LegalOperator

THE RIGHT ANSWER. RIGHT NOW. BETTER OUTCOMES.

COMPANION TO EBOOK ONE

The Scorecard

Build Blueprint

How to build the page, on the systems you already have.

CRAIG RAEBURN · CO-FOUNDER, LEGALOPERATOR

BEFORE YOU START

A build guide, not a read-through

The ebook, *How to Build a Board-Level Legal Scorecard*, designed the one-page Scorecard your board sees. This is the how-to-build companion: it takes you from that designed page to a running one, on the data you actually have. It is tool-agnostic, and it meets you wherever you are today, including if you have no system that can assemble this yet.

1 **Part One: The Data Model Underneath the Page**
The small set of records every board line feeds from, one status map that makes every system speak the same language, and how to tag matters by category.

2 **Part Two: Build the Eight Lines**
Each board line as a build sheet: the inputs it reads, the math in plain arithmetic, how often to refresh it, who owns it, and how to set what good looks like.

3 **Part Three: Stand It Up on What You Have**
Three tracks to build the page depending on how connected your systems are, and how to move up without redoing the design.

Part Four, keeping it honest and the operating cadence, follows once the page is running.

The page you are building

If you arrived here without the ebook, this is the page in one paragraph. One page. Five questions the board and CEO are really asking. Eight numbers that answer them, each with what good looks like and the move when it is off. To the right of every line: a status (on track, watch, or off), a trend (which way it is moving), and a short read, a few plain sentences that interpret the number and connect it to the others. Software drafts the read from the numbers; a person edits it and owns it. Where the data is thin, the line carries a coverage or confidence flag instead of false precision.

- **Are we on top of the whole workload?** Work in vs. work cleared; Shadow demand.
- **Is the business unblocked, and are matters moving?** Time to first response; Stalled-matter rate.
- **Are we protected?** Critical-date coverage.
- **Are we growing capacity faster than headcount?** Self-serve rate; Load per lawyer.
- **Is spend under control?** Outside counsel spend vs budget, with forecast.

Two more terms. The front door is the stream of requests the business brings to legal, and the intake tool or inbox that catches them. Two of the eight lines, response time and self-serve rate, come straight from it. The data layer is the small model in Part One: the records, the status map, and the matter-category tag. Three lines wait on it.

PART ONE

The Data Model Underneath the Page

The ebook told you the data comes first. Here is the data model, and the fastest route through it.

The ebook left you with a designed page: eight lines, each mapped to where its number lives, honest about the gaps. Its data chapter told you what has to be true underneath that page. Sort your matters by the kind of work they are. Get every system to describe where a matter stands in the same words. Count the gaps on purpose instead of reporting them as zero. Those are the principles. This is where they become a worksheet.

Part One builds the data model the whole page sits on. It is small on purpose: a short list of records, one translation table, and your matters tagged by category. Get it right once and the parts ahead are mostly wiring. Part Two builds each of the eight lines on this model, field by field. Part Three shows three ways to stand it up depending on what you have today. Neither asks for anything you do not define here.

Want a page in two weeks? Read the record table below, then go straight to Path 1 in Part Two and stand up the three single-system lines. Come back here for the status map and the category tag before you start Path 2. That is the ebook's front-door advice, applied.

The small model every line feeds from

The eight board numbers look like a lot to maintain. They are not. They all draw from a handful of records with a few fields each. The last column shows where each record usually already lives, so you are connecting to systems you have, not starting from a blank page.

What you track	The few fields that matter	Where it usually lives	Board lines it feeds
Request	Received time; first human-touch time; how it resolved (self-serve, became a matter, referred); who asked	*Your intake or legal front-door tool; a shared inbox or ticketing queue if you have no intake system*	*Time to first response; Self-serve rate; Shadow demand (the reconciliation set)*
Matter	Open date; close date; native status; last-activity date; matter category; practice area; owner; assigned lawyer	*Matter management (with eBilling, your ELM)*	*Work in vs. work cleared; Stalled-matter rate; Load per lawyer*

What you track	The few fields that matter	Where it usually lives	Board lines it feeds
Matter-category tag	Project, portfolio, or advisory	*A field you add in matter management, or a column in the workbook*	*Scopes Work in vs. work cleared, and Stalled-matter, to the right work*
Status map	Native status to one of the five standard statuses	*A mapping you keep once, in the platform or a workbook tab*	*Work in vs. work cleared; Stalled-matter; every line that needs one language*
Owner	The one person accountable for the matter or the date	*Matter management; docketing or CLM for a date's owner*	*Critical-date coverage; stalled follow-up*
Critical date	Each date is its own record, and a matter can have several. Date; type; the matter; owner; condition (on track, at risk, or met); source system	*IP docketing; court dockets; CLM for contract renewals; a regulatory or entity calendar*	*Critical-date coverage*
Invoice and accrual	Invoiced to date and accrual, per matter; and the department outside-counsel budget for the year, ideally phased by month or quarter	*eBilling for invoices and accruals; the budget set by legal with finance*	*Outside counsel spend vs budget, with forecast*
Roster	Lawyer; whether the matter is run in-house or carried by an outside provider	*Matter management, or a team roster you keep*	*Load per lawyer*

