
WHAT WE CAN DO WITH MCP TOOLS

Your Systems, Now Something You Can *Talk To*.

We put a governed AI layer over the platforms we already build and run for you. It reads the documents your business receives, lifts out the terms and figures that used to be keyed in by hand, and turns the whole system into a conversation, inside the assistants your team already uses. Same data, same permissions, every write approved first.

Reads

CONTRACTS, SCANS
& SOURCE DATA

Answers

IN PLAIN ENGLISH,
ACROSS EVERYTHING

Acts

ON YOUR CONFIRM,
NEVER BLIND

Brands

CHARTS, WORKBOOKS
& REPORTS

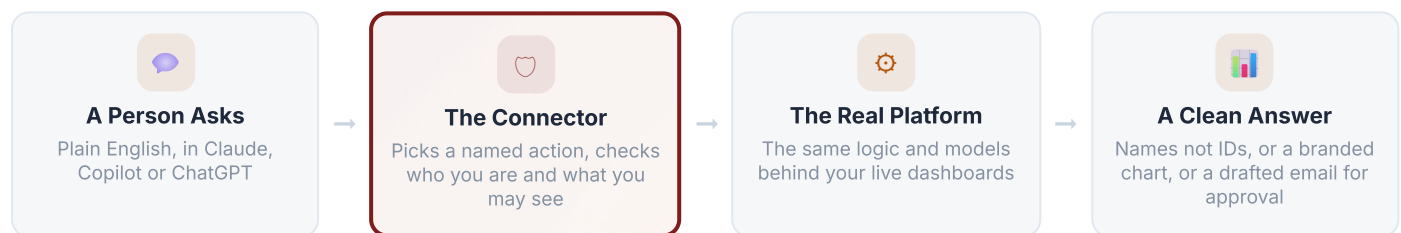
01

A Connector That Turns Your Platform Into a Conversation

MCP, the Model Context Protocol, is the standard way to give an AI assistant a set of named, permission-checked actions against a real system. We build one for each platform we run. Here is what that actually means.

Every business we work with already has a system of record: a fund-compliance platform, a legal practice-management system, a distributor's ERP, an accountancy reporting portal. Those systems hold clean, reconciled, trusted data. The problem is never the data. The problem is that using it means someone opening the right screen, on the right day, for one record at a time, and typing the answer into a report.

An **MCP connector** sits on top of that same system and exposes a fixed menu of safe actions to an AI assistant. The assistant does not get a database login, it does not write queries, it does not see raw tables. It gets a list of tools with names like **covenant_health**, **overdue_report** or **court_search**, and it can only do what those tools allow. Every call runs as the real person who signed in, inside their exact permissions, and every call is logged.



The one-line version: we do not build a new AI product bolted onto the side. We open up the platform you already trust, through a narrow, governed door, so your team can ask it anything and, when you approve, ask it to act.


Why this beats a dashboard, a chatbot, or raw database access

- **Better than a dashboard**, because a dashboard is pull-based and fixed. It answers the questions someone designed a page for. The connector answers the question you have right now, across every record at once, with no new page to build.
- **Better than a generic chatbot**, because it is wired to your real system. The numbers it returns are the same numbers as your reports, from the same models, so the AI and the dashboard can never disagree.
- **Safer than giving the AI a database**, because it never touches one. There is no AI-written SQL to inspect or repair. The assistant chooses from a menu; the platform composes every query itself.

02

One Governed Path, No Shortcuts

Every request, in every domain, follows the same route. Understand this once and you understand all of it. Nothing about it is a black box.



← A curated answer flows back: internal record IDs stripped, names only, every call logged to the real user

No direct path. The assistant cannot reach the database, open a connection, or run a query. Every request is mediated, permission-checked and recorded by the platform.

The five guarantees that make it safe by design

GUARANTEE	WHAT IT MEANS IN PRACTICE
Runs as the real user	Each request carries the signed-in person's identity into the platform's own authorization. They see exactly what their role allows, no more, whether that is one client or the whole book.
Named actions only	The assistant picks from a fixed list of defined tools. There is no free-form query, so there is nothing to inspect, escape or repair.
Same numbers as the reports	Every tool calls the same application services behind the live dashboards. The AI answer and the report are computed by identical logic.
Names in, names out	You refer to a borrower, a client, a customer by name. Internal database IDs are scrubbed from every response before it leaves the platform.
Audited and attributable	Every single call is written to a per-call log against the real user, alongside the platform's existing audit trail.

Where your team actually uses it: Claude, Microsoft Copilot and ChatGPT all speak MCP. We build the connector once and the same governed layer serves all three. Nobody learns a new tool; they ask the assistant they already have open.

03

It Reads Freely. It Only Writes When You Say So.

