

FOR ENTERPRISE AND MULTI-OFFICE FIRMS

Governance, not just capability.

Granular access across offices, approvals that follow the organisation chart, firm-wide process campaigns, and a platform your own engineers can build against in both directions.

What this is

Two categories of legal software usually sit side by side and never meet. Practice management holds your matters, time and billing but knows nothing about the law. Legal AI answers questions about the law but knows nothing about your matters.

ModulawAI is both. The assistant that researches a point of law is the same assistant that can see the matter, the deadline, the documents on file and what has already been billed. That is the whole design.

Research across eight jurisdictions

Most legal AI is trained on one body of law. ModulawAI routes eight jurisdictions, each to its own source rather than a single blended corpus. Seven are searchable today.

JURISDICTION	CASE LAW SOURCE	DEPTH
Nigeria	ModulawAI's own case law and statutes database	Deepest coverage
Ghana	GHALII	Growing
Kenya	Kenya Law	Growing
South Africa	SAFLII	Unavailable, see below
United Kingdom	BAILII	Broad
United States	CourtListener	Broad, federal and state
Canada	CanLII	Broad
European Union	EUR-Lex	Good

South Africa is the one exception right now. SAFLII has put its search behind a bot challenge our retrieval cannot pass, so South African search is unavailable until we work around it. It is listed because the jurisdiction is selectable and the fix is ours to make, not because it works today.

What these sources are, and why it matters

Apart from our own Nigerian database, these are the national and regional legal information institutes plus two open APIs. None is a bespoke integration built for one country. SAFLII is the Southern African institute rather than a South African one, and BAILII and CourtListener each carry far more courts than we currently query.

So if you practise somewhere not in the table, that is a routing gap rather than a missing source. Adding a jurisdiction usually means pointing an existing adapter at the right court databases. Ask us, because the answer is often a short lead time rather than a no.

Why this fits a firm at scale

At your size the capability question is largely settled: any credible vendor can summarise a document. The questions that decide a rollout are different, and they are about control, governance and whether the thing will still be yours in three years.

Access that survives an audit

Documents and folders grant to an individual, a role, a whole department or the entire workspace, at viewer, commenter, editor or manager level. Reviews share to a workspace, department, role, job title or named people. Share links carry expiry, password and use limits, and only a SHA-256 hash of the token is stored, so a database compromise cannot reconstruct live links.

Approvals that follow the organisation chart

Department heads resolve as approvers, so policy is expressed structurally rather than as a list of names that goes stale. Approval policy is per matter, and requisitions route through a sequential signed chain that prints with signature and stamp blocks.

Firm-wide processes, run as campaigns

A campaign launches a process across the whole firm, a role, a department or a named list, spawning one run per person and reporting across all of them. Appraisal cycles, annual declarations and compliance attestations stop being a spreadsheet somebody chases.

Not a closed box

ModulawAI is programmable in both directions. Your engineers can drive it from their own tooling over MCP or the REST API, and the assistant can reach out to MCP servers you run. Section on developers is written for them and does not simplify.

Tabular review

At scale the review question is not whether extraction works but whether its output is governable: who ran it, who can see it, and whether each answer can be traced.

A tabular review is a spreadsheet over a set of documents. Rows are your documents. Columns are questions in plain English. Every cell is an answer the assistant extracted from that specific document, with the quote it came from and the page it sits on.

What it looks like for a large review

COLUMN	THE QUESTION YOU TYPE
Document type	What kind of document is this?
Parties and counterparty	Who are the parties?
Governing law	What is the governing law and jurisdiction?
Change of control	Does a change of control trigger termination or consent? Quote it verbatim.
Liability cap	What is the cap on liability?
Key dates	What dates does this create or record?
Escalate	Is there anything here that should be escalated to the responsible partner?

Because every cell carries a verified quote and a page, the grid is defensible under review rather than an assertion. Sharing is segmented, so the deal team sees the grid and the firm does not.

The part that matters for legal work

Extraction at scale is worthless if you cannot trust it. Two things address that directly.

- **Every quote is checked against the source.** After the model answers, the quoted text is searched for in the passages it was given. If the quote cannot be found, the cell's confidence is downgraded and it is flagged. A fabricated citation cannot quietly sit in your grid looking like the others.
- **Every cell carries its provenance.** The quote and the page number travel with the answer, so checking a cell means reading one paragraph rather than reopening the document.

Cells carry a confidence of high, medium or low, and a reviewer can override any value and mark it reviewed. The grid is a starting point that a person signs off, not an answer you are asked to take on faith.