That is the whole list. Most of these records already exist in the systems you run: intake, matter management and eBilling (together, ELM), your IP docketing, and your contract system (CLM). The build is not inventing records. What it does need is a place to keep all of this together and a reliable way to pull it, whether that is a platform that assembles it for you or a disciplined workbook. Part Three covers both. The work here is making the few fields that matter trustworthy and connecting them so the page can read across all of them. Two records carry most of the weight, and they come next: the status map and the matter-category tag.

WHERE A PLATFORM DOES THE HEAVY LIFTING

Keeping all of this together and up to date is the part teams dread doing by hand. It is also the part a platform is built for. LegalOperator captures this data as work comes in, across every channel the business uses to ask for help, and where it connects to the systems you already run it pulls the rest in rather than asking you to re-key it. Data is King, AI is Queen: the software does the first pass, a person stays in control.

- **A first pass you verify.** AI proposes, a person confirms, and the confirmed answer is protected from being overwritten. LegalOperator runs that pattern in production today on matter complexity and impact. The status map and the category tag are the next two fields it applies to.
- **An assistant on every matter.** Activity from email, Slack, and connected systems lands on the matter and is summarized as it arrives, so the last-activity date is real rather than remembered. Stalled-matter rate reads from that date instead of from a monthly reminder to update the record.

This does not replace the design work in this Blueprint. It removes the manual upkeep around it. The platform handles the capture and the maintenance. The judgment in the pages that follow, what to measure and how to read it, stays yours.

Fill in the status map

Your systems already track status. The problem is that each one tracks it in its own language, so no movement, or on hold, or done means something different in every tool. The status map is the one translated layer that fixes it. You map each native status your systems use into a small, fixed set of standard statuses, and you do it once.

Five standard statuses cover a whole department:

- **Intake.** It has arrived. Work has not started.
- **Active.** Someone on your team is moving it right now.
- **Waiting.** It is open, but the next move belongs to someone else: a court, a counterparty, a signature, a government office.
- **Closed.** It is done.
- **Dormant.** Open by design and resting between deadlines: a registered mark in maintenance, a standing agreement, an entity in good standing. Never mistaken for stuck, and never counted as cleared.

Here is that map filled in across the practice areas whose dialects differ most. Copy the pattern for the systems you run. Underneath, IP keeps saying office action and litigation keeps reading its docket. Nothing changes in the source systems. You are adding one layer on top that translates.

Native status (the dialect)	Standard status	What it means, and any flag
IP: patents and trademarks		
Docketed, not yet filed	**Intake**	*In the queue, work not started*
A filing or response being prepared, for example an office action response	**Active**	*Our team is moving it now*
Filed, awaiting examination or an office action	**Waiting**	*The ball is with the patent or trademark office*
Allowed, awaiting issue or registration	**Waiting**	*Waiting on the office, not on us*
Issued or registered, in maintenance	**Dormant**	*Open by design, resting between renewal deadlines. Never read as stuck*
Abandoned, lapsed, or expired	**Closed**	*Off the active list*
Litigation		
Pre-suit or demand received	**Intake**	*Arrived, not yet worked*
Active discovery, motions, or trial prep	**Active**	*Our team is moving it now*
Awaiting a ruling, a hearing date, or the other side	**Waiting**	*The ball is with the court or opposing counsel*
Stayed or on a long hold	**Waiting, or Dormant if the stay runs long**	*A short stay is waiting; a multi-year stay goes Dormant so it does not read as stuck*
Settled, dismissed, or judgment entered	**Closed**	*Off the active list*
Contracts		
Request received, not started	**Intake**	*In the intake queue*
Drafting, review, or negotiation	**Active**	*Our team is moving it now*
Out for signature or awaiting the counterparty	**Waiting**	*The ball is on someone else's desk*
Executed	**Closed**	*The drafting matter is done. The renewal or termination date becomes a critical-date record*
Standing agreement tracked for renewal	**Dormant**	*The obligation lives on with a date; it is not an open matter that stalled*
Expired or terminated	**Closed**	*Off the active list*
Employment		
Complaint or claim received	**Intake**	*Arrived, not yet worked*

