HOW TECH TRANSFER OFFICES MOVE FROM AUTOMATING TASKS TO REDESIGNING THE WORK

Is your TTO using AI to its fullest potential?

There is a version of AI adoption in tech transfer that looks successful and isn't. The market landscapes come back faster, the outreach lists build themselves, the first-pass evaluations take hours instead of days — and then, six months later, the office is doing the same work it always did, just with less friction.

The difference between that version and one that changes how a TTO operates comes down to a single question: is AI being used inside an existing workflow, or has the workflow been redesigned around what AI makes possible?

This guide works through that question in practice. It draws on a panel of three TTO leaders at different stages of the transition — Priya Sinha Cloutier at University of California, San Diego; Alecia Muth, Strategic Partnerships and Ventures Manager at Gladstone Institutes, an independent nonprofit biomedical research organization with 33 labs; and Jennifer Souter, who manages tech transfer for the Texas Tech University System across five component institutions. The panel was moderated by Joe Fiore, Executive Vice President at Tradespace and former Vice President at the Medical University of South Carolina's tech transfer office.

You can read this guide alongside the webinar replay or independently — the frameworks here stand on their own.

Watch the webinar replay →

When faster isn't different: the gap between automating tasks and redesigning work

Task-level automation is a legitimate starting point. Running AI on outreach lists, market landscapes, and first-pass research saves time that can be redirected. But offices that stop there tend to find that the time savings are tangible without being transformative — because the underlying workflow hasn't changed, only its pace.

The gap shows up most clearly in what happens to decision quality. An office doing things faster but making filing decisions on the same subset of disclosures it always evaluated ends up with the same portfolio. The backlog clears faster, but the coverage gaps remain.

tradespace®

What changes when AI is built into the workflow is the range of what the office can act on. Priya at UCSD files on roughly 80% of provisionals and converts nearly all of them — with 500 disclosures a year and a small team. Standard practice would be to do thorough landscape work on a subset, prioritized by early signal. With AI embedded in the evaluation process, her office can run a preliminary white space analysis on every incoming disclosure regardless of volume. Every technology gets the same first look, and filing decisions that follow are informed by consistent data rather than by which disclosures happened to get attention.

At Gladstone Institutes, the change showed up in commercialization. Where the team previously built opportunity pages by hand, pushed them to a list of known contacts, and hoped for responses, they now move from disclosure to targeted partner outreach inside one workflow — with AI surfacing suggested partners, relevant financial context, and strategic fit signals that enable outreach at a specificity level that wasn't viable manually.

Jennifer at Texas Tech framed the test plainly: if time saved is the only metric, the office is optimizing the tasks it already has. The question she uses to gauge whether AI has changed the work is whether the office is making better decisions earlier and across more of the portfolio.

How the workflow changes — and what makes it harder than it sounds

Offices seeing the best results didn't start by selecting a tool. They started by mapping the existing workflow — where decisions get made, who's accountable, where information needs to flow — and then identified where AI could sit inside that map. Starting with a tool and designing the workflow around it tends to produce faster task completion rather than a different way of working.

Alecia at Gladstone described the transition in two stages. The first was bolting on: using AI for discrete tasks without changing the underlying process. The second was building in: stepping back to ask where in the process AI was touching things, where it wasn't, and whether the workflow itself should change. The IP assessment that now runs on every Gladstone disclosure isn't a tool added alongside the existing process — it's a standardized first step that feeds into a structured group decision covering inventor interest, spinout potential, and licensing landscape, with human judgment layered in after the data exists.

Standardization is what makes the portfolio legible. When every disclosure goes through the same evaluation, decisions can be made on consistent data, comparison across technologies becomes possible, and portfolio-wide visibility becomes achievable rather than aspirational.

The tool stack question sits inside the same redesign problem. Most offices accumulate disconnected tools — a CRM, a docketing system, an AI evaluation product — and discover that the handoffs between them create friction rather than removing it. Jennifer was direct about the downstream consequence: when AI expands the front end of the funnel without a corresponding redesign of what follows, the bottleneck doesn't disappear. It moves. If AI enables an office to evaluate three times as many disclosures, the question immediately becomes what happens to those disclosures afterward. Offices that have redesigned only the intake stage will feel that pressure quickly.