The large majority of tools are read-only and cannot change anything. A small number can write, and those are wrapped in a gate that no assistant can talk its way past.

Reading is the easy half: a read tool computes an answer and returns it. It cannot alter a record, so there is nothing to approve. **Writing is different, and we treat it differently.** A write tool never commits on the first call. It runs a dry run, shows you exactly what would change, and hands back a one-time **confirm-token**. Nothing is written until a human calls the same tool again with that token. The assistant cannot invent the token, cannot skip the step, and cannot batch it away.



Nothing is committed blind. Every write is dry-run first, shown to you item by item with the source beside it, and applied only when you supply the token. It goes through the same commit path as the app's own screens, never a side door.

WHAT A READ TOOL CAN DO

- ✓ Compute any figure the dashboards show, live, on demand
- ✓ Answer across the whole portfolio, client base or catalogue at once
- ✓ Rank, compare, trend and drill into the detail behind a number
- ✓ Read a stored document and extract its terms into a preview
- ✓ Render a branded chart, workbook or report from the result

WHAT NO TOOL CAN DO WITHOUT YOU

- ✗ Write a single record on the strength of a conversation alone
- ✗ Commit an extraction before you have seen it clause by clause
- ✗ Send an email, raise a filing or issue an invoice unconfirmed
- ✗ Reach past its named actions to the database underneath
- ✗ See or return an internal ID, or act outside the user's permissions

The net effect: your team gets the speed of "just ask" for everything that only reads, and the safety of "show me first, then I approve" for everything that changes the record. The AI proposes; a named person disposes.

04

Your Data, In Your Brand, Ready to Send

An answer in a chat window is useful. A chart on your letterhead, a workbook with your colours, or a client-ready report that came out of the same question is what actually leaves the building. Every connector can produce all three.

Alongside the data tools, each connector carries a set of **rendering tools**, named things like `*_branded_chart`, `*_branded_excel` and `*_branded_html`. Ask a question, then ask for it "as a branded chart" or "as a workbook for the board", and the platform renders the result in your identity: your logo embedded, your palette, tabular numerals, contrast-checked colours, a self-contained file that opens anywhere with nothing to install.



Branded Chart

SELF-CONTAINED IMAGE

- Server-rendered, no external scripts to load
- Distinct, accessible colours by default
- Drops straight into a deck or an email



Branded Workbook

EXCEL, READY TO WORK

- Logo, two-tone headers, right-aligned figures
- Real cells, not a picture of a table
- The board pack or the weekly review, formatted



Branded Report

A4 HTML TO PDF

- Print-ready document in your house style
- Commentary, tables and charts in one file
- The client-facing pack, generated on demand

And inside supporting assistants, the answer can be interactive

Some connectors go a step further with **inline application panels**. Instead of a static reply, a tool can return a small, interactive board that renders right inside the assistant: a covenant-health grid you can sort, a portfolio-compliance view you can filter, a trend you can scrub. It is the dashboard experience, summoned by a sentence, built from the same live data, with the same permissions applied.

Why this matters: the distance from "what is going on" to "here is the document I am sending about it" collapses to a single conversation. The person never leaves the assistant, never reformats anything, and what comes out already looks like you made it, because, in every way that counts, you did.

Chart

SELF-CONTAINED IMAGE

Excel

LIVE CELLS

HTML

PRINT-READY A4

Apps

INTERACTIVE PANELS

05

Four Very Different Worlds, The Same Governed Layer

We have built and deployed this pattern across four industries. Each has its own tools and its own language, but the architecture on the previous pages is identical in all of them. Here is who uses it and for what.

THE BUSINESS	WHAT THEY RUN ON	WHAT THE CONNECTOR GIVES THEM
A specialist fund manager debt to companies across African development-finance sectors	A fund-compliance and portfolio platform we built	Ingest signed facility agreements, extract covenants and reporting duties, and ask the whole portfolio how it is tracking
A law office litigation and advisory practice	A practice-management system, plus public legal sources	Search national case law and legislation, build the argument, draft on letterhead, compute deadlines, and send, on approval
A regional distributor packaging, several countries	An ERP and business- intelligence warehouse	Meeting prep, reorder forecasts, margin and debtor watch, and branded weekly packs, per company or consolidated
An accountancy practice hospitality-focused, many clients	A live, branded reporting portal we run for them	Read leases and facilities into structured data, and ask across every client's finance pack at once

The reuse is the point. When a new client comes to us, we are not inventing the safety model, the sign-in, the dry-run gate, the logging or the rendering tools. Those are ported from a system already in production and already through a third-party security review. What is new each time is the short list of business questions worth wiring up.

The next fifteen pages walk two of these worlds in real depth, the fund manager and the law office, because they show the full range: reading a forty-page contract into structured data and writing it back to production, and searching a national body of case law to shape and file a legal argument. The distributor and the accountancy practice then show the same layer applied to everyday commercial and reporting work.