Column types

A column can be free text, verbatim (the exact words from the document, nothing paraphrased), a date, a currency amount, a yes or no, or a list. The type constrains what the model may return, which is what makes a column sortable and comparable rather than six different phrasings of the same fact.

Running it without burning the budget

Sample mode runs the first few documents so you can see whether your questions are the right questions before committing to the full set. Credit cost is estimated up front, a run can be cancelled while in flight, and a single column can be re-run on its own after you reword it. That last one matters more than it sounds: the first draft of a question is rarely the right one.

Where the documents come from

Upload them directly, or pull from documents already on a matter, in the workspace library, or on a project. Nothing is duplicated: a review points at the document you already have.

Sharing a review

A review belongs to whoever made it and can be shared to the whole workspace, a department, a role, a job title, or named people. That is finer than most tools allow, and it is the difference between a partner sharing a due diligence grid with the deal team and sharing it with the entire firm.

For your lawyers

Research that shows its working

Ask a question in plain English. The assistant searches the case law database for the jurisdiction you selected, shows the steps it took, and cites the authorities behind its answer. Each authority appears as a card carrying the case name, a plain English summary of the holding, the citation, the court and the date.

- **Provenance is labelled.** Anything retrieved from the database is tagged as such. Anything the model produced from training data alone is tagged unverified.
- **Verification re-checks the source.** For US and Canadian authorities it re-queries CourtListener and CanLII. For EU it checks EUR-Lex. For UK, Ghana and Kenya it re-checks the source legal information institute. Nigerian cases are confirmed against our own database.

One distinction worth understanding. For a Nigerian authority, "verified" means we hold that case in our records. It is not an independent re-check against an external report, which is what it means for the other jurisdictions. That difference matters to anyone relying on the badge, so we document it rather than flatten it.

Contradiction checking

Ask whether an authority has been contradicted, distinguished or criticised by later decisions. The kind of check that is easy to skip under deadline.

Citation networks

The knowledge graph shows how cases relate: which authorities cite which, which judges rely on what, and how much precedential weight a decision carries.

Drafting

The assistant produces long form documents in a side by side editor called the Canvas: opinions, contracts, court documents, memoranda, disclosure schedules, formal letters and reports. Edit directly or ask for revisions in plain English, and the parts you did not ask about are left alone.

The sequence that produces usable drafts is research first, then draft. Ask for the authorities, then ask for the document. It writes with the cases it just found rather than inventing citations.

Export to Word or PDF, or save the draft against a matter as a note or a task.

House style, taught once

Skills are instruction bundles the assistant applies automatically when the situation matches. Written once at firm level, they apply to everyone. Twenty one ship as standard, and these are the ones that bear on your work:

AREA	BUILT IN SKILLS
Court documents	Brief of argument, motion drafting, affidavit drafting
Contract work	Contract review checklist, NDA drafting
Engagement	Engagement letter and conflict check
Tax and corporate	Transaction tax, CAC incorporation
Quantum and costs	Damages quantum, bill of costs, employment computations
Deadlines and citation	Limitation periods, court timetables, Nigerian citation format, IRAC case briefing

At your scale the built in skills matter less than the firm-written ones. A skill written once at firm level shapes every draft in every office, which is the only mechanism we know of that makes house style hold across a distributed practice.

Working with documents you already have

Attach a document from your computer, from a matter, or from connected cloud storage and ask about it. Summarise it, pull out every termination and notice provision with clause numbers, or compare two drafts and report only the substantive differences. Scanned documents and photographs are put through text recognition first, so a photographed court order can be read.

Inside Microsoft Word

Most drafting happens in Word, so there is an add-in. Research, redlining and your firm's house style, inside the document you are already working in.

Voice

A live voice mode for when typing is the wrong tool. You speak, it answers aloud, and the exchange is written into your chat history. Better for thinking out loud than for precise citations.

Your document estate

Most firms have precedent scattered across cloud storage, personal drives and individual memories. Three capabilities address that.

Bring it across in bulk

Import folders wholesale from OneDrive, SharePoint, Google Drive or Zoho WorkDrive. Documents arrive searchable and available to the assistant as precedent, rather than sitting in storage nobody browses.

Access control that matches how a firm works

The library behaves the way a modern drive behaves, which is to say access is granular rather than all or nothing.

YOU CAN GRANT TO	AT THE LEVEL OF
An individual member	Viewer
Everyone holding a given role	Commenter
A whole department or practice group	Editor
The entire workspace	Manager
A guest outside the firm, or a link	

Grants apply to a document, a folder, or a case document. Share links carry an optional expiry, an optional password, a maximum number of uses and revocation. The raw token is shown once and never stored: only a SHA-256 hash is persisted, so a database compromise cannot reconstruct live links.

