



# Legal AI Playbook For Higher Education

## How University Legal Teams Can Move From Reactive to Strategic With AI

### Featuring:

- A practical framework built for university legal teams
- The full demand picture: contracts and beyond
- AI-powered self-service for policy and compliance questions
- Matter management, reporting and analytics
- Use case prioritization guidance
- Vendor evaluation criteria for higher education
- A quick-start action checklist

## CONTENTS

# Inside this playbook

<b>01</b>	Why University Legal Teams Need a Front Door	03
<b>02</b>	Start With Your Institution’s Pain Points, Not the Technology	06
<b>03</b>	Choosing the Right Starting Point	08
<b>04</b>	What This Looks Like in Practice	11
<b>05</b>	Measuring What Matters	14
<b>06</b>	Driving Adoption in a University Environment	17
<b>07</b>	Evaluating Vendors	20
<b>08</b>	Implementation Guide for University Legal Teams	24
<b>09</b>	Your Quick-Start Checklist	27

# Why University Legal Teams Need a Front Door

University legal offices operate at the intersection of contracts, compliance, sponsored research, faculty matters, and policy guidance – most of it arriving through email, Teams, and informal asks. There is no queue. No priority signal. No record of how long a request has been waiting, or whether it actually requires an attorney.

Work requiring legal judgment competes for attention with questions a policy document could have answered, requests that should have been forms, and threads that needed six more pieces of information before legal work could even begin. This is not a staffing problem. It is an infrastructure problem.

“Legal departments don’t generate revenue, so every time we ask for new technology it comes under heavy scrutiny. We end up adding more people instead of fixing the process. Meanwhile, 75% of our work is coming through the inbox – so we don’t even have a complete picture of what legal is doing. Then every quarter, we’re asked to report on our time, but it’s all based on memory and estimates. It’s not real data.”

– Chief Regulatory Counsel, large public university

## Why a ticket portal isn’t the answer

The most common first response is a ServiceNow or Jira portal. On paper, it sounds like progress. In practice, it produces a structured version of the same problem: the requestor fills in three required fields, the attorney needs six more, and the back-and-forth begins. A form is not triage. A portal that captures the minimum and routes everything to the same queue is a slightly more organized inbox.

## What a legal front door does

A front door is the single entry point for every legal request, whatever the channel or matter type. The difference between a front door and a ticketing portal is functional, not cosmetic. A front door understands the request through conversational AI intake, collects what is needed, classifies it, and routes it to the right resolution – sometimes an instant policy answer, sometimes a self-service workflow, sometimes a matter assigned to the right attorney with full context already attached.

### For universities, a legal front door does three things together:

- 1. Handles the full universe of legal demand.** Not just contracts. Policy questions, FERPA guidance, compliance queries, faculty research inquiries, vendor agreements, employment intake – every type of request enters the same door and is routed intelligently.
- 2. Integrates with what you have and fills what you don't.** If your institution runs a contract lifecycle system, the front door routes contract work to it. If it doesn't, the front door delivers that experience natively. Either way, legal gets a complete operational picture instead of a fragmented one.
- 3. Is owned and operated by legal, not IT.** Workflows, intake logic, routing rules, and self-service responses are configured and updated by legal directly. IT's role is limited to the institutional gates that exist for any new software: security review, procurement approval, and SSO. After that, legal builds and maintains everything.

## The cost of inaction

The cost of operating without a front door does not arrive as a single visible failure. It accumulates – in research delays, in decisions made without legal involvement, in faculty learning to route around legal because it is faster, in institutional credibility eroded by a function that cannot demonstrate its own value. By the time the pattern is visible, a great deal has already been lost.

## What shifts when you have a Legal Front Door

FROM	TO
Requests land in the inbox with no priority signal	Every request enters with full context at intake
Portals capture the minimum and trigger back-and-forth	An AI agent collects what's needed and routes intelligently
Attorneys answer the same policy questions repeatedly	An AI agent handles policy and compliance questions instantly
Policy answers sit in a SharePoint site no one can navigate	The AI agent surfaces the answer at the moment of need
Contract work scatters across email, Teams, and meetings	Contract work routes to your CLM, or is handled natively
Requests that fall through the cracks leave no record	Every request is captured, tracked, and accounted for
Triage is informal and based on who asks loudest	Triage is structured – based on risk, deadline, and matter type
Workload is invisible across the team	Every attorney's workload is visible in real time
Quarterly reports are built on memory and estimates	Reporting is built on real data captured at intake
Legal is a bottleneck the institution works around	Legal is a service the institution can predict and trust

## Start With Your Institution's Pain Points, Not the Technology

The wrong sequence starts with a vendor demo. The right one starts with the institution – an honest look at where legal time is going, where demand is coming from, and what it costs when legal cannot respond at the speed the institution needs. The technology conversation comes after that, every time.