Native status (the dialect)	Standard status	What it means, and any flag
Investigation underway	**Active**	*Our team is moving it now*
Awaiting the employee, an agency, or outside counsel	**Waiting**	*The ball is on someone else's desk*
Resolved or closed	**Closed**	*Off the active list*
Immigration and visas		
Case opened, gathering documents	**Intake**	*Arrived, not yet worked*
Petition or an evidence response being prepared	**Active**	*Our team is moving it now*
Filed, awaiting a decision or with an evidence request pending	**Waiting**	*The ball is with the government*
Approved, renewal window tracked	**Dormant**	*Standing status; the renewal before it lapses is a critical date, not a stall*
Denied or withdrawn	**Closed**	*Off the active list*
M&A and corporate deals		
Term sheet or letter of intent under discussion	**Active**	*Our team is moving it now*
Due diligence and negotiation underway	**Active**	*Our team is moving it now*
Signed, awaiting regulatory approval or closing conditions	**Waiting**	*The ball is with a regulator or the other side*
Closed	**Closed**	*The deal matter is done. Post-closing and integration dates become critical-date records*
Post-closing obligations tracked	**Dormant**	*Open by design until the obligations run out; not a stalled deal*
Corporate and entity management		
Entity in good standing	**Dormant**	*Open by design. Its annual filing and renewal dates are the critical dates to watch*
Formation, change, or dissolution in progress	**Active**	*A project inside the portfolio; our team is moving it now*
Awaiting a state or registry filing to process	**Waiting**	*The ball is with the state or registry*

ACTIVE OR WAITING: THE CALL THAT CARRIES THE MOST WEIGHT

Most native statuses map themselves. The one worth slowing down for is the line between active and waiting, because that is the line the stalled-matter number rides on. Use one test.

- **Whose desk is the next move on?** If the next action is ours and it is not happening, the matter is active and at risk of stalling. If the next action belongs to a court, a counterparty, a government office, or a signature, it is waiting, and waiting is not stalled.

Get this one call right and a portfolio matter resting between deadlines never gets chased, and a project the business is waiting on never hides behind the word open.

Tag your matters by category

The status map tells you how to read where a matter stands. It only works once every matter also carries a category, because category decides which lines a matter belongs to at all. A patent and a lawsuit are both open, but the moving-and-progress lines were built for the lawsuit, not the patent. So you tag your matters. One word on the name. Most systems already have a matter type, usually a sub-type under practice area, and someone worked hard on those. Category is a different field, one level up, and it leaves your matter types alone. The fear is a six-month classification project. You do not need one.

Start with the default rule. It sorts most of your matters on the first pass:

- **Portfolio.** Open by design, with no natural end, run on deadlines rather than to a close. Patents, trademarks, standing entities, ongoing compliance programs, and executed contracts you track for renewal.
- **Advisory or quick-turn.** Fast work, often a question or a simple review turned around in days. Some of it stays a logged request; some becomes a matter the moment it needs a tracked owner and a deliverable, like a question routed to outside counsel or a simple contract. Advisory is a matter category, not a way of saying not a matter.
- **Project.** Everything else. It begins, it ends, it runs on a clock. Contracts in negotiation, litigation, deals, investigations. This is the default bucket, and the one work in vs. work cleared, stalled-matter, and response time are really about.

Underneath all three sits one rule you have to hold: what makes something a matter at all. A quick question resolved the same day is a logged request. The moment work needs a tracked owner and a deliverable, or it goes to outside counsel, or it carries real risk, it becomes a matter, even a small one. Write that rule down and apply it the same way every time. The exact line matters less than holding it, because an inconsistent one quietly bends work in vs. work cleared, self-serve rate, and shadow demand all at once.

A FOURTH CATEGORY, IF YOU NEED IT: ADVICE

Some departments will want a separate category for advice: the guidance that never becomes a matter at all, or that lives in a standing matter kept open all year to collect it. The three categories above still cover most departments, and the ebook keeps to three on purpose. But if advice is a big share of your work, or your Legal Ops team is already asking, split it out. Treat it like portfolio for counting: it stays out of work in vs. work cleared and stalled-matter, so it never distorts either line. Add the label to the prompt below and the default rule still holds.

TAG EVERY MATTER IN AN AFTERNOON

You do not classify matters one at a time. You map in bulk, then hand-check only the buckets that mix.

- **Map by what your system already files.** Patent and trademark records go to portfolio, litigation to project, standard contract and NDA requests to project. Whole categories move at once.
- **Hand-check only the mixed buckets.** The categories where project and portfolio live side by side, mainly contracts and IP, are the only ones worth a closer look.
- **Default the leftovers to project and flag them.** A flagged default is visible and easy to fix later.

Then handle the edges. Four cases actually mix, and each has a clean rule:

- **The executed contract with a renewal.** The negotiation was a project and it closed. The standing agreement and its renewal date are a portfolio obligation and a critical-date record. One record, two lives: project until executed, then it leaves the moving lines and lives on as a date.
- **The portfolio matter that flares into active work.** A patent draws an office action, a trademark draws an opposition. It is still portfolio for counting, but it is active on the status map for the duration. Category and status are two different tags. Do not collapse them.
- **The advisory that quietly grew.** An advice request that turned into a real negotiation. Re-tag it project once it meets your matter rule. A periodic sweep catches these before they distort the count.
- **The matter with no obvious home.** Default it to project and flag it. An uncategorized matter silently drops out of every scoped line, which is worse than a default you can see and correct.