Versions, and comparing them

Upload a revised file under the same name into the same folder and it groups with the original as another version of one logical document. The row shows how many versions exist, history expands in place, and any two versions open side by side to see what changed. Office formats open for editing in Office Online from the same menu.

Running the firm

Each matter opens into a workspace with fifteen tabs, grouped by what they are for.

GROUP	TABS
The matter	Details, Tasks, Notes, Documents, Calendar, Timeline
Money	Expenses, Billable Time, Invoicing, Purchase Orders
Process	Forms, Processes, Approval Policy, Pending Approvals
The client	Communication

Time and expenses recorded against a matter feed invoicing directly, so a bill is assembled from what was logged rather than reconstructed at month end.

Approvals

Approval policy is set per matter: what needs signing off and by whom. Department heads resolve as approvers, so the policy is structural rather than a list of names, and it survives promotions, departures and reorganisations. Pending approvals is the queue, and requisitions route through a sequential signed chain.

Clients

Client records carry custom fields your firm defines, lifecycle stages your firm configures, and a conflict check status of not required, pending, cleared or conflicted. That last field answers "has this been cleared?" without anyone asking around. Records import from a spreadsheet or sync from Zoho CRM.

The client portal

Clients get a private space to follow their matter, exchange documents and message your team. Sign in is passwordless: a secure link, valid for fifteen minutes by default and usable once, both configurable. No password for a client to lose, and none for you to reset.

Internal collaboration

Channels, direct messages, threads, and audio and video calls with screen sharing and recording. The assistant can be mentioned in any channel, so research happens in front of the team rather than in one person's private chat.

Unattended automation

Workflows run on a manual trigger, a schedule, or one of six workspace events: case created, document uploaded, task completed, deadline approaching, case status changed, intake submitted. Eighteen templates ship as standard. Workflows never wait on a person. When a person must act, that is a process, and processes are next.

Forms and processes

At scale, procedure that is not written down is procedure that differs by office. This is the layer that fixes that, and it is also how you run anything firm-wide.

Forms

Three levels, deliberately separated, because the same questionnaire usually needs deploying in more than one place.

LEVEL	WHAT IT IS
Template	The definition. Built in a drag and drop builder, or written by the assistant. Versioned, and scopable to the whole firm or to one department.
Form	The template deployed somewhere: the workspace forms area, every matter automatically, only named matters, a project, a department, or a public link.
Entries	What people submit. Each entry pins the template version it was filled under, so an old entry still renders as it was asked.

Twenty three field types, including currency, date and time ranges, person pickers, person sign-off, file lists, attachments with a remark, repeating groups, matrices with roll ups and computed fields. A form set to public gets its own link and needs no login, so client intake and conflict declarations can be collected from outside the firm.

Who sees the answers is set per deployed form: only the submitter, the matter's team, named reviewers or roles, or workspace administrators. An entry can require a signature, route into the approval engine, carry a deadline, and render out into a document filed against the matter.

Processes

A process is a sequence of steps where people fill forms, review and sign off, running across people and across time. The engine is durable and resumable: when a run reaches a human step it stops and waits, and it survives a server restart while waiting.

- **The assignee is resolved, not hardcoded.** The person who started it, the person it is about, a named user, anyone holding a role, the subject's manager, a field on the matter, or an external person.
- **Fan out.** One step, the same form, many people. Advance when everyone has filed, or when the first does.
- **Branching.** Approve and reject lead to different steps, and steps can carry conditions.
- **Deadlines per step,** with reminders to whoever is holding it up.

- **Campaigns** launch a process across the whole firm, a role, a department or a named list, one run per person, with a report across all of them.

Two engines, on purpose. Workflows are unattended and finish in one pass. Processes wait on people, branch, fan out and resume. If you are automating a notification, that is a workflow. If a partner has to sign something, that is a process.

Firm tools

A law firm is also a business, and the business half is usually the part software forgets. These sit at firm level rather than on a matter.

TOOL	WHAT IT HOLDS
Clause library	Pre-approved clauses by category and jurisdiction, each marked preferred, fallback or aggressive, and tagged as favouring your side, the counterparty or neutral.
Playbooks	Reusable emails, agreements, documents, prompts, checklists, letters, notices, phone scripts, meeting agendas, fee schedules and retainers, available from the chat composer.
Vendors and payables	Vendor records and their bills, with purchase orders as the commitment issued before an invoice arrives.
Requisitions	Multi line payment requests routed through a sequential signed approval chain and printed to PDF with signature and stamp blocks.
Leave	Requests, allocations, a team calendar and approvals.
Signatures	A saved e-signature per person, used on documents and approvals.