### The demand picture most offices have never seen

Ask the average general counsel how many legal requests their office received last semester. The honest answer is approximately. Which departments generate the most demand, which matter types take longest, which requests most often arrive without enough information – almost universally: we don't have that data.

This is not a failure of effort. It is a structural consequence of an intake model built on email and relationships. Unstructured work leaves no data trail. Without data, the Office of Legal Affairs cannot manage its workload strategically, demonstrate its value evidentially, or build a compelling case for the resources it needs.

The first job of a legal front door is not to make legal faster. It is to make legal visible – to the team, to leadership, and to the institution it serves.

### Every department has a dashboard. Legal has a best guess.

Sponsored research reports grant activation timelines. HR reports time-to-hire. Procurement reports vendor cycle times. Each walks into the quarterly meeting with evidence – specific, quantified, connected to outcomes the institution already tracks.

Legal walks in with a narrative assembled from email searches, memory, and attorney estimates. In a room full of evidence, narrative lands differently than everything that came before it. That gap is not just uncomfortable. Functions that cannot demonstrate value do not get invested in.

## **Eight questions to answer before any vendor conversation**

- 1. Which three request types consume the most attorney hours each semester?** Not the most complex – the most cumulative.
- 2. Through how many channels does legal demand arrive?** Email, Teams, walk-ins, phone calls, forwarded threads, departmental liaisons.
- 3. What is the current average turnaround on your highest-volume request type?** If you don't know, that is the finding.
- 4. How much of that turnaround is attorney work —** and how much is waiting, back-and-forth, and missing information?
- 5. What percentage of incoming requests could have been resolved through self-service** if the right infrastructure existed?
- 6. What does success look like to your provost or CFO — and can you measure it today?**
- 7. Who else in the institution is most affected by current response times?** They are your allies in building the case.
- 8. Does your institution have an existing AI use policy** that shapes what you can deploy?

### **Establish the baseline before anything else**

Before any workflow goes live, record where you are today. Write it down. Date it. Keep it. For most Offices of Legal Affairs, seeing the demand picture clearly for the first time is itself the breakthrough – because it changes the conversation about what legal needs, and why.

## Choosing the Right Starting Point

Teams that make real progress with AI don't launch transformation programs. They start with one workflow, build it well, prove it works, and use that proof to fund the next.

The right first use case delivers a visible result within a single semester, generates a before-and-after metric leadership understands, and builds infrastructure the next workflow can sit on top of.

### Two prerequisites before you choose a tool

**Legal must own the tool directly.** Not something IT builds and maintains on legal's behalf – something legal configures, updates, and expands without filing a request. When a policy changes or a sponsor requirement shifts, legal changes the system that day. This is the most common failure mode in university legal technology: a workflow that ages out of alignment because legal cannot touch it without a ticket.

**The platform must scale.** Easy to deploy is not enough. A tool that handles one workflow but stalls when volume grows leaves the same residue as a tool that was too hard to implement. The right platform is built to carry hundreds of workflows, multiple matter types, cross-departmental intake, AI-powered self-service, system integrations, and the reporting layer legal has never had.

### The right first use case

Two variables determine where to start: volume and risk. High-volume workflows compound time savings across every instance. Low-risk workflows are safer starting points while the team is still learning. The right first use case sits at the intersection – and is tied to an outcome that matters to someone outside legal.

At most research universities, that means research agreements and NDAs. They are high volume, repetitive, and most often delayed by process friction rather than legal complexity.

The other category – consistently underestimated – is policy and compliance questions. Attorneys spend their days answering the same questions by email, drawing on policy documents that exist but cannot be navigated by the people asking. An AI agent does not write new policy. It makes existing policy findable at the moment of need.

### The prioritization matrix

Use the matrix below as a sequencing tool. Start top-left. Expand deliberately. Never automate the bottom right.

