the IP office is trying to build.

The specific answers will be different for every office. The structure — active, visible, fast, relational — is what's transferable.

Test Tradespace on your own IP to see how you can save time, develop stronger disclosures and filings, and maximize value for the business.

[Get a demo →](#)