4

INDUSTRIES LIVE

1

SAFETY PATTERN

3

ASSISTANTS SERVED

Prod

NOT A PROTOTYPE

06

Fund Management & Compliance

A specialist investment manager runs a family of funds that lend to portfolio companies across agriculture and trade, off-grid energy, financial institutions and local-currency bonds. Every borrower owes periodic reporting and must hold to financial covenants. Tracking all of it is the platform. The connector makes the platform answerable and ingest-capable.

The shape of the problem: each fund lends to borrowers through deals; every deal carries reporting obligations (audited accounts, management accounts, compliance certificates, budgets, environmental and impact reports) and financial covenants (ratios like debt-service cover, gearing, current ratio and minimum cash, each with a threshold and a test frequency). All of that lives in one platform, and the dashboards show it beautifully. The connector adds two things the dashboards cannot: it **ingests the source agreements** that define those obligations, and it lets staff **interrogate the whole book by asking**.



Ingest a deal

GUARDED WRITE · PAGES 8-9

- ▶ Read the facility agreement, scan or digital
- ▶ Extract covenants and reporting duties
- ▶ Reconcile to the right borrower and deal
- ▶ Preview, approve, write to production



Interrogate the book

READ-ONLY · PAGE 10

- ▶ Covenant health and headroom, live
- ▶ Who is overdue, what is deteriorating
- ▶ Stress tests, forecasts, provisioning
- ▶ Board pack, in your brand, on demand

What makes this connector match the dashboards exactly


- **The tools call the same services as the app.** A covenant-health answer here and the covenant-health tile in the portal are computed by the identical code. There is no second, drifting version of the truth.
- **Fund access is scoped per user.** Someone who can see three funds in the app sees exactly those three through the connector. The whole-portfolio views are gated to whole-portfolio access.
- **Names in, names out.** You ask about a borrower by name and get answers by name. Internal identifiers never appear in a reply.
- **Every call is logged.** A per-call record ties each question and each import to the person who made it, with the import artefact path recorded for writes.

The everyday picture: a signed agreement arrives, and by that afternoon its covenants and reporting deadlines are on the borrower's record, checked and approved. Before a board call, an officer asks which borrowers are within ten percent of a covenant breach and gets it in a sentence, then asks for it as a branded workbook.

07

From a Signed PDF to a Structured Extraction

This is the flagship capability, and it runs in production today. A facility agreement goes in as a file. What comes out is every covenant and reporting obligation it contains, structured, checked and ready for sign-off. Here is the first half: get the document in, read it, extract it.



Step by step, with the tools that do it

STAGE	WHAT HAPPENS
Get a secure upload slot <code>create_deal_upload_url</code>	The platform mints a time-limited, access-controlled upload URL that lands the file in the fund's own folder in Azure Blob storage. Nothing is public.
Store the document <code>upload_deal_document</code>	The signed PDF or Word file is written to <code>deals/{fund}/</code> , encrypted at rest, filed against the fund, access-controlled.
See what is on file <code>list_deal_documents</code>	Lists the agreements already stored for a deal, so the operator works from the real, current set, not an emailed copy.
Read the content <code>read_deal_document</code>	Returns the text. Born-digital PDFs are read directly; scanned or photographed pages are put through OCR automatically, so a photocopied agreement is as readable as a clean export.
Extract the terms <code>(extraction skill)</code>	The assistant reads only the obligation and covenant clauses, plus the definitions they depend on, and lifts them into a strict deal-import format: each covenant with its ratio, operator, threshold and test frequency; each reporting duty with its document type, frequency and deadline.

Targeted reading, not a firehose. The extraction does not swallow the whole forty pages. It maps the document's structure first, then reads the clauses that carry obligations and the definitions behind them. That is what makes it accurate on long, cross-referenced agreements instead of vaguely summarising them.




Still nothing written. At this point the covenants and obligations exist only as a proposed extraction on the operator's screen. Not one has touched the borrower's record. The commit gate is on the next page.

08

Reconcile, Check, Approve, Write to Production

The second half is where the extraction meets the live system. It is matched to the right borrower, checked against what is already configured, previewed in full, and only then, on the operator's approval, committed to production through the same path the app's own screens use.