		LOW RISK	HIGH RISK
HIGH VOLUME	<p><b>START HERE</b></p> <ul style="list-style-type: none"> <li>Research NDAs</li> <li>Standard vendor agreements</li> <li>Data use agreements</li> <li>Certificate of insurance requests</li> <li>Policy &amp; compliance questions</li> <li>Routine MOU intake</li> </ul>	<p><b>EXPAND CAREFULLY</b></p> <ul style="list-style-type: none"> <li>Faculty employment agreements</li> <li>Union grievance intake</li> <li>Sponsored program compliance</li> <li>FERPA request routing</li> <li>Conflict of interest disclosures</li> </ul>	
LOW VOLUME	<p><b>AUTOMATE OPPORTUNISTICALLY</b></p> <ul style="list-style-type: none"> <li>Gift acceptance reviews</li> <li>Event liability waivers</li> <li>Intergovernmental agreements</li> <li>Standard lease renewals</li> </ul>	<p><b>DO NOT AUTOMATE</b></p> <ul style="list-style-type: none"> <li>Title IX matters</li> <li>Major litigation</li> <li>Federal audit response</li> <li>Board governance</li> <li>Accreditation legal review</li> </ul>	

## **A note on Title IX and faculty employment**

These will come up the moment AI is mentioned in a university context. Address them directly. The intake layer is appropriate. Everything beyond it is not.

For Title IX matters, the AI agent acknowledges receipt, collects initial information, and routes to the appropriate attorney or coordinator. It does not assess, advise, or communicate anything resembling a finding. The same boundary applies to faculty employment and tenure matters. Being explicit about these limits is not a limitation of the system. It is evidence of the judgment with which it was designed.

## What This Looks Like in Practice

Most descriptions of legal AI focus on what the technology does. This one focuses on what it changes – for the attorney, the faculty member, and the department administrator who has been routing around legal because it was easier than waiting. The scenario below is illustrative, drawn from patterns common across university legal teams.

### Before the front door

A research university. Mid-sized. Active sponsored research portfolio. Four attorneys handling agreements, employment matters, student affairs, procurement, compliance, and policy questions.

Incoming requests arrive by email, Teams, and walk-in. No intake system. No triage. Roughly forty requests reach the team in a typical week. Some are research NDAs. Some are policy questions answered in three minutes that consume thirty once the back-and-forth is factored in. Many arrive with so little information that the first action is to ask four follow-up questions and wait two days.

No one knows volume by department. No one knows the percentage requiring legal judgment. At quarter-end, the general counsel assembles a report from memory, email searches, and attorney estimates. It is the best available answer. It is not real data.

### The front door opens

The team deploys a front door with two workflows: a research NDA intake process and an AI agent configured to answer the thirty most common policy and compliance questions.

Both are built by legal. IT's role is limited to security review, procurement approval, and SSO. The NDA workflow takes two weeks to build and test. The AI agent takes one week to configure – fed with institutional policy documents, the team's standard guidance, and explicit escalation triggers. Every response is reviewed and approved by an attorney before anything goes live.

On day one, a faculty member submits an NDA through the intake portal instead of by email. The form collects what the attorney needs. The matter that would have taken four days to gather information for is ready for review in forty minutes.

The same day, a department administrator asks whether an unpaid advisory role at a startup requires a conflict of interest disclosure. The AI agent cross-references institutional policy and answers in forty seconds. No attorney involvement. No ticket. The office has a record that the question was asked, what answer was given, and that it was resolved without escalation.

### Ninety days in



The number that matters most isn't turnaround or deflection. For the first time, the Office of Legal Affairs has a complete record of what came in, what type of request it was, which department submitted it, how long it took, and whether it required attorney involvement.

### The conversation that follows

At the end of the semester, the general counsel walks into the quarterly review with a dashboard. Research NDA turnaround down 79%. Forty-four percent of incoming requests resolved through self-service. Policy questions answered in seconds rather than days. Sponsored research escalations due to legal delay: zero in the past quarter.

The provost asks a different question than usual: what do you need to keep going? That question has never been asked before. It was earned – not by working harder, but by making the work visible.

## **What comes next**

The NDA workflow and the policy agent were the foundation, not the destination. Data use agreements get their own workflow. Material transfer agreements follow. The agent extends to additional question types. Six months in, fifteen workflows are running across six matter types. Forty-eight percent of requests are resolved without attorney involvement. The team has not grown. The infrastructure between the team and the institution has.