LET AI TAKE THE FIRST PASS AT TAGGING

You do not have to tag by hand either. If your matters live in a system, let it propose the category from what it already knows: the practice area, whether there is a close date, whether the work runs on deadlines. You confirm the edges.

No system yet? Export your matter list to a spreadsheet and run it through your company's AI tool with a prompt like this:

> *You are classifying legal matters by lifecycle category. For each row, label it Project (begins and ends, runs on a clock, like a negotiation, a lawsuit, or a deal), Portfolio (open by design with no natural end, run on deadlines, like a patent, a trademark, a standing entity, or a compliance program), or Advisory (a quick question or review, usually resolved in a few days). Use the practice area, whether there is a close date, and the description. Return the label and one short reason for each row.*

You verify the output, fix the edge cases above, and every matter is tagged in an afternoon with no classification project. The same move works for the status map: the tool suggests, you decide.

One discipline holds all of this together: category and status are two separate tags, set independently. Category rarely changes; a project stays a project. Status changes constantly; active this week, waiting next. Keep them apart and both stay honest. Collapse them and you are back to the open-and-inactive trap the data chapter warned about.

What you have after Part One

A short list of records with a few honest fields each, and where each one already lives. One translation layer that makes every system speak the same five words. Every matter tagged by the kind of work it is. That is the data model, and it is deliberately small. Everything on the board page is built from it, whether a platform assembles it for you or you hold it in a workbook. Part Two takes the eight lines one at a time and builds each on this model, in plain math, with the refresh cadence, the owner, and how you set what good looks like. Nothing there needs more than what you just defined.

PART TWO

Build the Eight Lines

Part One gave you the data model: the records, the status map, the category tag. Part Two builds each of the eight board lines on it, one at a time. Every line gets the same core fields: the inputs it reads, the math in plain arithmetic, how often you refresh it, who owns it, and how you set what good looks like, plus where AI takes work off your hands. Nothing here needs data you did not define in Part One.

The eight lines fall into three build paths, easiest first. Three come straight from one system, and two of those three are the front door the ebook told you to start at. Three pull from several sources, or from the status map, into one view. Two you cannot measure directly at all, so you estimate them with care. Build them in that order and you have a first page fast, then make it complete.

One word on what good looks like before we start. A threshold is not a benchmark you import from someone else. You set it from your own promise and your own baseline. Measure where you are now, agree the level the business needs, and set the line there. The Scorecard flags movement against your bar, not an outside number that may not fit your business, your industry, or your department.

Path 1: lines that come from one system

The fastest lines. Connect one source, define the one term that trips people up, and the line is live. If you do nothing else this month, do these three. They stand a first page up on their own, one you can show the CEO with the matter lines marked as pending.

PATH 1 Time to first response

How long from a request arriving to a person actually helping.

INPUTS	The request record's received time and first human-touch time. One system, your intake or front door.
MATH	For each request, first human-touch time minus received time. Report the median across the period, not the average, so a few slow outliers do not hide a good typical experience. Illustrative: median 1.2 days against a 1.0 day target.
DEFINE THE ONE TERM	First human-touch, so an automated acknowledgment does not count. The clock stops when a person touches the work and moves it forward, not when the system says “we got it.”

WHERE AI HELPS	A tool that captures the request the moment it arrives starts this clock without anyone keying a timestamp, and logs first human-touch the moment a person actually acts, not when a bot acknowledges receipt.
REFRESH	Weekly. Response time moves fast and is an early signal of intake strain.
OWNER	The intake owner, or the ops lead who runs the front door.
SET THE BAR	From the promise you make the business, your intake response-time promise if you have one. If you do not, measure your current median, pick the level the business needs, and set the line there.

PATH 1 SELF-SERVE RATE

How much demand gets resolved without a lawyer opening a matter.

INPUTS	How each request resolved (self-serve, became a matter, referred). One system, intake.
MATH	Requests resolved by self-serve divided by total requests, over the period. Illustrative: 38%. An AI-generated answer counts as self-serve when it resolves the request; it counts as a deferral, not self-serve, when it escalates to a lawyer.
WATCH FOR GAMING	Self-served has to mean resolved, not quietly dropped. Track the re-submission rate. If self-served requests keep coming back within a short window, they were deferred, not handled. And bound it: self-serve covers only the categories a lawyer has pre-approved with a playbook. Anything outside them routes to a person, every time.
HONEST GAP	The self-serve that never became a tracked request is invisible, and the better your self-serve the more it handles that way. Count what the system can see, and note that the true self-serve rate is at least this high.
WHERE AI HELPS	The more of your front door a tool actually captures, the smaller that blind spot gets. You are counting more of what happened, not just what an inbox logged.
REFRESH	Monthly.
OWNER	The intake owner.
SET THE BAR	Read this one as a direction. Good is rising over time against your own baseline. Set a floor, but watch the trend.