GLADSTONE
INSTITUTES

How Gladstone cut triage time by 30—50%

After embedding AI into disclosure evaluation and partner identification, Gladstone Institutes reduced market research and triage time by 30—50%. The case study covers how a small office built the practice and where the gains showed up.

See the case study →

What the role looks like when AI handles triage — for staff and for inventors

When AI handles first-pass evaluation, partner identification, and market screening, the work that remains shifts in character. Less time goes to assembling information; more becomes available for the judgment calls AI can't make — what the inventor is trying to do, whether the science is ready, whether the right team is in place for a spinout, what the relationship with outside counsel should look like for a given technology.

Alecia at Gladstone described the effect as a reorientation. Work that previously required most of her day — building contact lists, writing opportunity pages, doing market research from scratch on unfamiliar technology areas — now gets handled upstream. What's left is work that was previously hard to reach: mapping publication records against investor data against IP portfolios to surface pre-IP opportunities, building tools in plain language, working through spinout formation questions with scientists.

Jennifer's framing was more measured. The shift in what the role requires changes what the office should hire for and where senior staff spend their time — but the leverage only lands if the team isn't spending recovered hours validating and correcting AI outputs. The measure she returned to: are we making better decisions across more of the portfolio?

The picture for inventors changes too. Priya at UCSD noted that Microsoft Copilot is available enterprise-wide at her institution. If a principal investigator uses it to draft portions of a disclosure, the legal questions that follow don't disappear because the TTO wasn't involved. The controls inside the office are the tractable part. Building the education, policy, and communication infrastructure to govern how the PI community uses AI — across a campus of researchers with varying levels of awareness — is work that runs in parallel and compounds over time.

One TTO's approach to reaching inventors before they reach you

University of Arkansas for Medical Sciences' BioVentures built a dedicated inventor outreach role and a proactive engagement program around it — including how to reach PI communities that don't know the TTO exists and how to govern AI use across a research campus.

Read the article →

What portfolio-wide coverage looks like — and what it takes to sustain it

The most important consequence of building AI into the workflow is coverage.

Under a traditional model, TTO attention gets rationed by volume. An office with a small team and 200 disclosures makes choices about which technologies get thorough evaluation and which get a quick pass. That triage is a resource decision, but it shapes the portfolio — because technologies that don't get evaluated tend to sit, and the ones that do tend to be the ones with existing momentum.

When AI runs a standardized evaluation on every disclosure, that constraint lifts. Every technology gets the same first-pass treatment. Technologies from departments that have historically underperformed on commercialization get evaluated alongside those with established PI relationships. Filing decisions follow from consistent data rather than from what the team had time to look at.

Jennifer described the goal her office is working toward: moving from selective insight on initial disclosure activity to portfolio-wide visibility, with AI supporting assessment throughout the lifecycle from disclosure through maintenance and commercialization. Getting there requires solving the downstream question — what happens to all the evaluated technologies that follow — but solving it is what makes the coverage gains durable.

Three practices the best-performing offices share

Offices that have moved furthest from bolted-on to built-in tend to share three practices.

They map the workflow before selecting tools, identifying where decisions get made, who is accountable, and where information flows — then determine where AI sits inside that map rather than around it.

They standardize evaluation before scaling it, so that every disclosure goes through the same first-pass process, comparison across technologies becomes possible, and group decisions can be made on consistent data.

They build PI governance infrastructure alongside internal controls, starting the education and policy work early enough that it stays current with how rapidly AI tools are spreading across research campuses.

The panelists were asked at the close what one action an office should take in the next 30 days. Alecia's answer was to start somewhere simple — pick one task that doesn't involve confidential information, hand it to AI, and evaluate the output. Priya's was to map the pain points first, find the friction, and see where automation has something to offer. Jennifer's was to pick something harder than a simple task, to understand what AI can and can't do before building a practice around assumptions.

The variation in those answers is itself informative. There's no single entry point that works for every office. What the panelists share is the framing: they're not asking how to use AI to do what they're already doing. They're asking what the work should look like — and whether AI changes the answer.

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

Get a demo →