## Measuring What Matters

Legal has historically struggled to quantify its value. AI changes that, but only if measurement is built in from the start. Every workflow you automate is an opportunity to capture before-and-after data that connects legal performance to institutional outcomes leadership already tracks: grant activation, compliance, deal velocity, outside counsel spend.

### What having data actually changes

**Reputation.** When the sponsored research office tells the provost that NDA turnaround dropped from nineteen days to four, perception forms before the quarterly meeting ever happens. Evidence travels in ways narrative cannot.

**Resources.** A before-and-after table showing that 44% of requests are now self-resolved – over a period when request volume grew 20% – turns the budget conversation into arithmetic.

**Standing.** Leadership stops treating legal as an administrative cost to be contained and starts treating it as a strategic function to be invested in. That shift doesn't happen because legal argued for it. It happens because legal demonstrated it.

**Self-perception.** When attorneys see that half of what used to reach them no longer does, they understand their own capacity differently. They stop feeling perpetually behind.

#### The baseline that cannot be skipped

Six months in, NDA turnaround is down to four days and the AI agent is handling 60% of policy questions. The team feels the difference daily. But no one recorded the before – so the quarterly update is still a description of improvement rather than

a demonstration of it. Before anything is built, pull the last two semesters of completed matters and record:

- Average turnaround on your three highest-volume request types
- Total request volume by department and matter type
- Estimated percentage of requests that required attorney involvement
- Number of requests arriving without enough information to act on
- Compliance deadlines met versus missed
- Attorney hours spent on work that did not require an attorney

### Three types of value to track

Volume reduction	Productivity uplift	Compliance & risk
<p>Reducing the number of requests that reach attorneys through self-service, AI-powered responses, and intelligent routing. Legal only touches what requires legal judgment.</p> <p><b>Metrics to track</b> Requests resolved without attorney involvement; self-service completion rate; deflection rate by matter type; share of requests arriving with complete information.</p>	<p>Enabling attorneys who do touch work to handle it faster and more consistently – through structured intake, AI-assisted review, and workflow automation.</p> <p><b>Metrics to track</b> Average turnaround by agreement type; time from first contact to attorney assignment; number of clarification rounds; attorney hours per matter type.</p>	<p>Reducing missed deadlines, inconsistent positions, and matters that fall through the cracks – the value type most compelling to leadership and least tracked without infrastructure to capture it.</p> <p><b>Metrics to track</b> Statutory and regulatory deadline compliance rate; research agreements executed before grant start; consistency of contract positions; matters escalated due to missed response windows.</p>

## The metrics that move leadership

In a budget conversation with a provost or CFO, the following metrics land. Others are useful internally. These change the room.

- **Grant activation timeline.** Research agreements executed before the sponsor's required start date. Connects legal directly to research revenue and faculty productivity.
- **Self-service rate by department.** Turns departmental stakeholders into advocates.
- **Outside counsel spend.** When internal capacity increases, routine work returns in-house. A line item the CFO already watches.
- **Compliance deadline performance.** Specific, named, before and after.
- **Faculty service level.** A proxy for the institution's ability to recruit and retain productive researchers.

# Driving Adoption in a University Environment

Building the front door is the technical challenge. Getting the institution to walk through it is the human one. Universities that operationalize AI treat adoption as a design problem, not a communication problem. They build workflows that are easier to use than the alternatives that already exist.

## 01 MEET PEOPLE WHERE THEY ALREADY WORK

Faculty submit requests by email because email works. Administrators use Teams because Teams is where their work lives. The front door has to come to the institution, not the other way around. The AI agent is accessible from Teams. The intake workflow can be triggered from an email. When the front door is embedded in existing channels, the question stops being “will people adopt the new system?” and becomes “did people notice that legal got better?” For unionized institutions, changes to how administrative work is performed may require advance consultation with labor relations – engage HR early.

## 02 DESIGN THE AGENT BEFORE IT GOES LIVE

The fear that stops most universities from deploying a policy AI agent is well founded. A generic agent drawing on broad training data can produce responses that are plausible, confident, and wrong – at scale, with no attorney aware. The right AI agent draws only on what legal builds into it: institutional policy documents, the team’s standard guidance, approved language for each scenario, and explicit fallback instructions for questions outside scope. Every response is tested and signed off in a controlled environment before go-live. The result is an agent legal can stand behind professionally – because legal built it, tested it, and knows exactly what it says.