STAGE	WHAT HAPPENS
Match to the borrower <code>reconcile_borrower</code>	Borrowers exist under short names and codes already. Reconciliation matches the extraction to the correct existing borrower and deal, so the import attaches to real records and does not create duplicate assignments.
Dry-run the import <code>validate_deal_import</code> <code>validate_covenant_import</code>	The proposed assignments and covenants are checked against the live configuration. Anything already present, in conflict, or malformed is surfaced now, before any write.
Commit, on approval <code>import_deal_files</code> <code>import_covenants</code>	Guarded writes. They preview first and return a confirm-token; only a second call with that token commits. The write goes through the same services as the app's own configuration screens.
Confirm it landed <code>check_deal_ingested</code> <code>verify_extraction</code>	After the commit, the operator confirms the covenants and obligations are live on the borrower, exactly as approved.
Undo cleanly if needed <code>remove_covenant /</code> <code>remove_assignment</code>	Report-first removal: it shows what it will deactivate before it does, at borrower or fund scope, through the same deactivation path the UI uses.

What actually lands on the borrower's record

FROM THE AGREEMENT	BECOMES, IN PRODUCTION
Financial covenants	Covenant assignments with ratio, operator, threshold and test frequency, at the right scope, monitored from the next reporting period
Reporting obligations	Reporting-requirement assignments with document type, frequency and deadline, which drive the pending-submission and reminder engine automatically

Why this is the flagship: obligations that used to be read off a long agreement and held in one person's head become structured, dated and monitored, in minutes, with a full preview and a human sign-off, and the reminders key off what the platform read, not what someone remembered to type.

09

Ask the Whole Portfolio Anything

With the deals configured, the reads open up. Every one of these is a named tool that computes from the same live services as the dashboards, scoped to what the person is allowed to see, returned by name.



Compliance & Covenants

LIVE STATE

- covenant_health, covenant_trend
- portfolio_compliance, overdue_report
- waivers, risk_classification_history



Early Warning

WHAT IS MOVING

- deterioration_scan, covenant_runway
- portfolio_watchlist, risk_snapshot
- data_quality_alerts, market_news



Treasury & Stress

FORWARD VIEW

- repayments_due, expected_cashflows
- stress_fx, stress_default
- provisioning_inputs



Portfolio Analytics

THE SHAPE OF THE BOOK

- top_borrowers, portfolio_breakdown
- growth_trend, exposure, limits_report
- irr_summary, deal_status, taf_summary



Reporting & Detail

DOWN TO THE CELL

- borrower_briefing, borrower_profile
- form_submission, metric_trend
- analyze_submissions, approvals_queue



For the Board

IN YOUR BRAND

- board_pack, fund_financials
- fund_commentary, screening_context
- cygnum_branded_chart / excel / html

What that sounds like in the assistant

Covenant headroom	Which borrowers are within ten percent of breaching a covenant this quarter, worst first?
Deterioration	What has deteriorated across the portfolio since last quarter, and why?
Overdue	Who is overdue on reporting right now, and is there a waiver on file for any of them?
Stress	If the local currency moved fifteen percent, which positions take the hit and how much?
Board pack	Build the fund's board pack for this month and give it to me as a branded workbook.

The qualitative layer matters too. Before the assistant narrates a breach or an overdue as a problem, it can check **waivers** and **market_news**, so it tells you "overdue, but there is an approved extension", not a false alarm.

10

The Fund Connector, In Full

For reference: the tool categories in the production fund-compliance connector. The overwhelming majority are read-only; the guarded writes are the deal, covenant and removal tools already described.

CATEGORY	ACCESS	REPRESENTATIVE TOOLS
Roster & profile	Read	list_funds, list_borrowers, borrower_profile, borrower_briefing, exposure
Compliance & covenants	Read	covenant_health, covenant_trend, portfolio_compliance, overdue_report, waivers, risk_classification_history
Early warning	Read	deterioration_scan, covenant_runway, portfolio_watchlist, risk_snapshot, data_quality_alerts, market_news
Treasury & stress	Read	repayments_due, expected_cashflows, stress_fx, stress_default, provisioning_inputs
Portfolio analytics	Read	top_borrowers, portfolio_breakdown, growth_trend, limits_report, irr_summary, deal_status, taf_summary
Management accounting	Read · admin	management_pnl, gl_lines, xero_journals, fund_financials, fund_commentary
Reporting detail	Read	form_submission, document_submissions, metric_trend, analyze_submissions, approvals_queue
Board & screening	Read	board_pack, screening_context
Branded rendering	Read	cygnum_branded_chart, cygnum_branded_excel, cygnum_branded_html
Interactive panels	Read	covenant_health, covenant_trend, portfolio_compliance (render inline apps)
Deal document loop	Read	create_deal_upload_url, upload_deal_document, list_deal_documents, read_deal_document
Reconcile & validate	Read	reconcile_borrower, validate_deal_import, validate_covenant_import, find_similar_borrower / metric / requirement
Deal & covenant import	Write · gated	import_deal_files, import_covenants
Report-first removal	Write · gated	remove_covenant, remove_assignment, remove_fund_covenant, remove_fund_assignment
Submission backfill	Write · gated	backfill_reporting_submissions
Facility-to-workbook	Read	d1m_extraction_spec, validate_d1m_extraction, render_d1m_extraction

Reading the access column: "Read" tools cannot change anything. The "Write · gated" tools are the only ones that can, and none commit without a dry-run preview and a human-supplied confirm-token.