Departments, and why they matter for approvals

A department is typed: practice, finance, HR, IT, marketing, admin or committee. Practice groups are the only kind that reach into matter assignment. Each department can have a head, and that head is a resolvable approver, so approval routing follows the organisation chart rather than a list of names that goes stale when someone is promoted.

Firm tools, financials, reports and AI workflows sit on Professional and above.

For your developers

ModulawAI is programmable in both directions, which is what lets a firm treat this as infrastructure rather than an application it is locked into.

Inbound: assistants into your workspace

ModulawAI runs a Model Context Protocol server. Point Claude, or any MCP capable client, at a workspace and it works with real matters. Twelve tool families are exposed: cases, tasks, clients, intake, documents, notes, calendar, financials, forms, processes, workspace admin and legal research.

- **OAuth 2.1** authorisation server with discovery documents. No shared secret pasted into a third party tool.
- **Device authorisation grant** (RFC 8628) for command line clients that cannot open a browser.
- **Tokens are workspace bound.** A token issued for one workspace cannot reach another, and scope enforcement fails closed.
- **Confirmation gate** on destructive operations, and revocation takes effect immediately.

Outbound: your tools into ModulawAI

The same protocol runs the other way. Connect an external MCP server by URL from the chat composer and its tools become available to the assistant while a lawyer works. Both streamable HTTP and the legacy SSE transport are supported, and credentials are encrypted at rest. If something you run exposes an MCP server, the assistant can use it without us building an integration first.

REST API

Scoped credentials, read or read and write per key, an optional workspace pin, per key rate limiting that is Redis backed and distributed, and independent revocation. The API and the MCP server share one authorisation policy, so a scope means the same thing in both.

Programmatic access requires an active subscription and is not available on a trial.

Choice of model

Eight models are selectable per conversation, across Bedrock and Fireworks. GPT-OSS 120B is the default for its speed. Claude Haiku 4.5, Amazon Nova Pro, Kimi K2.6 and DeepSeek V4 Pro sit at the standard credit rate. Claude Sonnet 4.6, GLM 5.2 and MiniMax M3 are premium options at twice the credit rate on Professional and above. MiniMax M3 has native image understanding, which matters for scanned exhibits, and GLM 5.2 carries a very long context window for large document sets.

What it does not do

Every vendor's brochure lists capabilities. Fewer list limits, which is precisely why the limits are worth reading.

LIMIT	WHAT IT MEANS IN PRACTICE
It does not file with a court	The platform prepares the filing package and can drive your own browser on the desktop application, but a person reviews every screen and submits. Filing, and the deadline, remain yours.
It is not a lawyer	Answers are research and drafting output, not advice. Every citation should be opened and read before it is relied on. Models can attribute a holding to the wrong case.
Paywalled databases are out of reach	No access to LawPavilion Premium, Westlaw or LexisNexis. Coverage is our own Nigerian database plus the public sources listed earlier.
Eight jurisdictions are wired today	Anywhere else is not selectable yet. Because the sources are regional institutes rather than per-country integrations, adding one is usually configuration. Ask for a lead time on yours.
Workflows do not wait or branch	The unattended engine runs in one pass. Anything needing a wait or a branch is built as a process instead, which does both.
Tabular review has no spreadsheet export yet	The grid lives in the platform, where the citations stay attached to the answers. If you need the numbers in Excel today, tell us and we will tell you where that sits.

We would rather you find these here than three months into a rollout. If any is disqualifying, tell us and we will say plainly whether it is on the roadmap or not.

How firms start

A sensible first month at your scale.

Weeks one and two: one practice group in one office

Resist the firm-wide pilot. One group running live matters tells you more in two weeks than a broad rollout tells you in two months.

Week three: the library and access model

Import one group's precedent estate and configure grants by department and role. This is the piece your risk function will want to see before anything wider.

Week four: your engineers

Have them connect an MCP client to a sandbox workspace with read only scope, and exercise the REST API. That answers the integration and lock-in question with evidence rather than a discussion.

Then decide

Firm-wide skills and a process campaign are the two things that only make sense once the pilot has convinced people. Do them in month two, not month one.

Plans

Programmatic access, AI workflows, financials, reporting and firm tools sit on Professional and above. Current plans and pricing are at **modulaw.ai/pricing**, which is the single source of truth for figures. Enterprise adds single sign on, custom integrations and dedicated support.

Talk to us

The app is at **app.modulaw.ai**. Guides for everything in this document are at **modulaw.ai/support**. For a demonstration against your own matters, at whatever scale you want to test, reach us through **modulaw.ai**.