### **03 KEEP THE HUMAN WHERE THE HUMAN BELONGS**

AI does not replace attorney judgment. It protects it – by ensuring attorney time is reserved for work that requires it. The agent handles the conflict of interest question being asked for the fourteenth time this semester. The workflow handles the NDA the counterparty was willing to sign as-is. The attorney handles what requires them. The hard lines – matter types where AI supports intake and nothing beyond – are built explicitly into the workflow, not left to case-by-case discretion. Faculty and staff who trust that legal has drawn those lines carefully are more likely to engage.

### **04 WIN EARLY AND WIN VISIBLY**

The first successful workflow is operational and political. Choose the first use case with that political dimension in mind: not necessarily the one that saves the most attorney time, but the one whose improvement is most visible to the people with the most influence over legal's resources. Research NDA turnaround is visible to the sponsored research office – which talks to the provost. Policy response time is visible to every department that has been waiting days. When the first result lands in someone else's quarterly update, it stops being a legal technology story and becomes an institutional improvement story.

### **05 POSITION LEGAL AS THE INSTITUTION'S AI AUTHORITY**

Universities are increasingly asking legal to advise on AI governance – acceptable use policies, data privacy frameworks, procurement standards for AI tools, risk assessment for AI-assisted decisions. These are legal questions, and they belong in the Office of Legal Affairs. The team that has built and operated its own AI-powered workflows is far better positioned to answer them credibly than a team observing AI from the outside. Operational experience creates institutional authority. The office that deploys AI thoughtfully and governs it carefully stops functioning as a service department and starts functioning as a strategic institutional resource.

## **The thread that runs through all five**

Teams that get adoption right are deliberate about every dimension of it – where the front door lives, what the AI agent says, where the human line is drawn, which result gets shared first, and how the team builds authority over time. From the outside, those teams look like they got lucky. It is rarely luck. It is design.

# Evaluating Vendors

The legal technology market has never been more crowded. Every platform promises AI. Every demo is impressive. For an Office of Legal Affairs evaluating under procurement rules that require defensible decision-making, the noise is exhausting – and the cost of choosing wrong is political as well as financial.

This is not a feature checklist. It is a set of questions that reveal whether a vendor understands the problem a university legal team is trying to solve, and whether their platform is configurable enough to solve it.

## Dimension 1 – Does the platform solve the right problem?

The market contains two fundamentally different categories of tool. The first manages work that has already arrived – matter management, contract lifecycle platforms, document repositories. They organize the internal workspace. They do not change the experience of the person submitting the request.

The second changes what arrives and how – a legal front door, an AI agent that resolves questions before they reach an attorney, workflow automation that routes work from first contact, and reporting that captures data as a byproduct of every interaction. Before evaluating any platform, ask which category it belongs to.

## Dimension 2 – Is the platform configurable for universities?

University legal work has specific characteristics a platform must accommodate through configuration – not custom development, not professional services that create IT dependencies. Ask every vendor:

- Can the platform handle the full universe of university legal demand, not just contracts? Policy questions, compliance queries, faculty research inquiries, student affairs, employment intake.
- Can the AI agent be configured for your institution’s policies and language? Source materials defined by legal, responses crafted and approved, scenarios tested in a

controlled environment before go-live, explicit fallback behavior for out-of-scope questions.

- Can legal own and update workflows without IT involvement after initial setup? IT handles the institutional gates. After that, legal makes changes directly – or it doesn't.
- Does the platform integrate with your actual systems – Workday, Banner, Quali, Cayuse, PeopleSoft? Ask what is native and what requires work.
- Does the platform meet data governance and security requirements? FERPA, SOC 2, state data residency, whether your data trains the vendor's models, and whether the vendor will sign your standard DPA.

### **Dimension 3 — Can you trust the partnership?**

A platform is not a product purchase. It is a relationship that will shape how the office operates for years. Evaluate:

- Does the vendor understand your environment without being taught? A vendor who asks about your research portfolio, FERPA obligations, and AI use policy before showing a demo has done the work.
- Is the vendor honest about what the platform cannot do? Vendors who acknowledge gaps and explain how they're addressed are telling you something important about how they'll behave after the contract is signed.
- What does implementation actually look like – and how long before the first workflow is live?
- Does the vendor's roadmap align with where you need to go? Ask for specifics. Reassurances are not commitments.