11

The Law Office That Researches, Drafts and Files by Asking

A litigation and advisory practice runs its matters, clients, deadlines, hearings and billing on one system. Beside it sits a second connector over the national public legal sources. Together they take a case from a legal question to a filed, deadline-tracked pleading, with the lawyer approving every step.

There are two connectors here, and the split is deliberate. One is the **practice-management** layer: matters, clients, deadlines, tasks, hearings, documents, time and billing, the office's own working data. The other is the **legal-research** layer over public sources: the national court-decisions portal, the official gazette of legislation, and the consolidated statute texts. The lawyer moves between them in a single conversation: research the law, then act on the matter.



Research the law

PUBLIC SOURCES · PAGES 13-14

- Search national case law, millions of decisions
- Read the strongest, see who cited them
- Pull the authoritative statute text
- Verify and format every citation



Run the matter

OFFICE DATA · PAGE 15

- Draft on the office letterhead
- Compute procedural deadlines exactly
- File the email into the matter
- Log time, raise the cost statement

The non-negotiables built into this domain

- **Every legal citation traces to a real source returned by a tool.** Nothing is invented. Each result carries its origin, the court-decisions portal, the official gazette, the consolidated-text service, and the source is shown with the answer.
- **Authoritative text comes from the official gazette,** not the convenience of a consolidated copy, whenever the exact wording of the law matters.
- **Deadlines are computed, never estimated.** Procedural time limits are calculated against the civil-procedure rules and the calendar of non-working days, and a statutory forfeiture deadline is always kept distinct from an internal one.
- **Every change is two-step.** Create a client, add a deadline, log time, send an email, each previews first and commits only on the lawyer's explicit confirmation.

The through-line: the connector does not replace legal judgement. It does the searching, the reading-at-scale, the deadline arithmetic and the paperwork, and hands the lawyer a sourced draft to approve. Faster to the argument, and safer, because every citation is traceable and every deadline is calculated.

12

Search Millions of Decisions, Keep the Ones That Win

This is the legal flagship. The lawyer describes the dispute; the connector searches the national body of court decisions, filters to what is on point, reads the strongest, and checks how each precedent has been followed since. The argument then gets built toward the line of authority that actually holds.




STAGE	WHAT HAPPENS
Search the case law <code>court_search</code>	Queries the national court-decisions portal across a very large body of published rulings, filtered by court, by area, by date, by the legal question in play. Returns candidates ranked by relevance, each with its court and its source.
Read the strongest <code>read_decision</code>	Opens a decision in full so its reasoning, not just its headnote, can be weighed. This is where the ones that genuinely support the case are separated from the ones that merely mention the topic.
Test how it held up <code>citing_decisions</code>	Shows which later rulings cited a given decision, so a precedent that has been consistently followed is preferred over one that was later distinguished or ignored. The argument leans on authority that lasted.
Anchor to the statute <code>legislation_search</code> <code>statute_consolidated</code>	Finds the governing legislation and pulls the authoritative wording from the official gazette, with the consolidated text alongside for reading. The pleading cites the law as enacted, not a paraphrase.
Watch for change <code>legislation_watch,</code> <code>legislation_diff</code>	Flags if a relied-upon provision has been amended, and shows exactly what changed, so an argument is never built on a superseded version.

The move that matters: the search is not "find something about this topic". It is "find the decisions that support this position, read them properly, confirm they still hold, and cite the statute behind them". Reading at a scale no person has time for, then narrowing to what a court will actually respect.

13

Write the Case in That Direction, Then File It

With the supporting authority in hand, the same conversation turns to producing the document: draft it on letterhead toward the winning line, verify the citations, compute the filing deadline, and send, on the lawyer's confirmation, filed straight into the matter.



STAGE	WHAT HAPPENS
Draft the document <code>generate_draft / generate_document</code>	Produces the pleading or letter on the office letterhead, structured around the supporting decisions and the statute, in the matter's context. It lands as a draft, not a sent document.
Verify the citations <code>citation_check</code> <code>citation_format</code>	Checks each legal reference resolves to a real source and formats it to the accepted citation convention, so the draft is clean before anyone reads it.
Lawyer approves <code>approve_draft</code>	The draft is reviewed and explicitly approved. Until then it can be revised or returned to draft; nothing leaves the office on the assistant's say-so.
Compute the deadline <code>deadline_calculate</code> <code>add_deadline</code>	The filing or appeal deadline is calculated against the civil-procedure rules and the non-working-day calendar, then diarised on the matter. A statutory forfeiture date is marked distinctly from an internal target.
Send and file <code>send_email</code> <code>file_email_to_matter</code>	Two-step: the send previews first and goes only on the lawyer's confirmation, then the outgoing message is filed into the matter's record. Time can be logged and a cost statement raised in the same flow.