PATH 1 Outside counsel spend vs budget, with forecast

Are we on plan, and where will we land.

INPUTS	The invoice and accrual record from Part One: invoiced-to-date and accruals per matter, and the department outside-counsel budget for the year (not per-matter budgets), ideally phased by month or quarter. Mostly eBilling, with the budget set by legal and finance.
MATH	Invoiced to date plus accruals, divided by the department budget, for where you stand. For the forecast, project the run rate to year-end and add known future work. Illustrative: 68% used, forecast 4% over.
THE HARDEST INPUT	Accruals, work your firms have done but not yet billed. Leave them out and spend looks lower than it is and the forecast is wrong, which is exactly the number the CFO will hold you to. If firms do not report accruals reliably, estimate from matter stage and historical billing lag, and label it an estimate.
WATCH FLAG	Open matters more than a month old with outside counsel assigned but no invoices yet. Newer matters are normal, since most firms bill monthly, so an older one with no invoice is spend on the way that has not landed.
WHERE AI HELPS	The strongest accrual data comes straight from your firms. Make it their job: a vendor portal where firms submit their own accruals and invoices, end to end, puts the hardest input on this line at the source instead of in an internal guess. The check to add on top is AI reading for holes: a matter with recent activity or a submitted invoice but no matching accrual, and an incoming invoice that does not fit your billing guidelines. The firms and your team supply the numbers; the software flags what is missing. Illustrative read it can draft: "Spend is on plan overall, but litigation is trending eight points over its allocation while IP is under. If litigation holds this pace, forecast finishes at 4% over budget by year-end." A person edits and owns it before it reaches the board.
REFRESH	Monthly, in step with your billing cycle.
OWNER	Legal ops, or the finance partner to legal.
SET THE BAR	The department budget is the line, not per-matter budgets. Good is on or under plan with the forecast holding. Set the alert margin, for example any forecast more than three percent over plan.

PRODUCT PROOF: ACCRUALS FROM THE SOURCE

This is the line the CFO holds you to, and the input the whole line depends on is the one legal has always had to guess. LegalOperator moves it to the source. Law firms submit their own accruals, invoices, and timekeeper rates through the LegalOperator vendor portal, end to end, in production today. Invoices, line items, allocations, currency, and accruals all land in one ledger, so the "invoiced to date plus accruals" half of this line is real the month you turn it on. What you supply is the budget it is measured against.

Path 2: lines that wait on the data layer

These three wait on the Part One data layer, then light up together. Work in vs. work cleared and stalled-matter need the category tag and the status map, so a portfolio matter resting between deadlines never reads as stuck. Critical-date coverage needs dates pulled from several systems into one view. Do that data-layer work once and all three come online. This is the integration project the ebook told you to put your resources into.

PATH 2 WORK IN VS. WORK CLEARED

Are we keeping up with the work we can see.

INPUTS	Matters opened (work in) and matters moved to the standard status Closed (work cleared), over the period, in the same unit on both sides. Scoped to project and advisory matters with the category tag. Requests that never became a matter are counted on the front-door lines, not here, so nothing is counted twice.
MATH	Work cleared divided by work in, or simply the two counts side by side. Illustrative: 104 in, 96 cleared, slightly behind. Above 100% you are catching up; below, the backlog grows.
SCOPE WITH THE CATEGORY TAG	Leave portfolio matters out. A trademark portfolio is open for years by design; count it as backlog and the number climbs forever and tells you nothing. The Part One category tag does this for you.
THE TRICKY TERM	A clean, consistent definition of cleared across every practice area. That is exactly what the status map gives you: Closed means the same thing whether it came from litigation, IP, or contracts.
WHERE AI HELPS	Once the status map and category tag are set, a platform reads Closed events the same way across every system, so this line refreshes on its own instead of a count pulled by hand each month.
REFRESH	Monthly.

OWNER	Legal ops.
SET THE BAR	Good is work in and work out roughly matched over time, measured against your own trailing baseline, not a fixed ratio.

PATH 2 Stalled-matter rate

How much of our open project work has gone quiet.