# Green flags and red flags

Green flags	Watch-outs
Asks about your institution before showing anything	Demo begins before they've asked what kind of institution you are
AI agent configurable by legal – source materials defined, responses approved, controlled testing, explicit fallback behavior	AI agent draws on broad training data and can't be limited to approved source materials
Understands the difference between a research NDA, data use agreement, and material transfer agreement	Cannot demonstrate a controlled testing environment where legal approves all responses pre-launch
Transparent about FERPA compliance and data handling	Vague about FERPA or data residency
Will sign your institution's standard DPA	"We'll figure it out" on DPA
IT involvement limited to institutional gates – legal owns everything after	Implementation model puts vendor or IT in control of ongoing workflow maintenance
Holds cooperative purchasing agreements your institution can leverage	Every university-specific question answered with "we can configure that" without demonstrating how
Realistic implementation timeline grounded in comparable deployments	Contract timeline significantly shorter than comparable implementations actually took

## **A word on procurement**

University procurement requirements exist for good reasons. Build them into your evaluation timeline from the start. Vendor evaluation and institutional approval should run in parallel, not in sequence.

Many legal technology vendors hold cooperative purchasing agreements through E&I Cooperative Services, Sourcewell, or NASPO ValuePoint. When a vendor holds a contract your institution can piggyback on, you skip the RFP entirely. Ask early.

Consider what the right platform does for procurement itself. When procurement officers have a self-service path for vendor agreements within pre-approved terms, they don't wait for an attorney. Legal solving its own intake problem and simplifying procurement for the rest of the institution at the same time may be the single most compelling benefit for a CFO or chief procurement officer.

# Implementation Guide for University Legal Teams

Most university legal teams are lean. No dedicated engineers, no large implementation budget, no months of runway. This section translates the framework into practical steps. The organizing principle: start with one workflow, build it fast, prove it works, use the result to fund what follows.

## Before you build anything

**Establish your baseline.** Without it, your first result is a story. With it, your first result is proof.

**Engage IT, security, and purchasing early** – at the start of evaluation, not after vendor selection. If the vendor holds a cooperative purchasing agreement your institution can use, that single conversation can compress procurement from months to weeks.

**Confirm the ownership model.** IT's role is security review, procurement approval, and SSO. Once cleared, legal builds, configures, and maintains everything directly.

**Define success before you build.** Not “faster NDA turnaround.” Eighteen days to four days. The metric you define before launch is the metric you report after.

## The implementation sequence

STEP

1

### IDENTIFY YOUR HIGHEST-VOLUME, MOST MEASURABLE WORKFLOW

Look at where legal time is being consumed and where data already exists or can be easily collected. Good candidates: research NDAs, standard vendor agreements, intake from a specific department.

STEP

2

### MAP THE CURRENT WORKFLOW IN DETAIL

Document every step before touching technology. Automation forces you to confront gaps and assumptions you've been working around. Be thorough.

STEP

3

### BUILD INTAKE FIRST, AGENT IN PARALLEL

Foundation work first – intake form, routing logic, matter management layer. Configure the AI agent in parallel, fed with institutional policy documents and standard guidance. Every scenario tested and approved before it goes live.

STEP

4

### PILOT WITH ONE DEPARTMENT BEFORE CAMPUS-WIDE

Brief the pilot department in advance. Launch. Monitor daily for the first two weeks. Refine before you expand.

STEP

5

### REPORT AT 30 DAYS, EXPAND FROM THERE

Pull the data. Compare to baseline. Share results with the sponsored research office, the provost, the CFO – framed as institutional improvement, not legal technology. The second workflow is faster to build than the first. The third faster still.

## Realistic implementation timescales

Workflow	Timeline post-procurement
Simple intake form with auto-routing and acknowledgment	<b>1–2 weeks</b>
NDA or standard agreement workflow with guided options and e-signature	<b>2–4 weeks</b>
AI agent for policy and compliance questions (initial configuration)	<b>3–6 weeks</b>
Full research agreement intake portal with matter management and reporting	<b>6–10 weeks</b>
Institutional procurement and IT approval	<b>4–16 weeks — begin in parallel</b>

**Practical tip**

The first workflow creates the time to build the next. Even an automation that saves five hours a week gives you back the bandwidth for something more complex. Start narrow. Prove the value. Reinvest the time you save.