Approve, then act. Draft, deadline and send are all confirm-gated. The assistant assembles and proposes; the lawyer signs off. Every legal claim in the draft traces to a source the research tools actually returned.

End to end: "a dispute about X" becomes a sourced, letterhead pleading built on decisions that have held up, with the filing deadline calculated and diarised, filed into the matter and the time logged, in one conversation the lawyer stays in control of throughout.

14

The Rest of the Office, By Asking

Research and drafting sit on top of a full practice-management connector. The whole running of the office, matters, deadlines, hearings, billing, becomes a conversation, with every change confirm-gated.


Matters & Clients
READ & TWO-STEP WRITE

- list_matters, matter_console
- client_profile, conflict_check
- create_matter, create_client


Deadlines & Tasks
COMPUTED, NEVER GUESSED

- upcoming_deadlines, limitations_watch
- my_agenda, daily_dashboard
- add_deadline, add_task, complete_task


Hearings
PREPARED IN ADVANCE

- upcoming_hearings, hearing_brief
- generate_hearing_brief
- schedule_hearing, record_hearing_outcome


Time & Billing
TO INVOICE

- log_time, log_court_cost
- generate_cost_statement, generate_invoice
- overdue_invoices, matter_billing


Documents & Mail
FILED TO THE MATTER

- upload_matter_document, read_matter_document
- route_incoming_mail, file_email_to_matter
- generate_document, generate_workbook


The Whole Office
AT A GLANCE

- office_overview, office_coverage
- team_workload, intake_queue
- generate_chart (branded)

A morning in the practice

Start the day	<i>Give me my agenda: today's deadlines, hearings and anything that has just become urgent.</i>
Limitation watch	<i>Which matters have a limitation period expiring in the next thirty days?</i>
Prep a hearing	<i>Prepare the brief for tomorrow's hearing on the Whitmore matter.</i>
Chase the money	<i>Which invoices are overdue, and draft the reminders.</i>

Same governance as everywhere else: entities are shown by name and matter reference, never internal IDs; every data change is a two-step confirm; and the whole office view is scoped to what the signed-in team member is allowed to see.

15

The Distributor Who Preps Every Sales Call by Asking

A regional packaging distributor operating across several countries runs its sales, stock and receivables in an ERP, fed into a business-intelligence warehouse. The connector turns that warehouse into a sales assistant: meeting prep, reorder timing, margin and debtor watch, per company or consolidated across the group.

Staff mostly ask in their own language, and the connector answers in kind. A single user sees only their company; a group user can split or consolidate across all of them. The aggregate tools insist on knowing which company is in scope before they answer, so a "how are sales" question is never silently answered for the wrong entity.



Meeting Prep

ONE CALL, THE WHOLE CUSTOMER

- get_customer_briefing
- get_sales_board_pack
- get_customer_watchlist



Sell More, Smarter

FORECAST & WHITESPACE

- get_reorder_forecast
- get_cross_sell_whitespace
- get_price_increase_simulation



Protect the Margin

ANOMALIES & DEBT

- get_margin_anomalies
- get_top_debtors, get_receivables_overview
- get_stock_cover, get_open_complaints

What the sales team actually asks

Meeting prep	<i>Prep me for the meeting with customer X: everything I need, and what else I could offer them.</i>
Who to call	<i>Who should I call this week, who is overdue an order by their own rhythm, and what?</i>
Price move	<i>What does a four percent increase on cartons do to us, and who is most exposed?</i>
Below cost	<i>Where did we sell below cost this month? Group it by salesperson.</i>
Weekly pack	<i>Build the weekly sales review for one company as a branded document to print.</i>

The branded weekly pack: the same question that answers "how did we do" also renders the answer as a branded chart, a workbook or a print-ready report, in the distributor's identity, ready to hand to a manager. Ask once, get the analysis and the deliverable together.

Grounded, always: the connector never states a figure it did not read from the system, always shows how current the data is, and resolves fuzzy names ("the brewery", "that stretch film") to the real customer and item before answering.

16

The Accountancy Practice That Reads the Paperwork

An accountancy practice runs a live, branded reporting portal for dozens of client businesses. The connector adds the two things a fixed report cannot do: it reads the documents clients send, and it answers questions across every client at once.

The portal is excellent at what it does, per-client finance packs that refresh live, with profit and loss against budget, site breakdowns, benchmarks and cash-flow. But it is pull-based: someone opens a workspace, on the right page, on the right day. And it cannot read a lease or a bank facility; those are still keyed in by hand. The connector closes both gaps with the same document pipeline and the same governed reads as every other world.