INPUTS	The matter record's standard status and last-activity date. Scoped to project matters with the category tag.
MATH	Project matters in standard status Active with no activity for longer than your threshold, divided by all open project matters. Illustrative: 9%. Waiting matters are excluded, because the ball is on someone else's desk.
WHY THE STATUS MAP MATTERS	This is why active and waiting had to be separated in Part One. A matter waiting on a court is not stalled. A project matter that should be moving and has not is. Portfolio matters resting between deadlines never count.
SET THE NO-MOVEMENT THRESHOLD	Pick the number of days that means quiet for your work, for example 30. Set it once and apply it the same way everywhere.
WHERE AI HELPS	Activity that lands on a matter, from email, Slack, tasks, and invoices posted against it, refreshes the last-activity date as it arrives, so the stall shows up at the next refresh instead of the next time someone remembers to update the record. Illustrative read it can draft: “9% stalled, up from 6% last quarter. Three matters account for the rise, all in commercial leasing, where two lawyers are out.” A person edits and owns the read before it reaches the board.
REFRESH	Monthly, or weekly if stalls are a live problem.
OWNER	The practice leads own their own stalled lists; ops owns the board number.

PATH 2 Critical-date coverage

Are the deadlines that can really hurt us owned and on track.

INPUTS	The critical-date record from Part One. Each date is its own record, and a matter can carry several. Date, owner, and condition (on track, at risk, or met), pulled from every system a date hides in: IP docketing, court dockets, CLM for contract renewals, and your regulatory or entity calendar.

MATH	Two numbers. The share of critical dates with a named owner (coverage), and the count in at-risk condition. The board line surfaces the nearest at-risk dates; the record holds them all. Illustrative: 100% owned, 2 at risk. Zero missed is the standard; the at-risk count is what you work down.
THE HARD PART, AND BE HONEST ABOUT IT	This line takes more plumbing than any other. Your dates do not live in one place. Until they are in one view, show coverage as the share of dates you can see, and raise it as each source comes in. Saying so on the page is a strength, not a weakness.
CAPTURING WHAT OUTSIDE COUNSEL HOLDS	When a firm runs a matter, the dates often live with them, not in your systems. Make reporting key dates part of how you expect firms to manage the matter, and pull the rest from your docketing, CLM, and calendars. This is the input teams most often miss.
WHERE AI HELPS	AI can read new contracts, filings, and the updates your firms send, and pull the renewal, response, and maintenance dates into the one view, so the plumbing is less manual. A person still confirms the date and its owner.
REFRESH	Weekly. A missed critical date is the one board number where a single miss is material.
OWNER	A single named owner for the consolidated view, even though each source has its own keeper.
SET THE BAR	There is no soft target. Every critical date owned, zero missed, at-risk worked to the floor.

Path 3: lines you estimate, not measure

Two lines have no clean feed. You cannot pull a report for work that never arrived, and legal does not track its own hours. So you build a defensible estimate, report a direction with a confidence flag, and improve it over time. Reported with its flag showing, an estimate belongs on the board page. Dressed up as precise, it is the fastest way to lose the room.

PATH 3 Shadow demand

How much legal work is happening without legal.

INPUTS	No system of record. You triangulate from traces: contracts executed in the business or procurement system with no matching legal matter; purchases over a threshold with no legal touch; signed NDAs or order forms never logged at intake; and a short anonymous survey of business teams asking how often they handled something legal without involving legal, run the same way each period so the trend is the signal.
MATH	Not a decimal. A reconciliation you run each period: count the traces, compare to prior periods, and report the direction. Illustrative: trending worse, low confidence. Trending worse is enough to act on.
HONEST FLAG	Report it as a direction with a confidence flag, and never let the read claim more than the traces support. It earns its place because it is the truest read on whether the business trusts the function. A department can look controlled on every other line and still be quietly bypassed.
WHERE AI HELPS	Point AI at the business system's executed contracts and your matter list and have it flag the agreements with no matching matter. It surfaces the candidates; a person confirms which are real shadow work.
REFRESH	Quarterly. It moves slowly and the reconciliation takes effort.
OWNER	Legal ops, with the business relationships to run the pulse.

PATH 3 Load per lawyer

How much is each lawyer carrying.

INPUTS	The matter record and the roster, both defined in Part One: the matter carries its assigned lawyer and standard status, the roster carries who is in-house. Count active matters per lawyer, split by work run in-house and work carried by an outside provider.
MATH	Active matters per lawyer, counted internally. Illustrative: about 22 each. Weight it if your data lets you, since 400 quiet patent matters are not the load of a dozen active lawsuits, but a plain count is a fair start.
HONEST FLAG	A proxy, labeled as one on the page. Legal does not track its own time, so this stands in for effort, not hours. Read it as a trend and a flag for who is underwater, not a precise workload.
THE REFINEMENT	Separate the matters your team runs from the ones an outside provider carries. A matter you supervise is not the same weight as one you work yourself.

WHERE AI HELPS	A platform that carries the assigned lawyer and status on every matter keeps this count current without a monthly tally. The in-house versus outside-carried split is one roster field you add, and once it is there the split reports itself.
REFRESH	Monthly.
OWNER	The GC or practice leads, for the staffing conversation it feeds.
SET THE BAR	There is no universal right number. Set your own band from a period when the team felt appropriately loaded, and watch for a sharp rise, the early signal of the hire-versus-tech decision.