# Your Quick-Start Checklist

Four phases. Each builds on the one before it. Do not move to Phase 2 until Phase 1 is complete. Every item is actionable by a lean team without dedicated implementation resources.

## PHASE 1 Diagnose

- Identify the three request types consuming the most attorney hours each semester
- Pull the last two semesters of completed matters and calculate average turnaround for each
- Estimate the percentage of requests that required attorney involvement vs. could have been self-served
- Map every channel through which legal requests currently arrive
- Identify which departments generate the most demand
- Document compliance deadlines that govern your highest-volume request types
- Review your institution's AI use policy for constraints on tool deployment
- Identify your IT security contact, privacy officer, and purchasing office
- Determine whether prospective vendors hold cooperative purchasing agreements
- Write down today's baseline – date it, keep it

## PHASE 2 Design

- Select one workflow as your first use case using the prioritization matrix
- Map the workflow step by step – every handoff, every point where time is lost
- Identify which steps require attorney judgment and which are administrative
- Define what information must be collected at intake so attorneys can begin without follow-up
- Identify the policy and compliance questions your office fields most frequently – the AI agent's starting inventory
- Gather the policy documents and standard guidance the agent will draw on, reviewed by legal
- Define your institution's standard positions, fallback clauses, and escalation triggers
- Confirm the integration path with any existing contract management system, or confirm native handling
- Document IT's role as limited to security review, procurement approval, and SSO
- Define success metrics with specific before-and-after targets
- Identify the pilot department and plan communication

### PHASE 3 Build and launch

- Complete IT security review, data governance assessment, and procurement approval – in parallel with design
- Build the intake workflow in a platform legal owns and can update directly
- Configure the AI agent using only approved source materials and legal-team-authored responses
- Test every scenario in a controlled environment – review every response, define escalation and fallback
- Confirm nothing goes live that the legal team has not reviewed and signed off on
- Configure the matter management layer so every submission is captured and visible to the general counsel
- Configure the reporting dashboard to capture the metrics defined in Phase 2
- Ensure the requestor experience includes acknowledgment & status visibility without follow-up
- Pilot with the identified department – monitor daily for the first two weeks
- Brief other departments before campus-wide launch
- Launch and monitor for the first thirty days

### PHASE 4 Measure, report, and expand

- Pull the 30-day data and compare to the baseline
- Build the before-and-after comparison in language tied to institutional outcomes
- Share results with the provost, CFO, and sponsored research office
- Document what the first workflow cost, how long it took, and what it delivered
- Identify the second use case using the prioritization matrix
- Extend the AI agent to the next most common question types
- Review AI agent outputs regularly – escalations and non-escalations alike
- Update workflows immediately when policies change – legal makes these changes directly
- Build the reporting layer to cover volume, deflection, attorney hours & compliance performance
- Present results to the board audit committee or relevant oversight body
- Evaluate Title IX intake routing, employment workflows & student affairs routing as next use cases
- Contribute operational experience to your institution's AI governance process


## **One final note**

Every item on this checklist is achievable by a lean university legal team without dedicated implementation resources, a large technology budget, or months of runway. What it requires is sequence. Discipline about starting narrow. Commitment to measuring from day one. And the willingness to share the first result in the right rooms before the momentum has a chance to dissipate.

The institutions that move fastest on this aren't the ones with the largest legal teams or the largest technology budgets. They are the ones where the general counsel decided that the cost of inaction was higher than the discomfort of building something new – and then built it deliberately, one workflow at a time.



**Checkbox** is the Legal Front Door trusted by in-house legal teams across the globe. From structured intake and intelligent triage to automated workflows, self-service portals, and cross-tool orchestration, Checkbox provides the infrastructure layer that makes every other investment in your legal tech stack work harder.



Trusted by over 100 enterprise organizations including **SAP**, **PepsiCo**, and multiple **Fortune 500 companies** to streamline legal intake, reduce administrative burden, and improve service delivery across the business.

Legal requests come from every direction.

Checkbox's AI Legal Front Door centralizes all legal work into a single, unified system of knowledge.



AI-powered intake across every channel



Self-service for routine low-value work



Intelligent triage and automated workflows



Real-time visibility into workload and risk