It reads the documents

THEN YOU APPROVE

- ▶ Commercial lease: rent-review and break dates, service charge
- ▶ Bank facility: covenants, thresholds, lender reporting dates
- ▶ Supplier contract: payment terms, price escalators, renewals
- ▶ Source data: mapped and reconciled into the finance pack



Ask across every client

ONE QUESTION, WHOLE BOOK

- ▶ Which clients are below budget EBITDA this period?
- ▶ Rank clients by wage-to-sales; who breaches the benchmark?
- ▶ Whose margin fell most versus last period?
- ▶ Which clients are behind on data this week?

Proven pattern, ported: the document-reading pipeline, the dry-run approve-before-commit gate, the sign-in and permission checks, the ID-stripping and the call logging are all lifted from the fund-compliance system in the first half of this brief. The practice does not get an experiment; it gets a working pattern applied to its own data.

The everyday picture: a manager asks which clients breach the wage benchmark before the partner call and gets it in a sentence. A new lease lands and its dates are on the client's file that afternoon. A client's commentary is drafted from their own numbers, ready to review. Live, inside each person's permissions, all logged.

17

Why We Can Put This Near Live Data

The reason a governed AI layer over production systems is defensible, and has passed a formal third-party security review, is that safety is not a policy bolted on afterwards. It is the architecture.

CONTROL	HOW IT ACTUALLY WORKS
Corporate sign-in	Users authenticate through a central identity broker with their own corporate account. The connector is a resource server that validates the token on every call; there is no shared or standing credential.
Staff gate	Beyond a valid sign-in, access is gated to authorised staff. Anyone without a mapped account is denied, fail-closed, not fail-open.
Runs as the real user	A context bridge carries the signed-in identity into the platform's own authorization, so every call inherits that person's tenant and permissions. The AI has no permissions of its own.
No AI-written queries	Tools call the application's own services. The assistant never composes a query; the platform composes every one. There is no query surface to exploit.
Names in, names out	Requests and responses use business names. Internal record identifiers are scrubbed from every payload before it leaves the platform.
Confirm-token writes	Every write dry-runs, returns a one-time token, and commits only on a second call with that token, through the same commit path as the app's own screens.
Per-call logging	Every call is written to a dedicated log against the real user, with the action and, for writes, the artefact that was committed. Fully attributable after the fact.
Feature-flagged	The whole layer sits behind a flag with a central off switch. It can be turned on for a pilot group and off instantly, without touching the platform it sits inside.

Independently reviewed: this exact pattern, the document pipeline, the governed connector, the sign-in and permission model, the ID-stripping and logging, has passed a formal third-party security review in a regulated financial-services setting. Every new deployment reuses the reviewed pattern rather than reinventing it.

The short answer to "is this safe": the assistant can only do the named things you gave it, only as the person who signed in, only within their permissions, never to the database directly, never writing without an approval, and every move is logged. It is a narrow, watched door onto a system you already trust.

18

No New Servers, One New Subscription

The connector ships inside the platform we already run for you. There are no new databases or licences on the platform side. The only recurring cost is the assistant subscription for the people who use it.

Prerequisites

- **An AI assistant that supports custom connectors.** Claude, Microsoft Copilot and ChatGPT all qualify. This is the only new recurring cost the layer introduces.
- **The platform reachable over HTTPS.** Already true of the deployments we run; the connector lives inside the same system, behind a feature flag with a central off switch.
- **A sign-in mapping per user.** Each pilot user's assistant sign-in is linked to their own account; anyone without a match is denied.

Where the team uses it, and the list prices

PLAN	WHO IT FITS	LIST PRICE	CONNECTOR
Claude Pro	One user; the fastest way to try a pilot	≈ £17 per user / month	Supported
Claude Team (recommended)	The pilot group and beyond, central billing and admin	≈ £22–£25 / user / month, min. 5 seats	Supported
Claude Enterprise	Organisation-wide rollout with SSO and data controls	Custom	Supported
Copilot / ChatGPT	Later phases; the same connector serves them with no rework	≈ £20–£25 / user / month	Same standard

What a pilot actually costs: five Claude Team seats come to roughly **£110–£125 per month**. Monthly billing means a pilot can start, prove itself or not, and stop, with nothing stranded.