What you have after Part Two

Eight lines, each with its inputs, its math, its cadence, and an owner. Build them in the three waves and the page fills in the same order: the single-system feeds give you a first page in a week or two, the status-map lines complete the picture once the Part One data work is done, and the two estimates round it out with their flags showing. Part Three takes these same eight lines and shows how to stand them up on what you actually have today, whether that is an integrated platform, a set of systems that do not talk, or a spreadsheet and some discipline.

PART THREE

Stand It Up on What You Have

Parts One and Two gave you the model and the eight lines. This part is about starting from where you actually are. There are three ways to stand the page up, and which one is yours depends on one thing: how connected your systems are today. The design does not change between them. The eight lines, the status map, the category tag, and the thresholds are the same in all three. Only how each line gets fed is different.

Pick your track by what you have now, and know you can move up later without redrawing anything. That is the point of having designed the page first.

Track	You are here if	How the page gets fed
Track A	Your systems already connect, or a platform assembles the page for you	Live feeds. Map statuses and categories once, then the page fills itself
Track B	You have the systems (matter management, eBilling, docketing, CLM) but they do not talk	Scheduled exports into one light reporting layer you own
Track C	You have no system that can assemble this, which is most teams starting out	One disciplined workbook, filled on a fixed monthly routine

Every track produces a board-ready page. They differ in how much effort stays manual and how fast the page refreshes, not in what the board sees. Start where you are.

Track A: an integrated platform, or systems that already talk

TRACK A the fastest path

Your sources connect, or a platform assembles the page from them. Your job is design and confirmation, not data entry.

- **Build it:** connect the feeds you have, intake, matter management, eBilling, docketing, and CLM. Map each system's native status to the standard status once, tag matters by category once, and set the thresholds from Part Two. Then the page fills itself and refreshes on your cadence.
- **Fill it in the three waves:** the single-system feeds light up first (response, self-serve, spend), the status-map lines follow once the mapping and tagging are

done (work in vs. work cleared, stalled-matter, critical-date coverage), and the two estimates round it out.

- **Where AI helps:** it takes the first pass on the busywork. It proposes the status mappings and the category tags for a person to confirm, keeps the last-activity date real from the activity that lands on each matter, matches executed contracts to matters for shadow demand, and pulls dates out of filings and firm updates. And for each of the eight lines, it drafts the first read, the plain-English sentence that says why the number moved and what to do about it, so the page assesses instead of just reporting. A person edits and owns that read before it reaches the board.

> WHAT LEGALOPERATOR DOES TODAY
>
> The front door is where LegalOperator is strongest, and it is where the ebook told you to start. Requests arrive by web, Slack, email, and Zoom and land in one place with their origin recorded, so the front-door lines start from real timestamps rather than re-keyed ones. Time to first response reports a true median. Self-serve rate reports the full three-way path: resolved by self-serve, became a matter, or referred out. Spend runs on LegalOperator's own eBilling, with firms submitting accruals and invoices through the vendor portal. And the page itself is built: a locked monthly snapshot, a print-ready one page, a trend on every figure, a read drafted on every row, and a chat tool to ask questions of it. Matter management, eBilling from other systems, and the rest connect as those integrations mature. You are adding a layer, not replacing what you run.

- **Where it must not overreach:** AI cannot invent data you never collected. It proposes; a person confirms. The read gives a direction when the data is thin and never claims more than the data supports, and the coverage and confidence flags stay on the page. Data is King, AI is Queen.
- **Stay integration-friendly:** you are adding a layer over what you already run, not replacing it. Connect what you have, in the order it pays off: the front door first, because that is where the fastest lines live, then spend, then the systems that feed the data layer.

Who owns it: legal ops owns the design and the mappings; the platform owns the plumbing. That split is the whole appeal of Track A.

Track B: systems that do not talk

TRACK B most mid-size departments

You have the systems, but they do not connect. In my experience this is most mid-size departments. The page is assembled by a light, disciplined process instead of a live integration, and it is a genuinely good place to run from.

- **Build it:** set a scheduled export from each system on a fixed cadence. Land those exports in one light reporting layer that holds the Part One model: a workbook, a reporting database, or a BI tool. Apply the status map and the category tag in that layer, and let the page read from it.
- **Hold the discipline:** same reports, same day each period, same steps, written down so anyone can run it. One owner for the layer. Mark any line you cannot yet trust.
- **Where AI helps:** it can do the joining that used to be manual, matching executed contracts to matters for shadow demand, matching invoices to accruals, mapping and classifying straight from the exports, and drafting the read. You verify the output.
- **It already migrates:** as any source gets connected, its live feed replaces its scheduled export. The reporting layer and the page do not change. You are on the path to Track A without a rebuild.

Who owns it: legal ops owns the reporting layer and the pull routine. Track B is Track C with the exports automated and Track A with the last connections still to make.

Track C: no system that can assemble it yet

TRACK C a first-class path, not a consolation

No platform, systems that cannot assemble this, maybe just intake in a shared inbox. This is where most teams start, and it is a real path. You build the page by hand, on purpose, and you make it trustworthy by treating one workbook like a system instead of a one-off.

The workbook is the Part One model made physical: a tab for each record, and one tab that rolls up to the page. The design you already drew is the workbook. Here is the layout.

Tab	What it holds	How you fill it each month	Feeds
Requests	Received time, first human-touch, how it resolved, who asked	From your intake tool, or a log of the shared inbox	*Response; Self-serve; Shadow demand*
Matters	Open and close dates, native status, last-activity, category, practice area, owner, assigned lawyer	An export or a kept list from matter management	*Work in vs. work cleared; Stalled-matter; Load*

Tab	What it holds	How you fill it each month	Feeds
Status map	Each native status mapped to one of the five standard statuses	Built once; updated only when a new status appears	*Every line that needs one language*
Critical dates	One row per date: date, type, matter, owner, condition, source	From docketing, court dockets, CLM, calendars, and firm reports	*Critical-date coverage*
Spend	Invoiced-to-date and accrual per matter; the department budget, phased	From eBilling or firm invoices; budget entered once	*Spend vs budget*
Roster	Each lawyer, and in-house versus outside-provider carry	From your team list	*Load per lawyer*
Scorecard	The one page: eight lines, status, trend, and the read	Formulas pull from the tabs above	*The board page itself*

That table is the whole build for Track C. Each tab is one record from Part One; the Scorecard tab is the page, fed by formulas that read the others. What makes it trustworthy is not the tool, it is the routine around it.

- **One owner, one cadence:** the page is only as current as its slowest pull. Pick a rhythm, hold it, and put one person in charge of it.
- **A written pull routine:** the same reports from each source, the same day, the same steps, so it comes out the same no matter who runs it.
- **Version control and honesty flags:** date each refresh and keep the prior one. Flag any line you cannot fully feed with a coverage note or a confidence label, so the board knows exactly how far to trust it.
- **A QA pass before it ships:** a quick check for the bad-data traps before the board sees it. Part Four is that checklist.

YOU DO NOT NEED A PLATFORM TO GET AI'S HELP

Track C is manual, but it is not cut off from AI. You do not need a system to let the software take the first pass.

- **Export a tab, run a prompt, paste back the verified result.** Classify matters by category, map native statuses to standard ones, or match executed contracts to matters for shadow demand. The example prompt in Part One works here unchanged.

The software does the first pass; you confirm it before it goes in the workbook.

This gets you there. What a platform buys back is not having to run this loop yourself every month.

Be honest about the trade. A workbook page is real if you treat it like a system, and fragile if you treat it like a one-off. It takes more manual effort than Track A or B. The payoff is that you ship a board page you can trust in weeks, on the data you have, instead of waiting for systems you do not have yet.

Move up without redoing the design

This is the idea that ties the three tracks together, and the reason the ebook made you design the page before building it. You design once. Moving from one track to the next never means redrawing the page.

A feed replaces a tab, not the page. When intake gets a real tool, its export replaces the Requests tab. When eBilling connects, it replaces the Spend tab. When a platform takes in your matters, it replaces the Matters and Status-map tabs. Each swap changes how one line is fed. It never changes what the line is, where it sits, or how the board reads it.

YOU ARE ALLOWED TO BE HALF-AUTOMATED

You do not migrate all at once, and you do not have to. A single page can mix live feeds and hand-pulled tabs while you connect one source at a time.

- **Mark which lines are live and which are still pulled by hand.** A line fed by a real feed and a line filled from a monthly export can sit side by side, each honest about how current it is.

Progressive, not perfect. A page that is half automated and fully trusted beats a fully automated page you are still waiting to build.

So the path is C to B to A, and you can stop wherever the effort and the value balance for your team. Track C stands the page up by hand. Track B automates the pulls into one layer. Track A connects the feeds so the page refreshes itself. Same page the whole way, gaining automation, never losing the design.

What you have after Part Three

A way to stand the page up no matter where you start, and a path to improve it without starting over. If you have an integrated platform, connect and confirm. If your systems do not talk, export into one layer you own. If you have no system yet, build the workbook and run it like one. Then move up when it pays to, a feed at a time.

SEE WHAT'S ALREADY BUILT FOR YOU

If you got this far by hand, on a workbook, on exports, on Track B or C, you now know exactly where the effort went.

Email hello@legaloperator.ai and tell us which track you're on. Two of the eight lines, time to first response and self-serve rate, fill in from your intake data today with no build on your end, and spend follows the month your invoices and accruals run through it. We'll show you those live, and exactly what it takes to light up the rest.

Part Four keeps the page honest and alive: the checks that catch bad data before the board does, the operating cadence that keeps it current, and the signal that you are ready for the reports that sit beneath the board page.