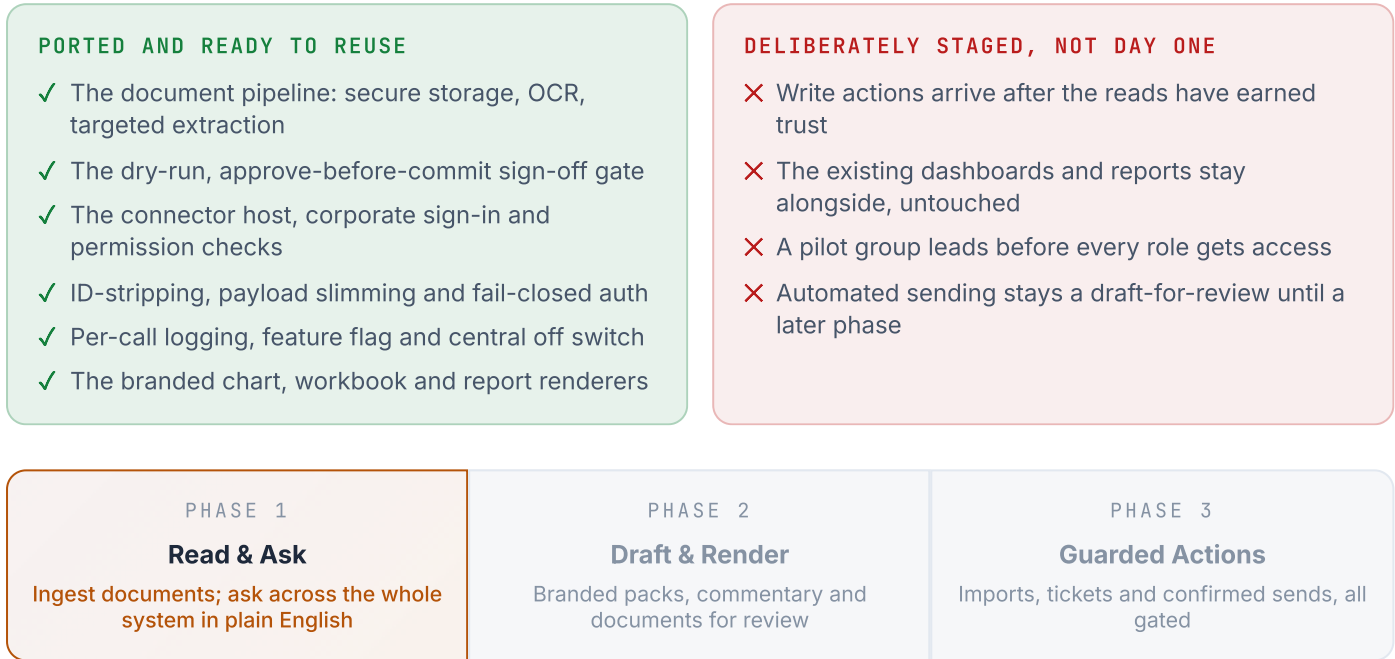
And no metered AI bill. Usage is covered by the flat subscription. There is no per-question API charge on top; the seat price is the entire recurring cost. Development to wire up each connector's action list is quoted separately, once that list is agreed.

Note on pricing: figures are vendor list prices as of mid-2026 and are confirmed at kickoff.

19

Not a First Attempt

Everything in this brief runs in production somewhere today. Bringing it to a new system is a porting exercise on a reviewed pattern, not a research project. Here is how a new engagement typically phases.



How a new engagement starts

- **Agree the action list.** A short workshop to pick the twelve to fifteen questions and one or two document types that save the most time.
- **Stand up the connector behind a flag,** off by default: the reads, the document reader if in scope, the sign-in, the logging and the off switch.
- **Pilot with a small group,** then review the call log and the extractions together, and decide what to add.
- **Add guarded writes only when the reads have earned it,** each one dry-run and confirm-token gated from the start.

20

What We Can Do, In One Page

The whole brief, compressed.

Read

ASK ANYTHING, LIVE

Ingest

DOCUMENTS TO
STRUCTURED DATA

Act

ON YOUR CONFIRM

Brand

CHARTS, WORKBOOKS,
REPORTS

- **We open your own platform through a narrow, governed door.** The assistant gets a fixed menu of named actions, never the database, never a query, never permissions of its own.
- **It reads the documents your business receives** and lifts out the terms and figures that used to be keyed in by hand, storing to secure cloud storage, reading scans by OCR, extracting into structured data, and writing to production only after a human approves it clause by clause.
- **It answers across everything at once,** from the same logic as your dashboards, so the AI and the report never disagree, scoped to what each person may see, returned by name, every call logged.
- **It produces what actually leaves the building:** branded charts, working Excel, and print-ready reports in your identity, generated from the same question.
- **It is proven and reviewed.** The pattern runs in production across fund compliance, legal practice, distribution and accountancy, and has passed a third-party security review.

The one thing to remember: we are not selling an AI product to bolt on the side. We are making the systems you already trust answerable, ingest-capable and able to act, safely, inside the assistants your team already has open. The AI proposes; your people dispose.

Next step: a one-hour workshop to choose the first dozen questions and the first document type for your system. From there, a flagged pilot with a small group, and a shared review of the log before anything grows.

Proceptio

The Agentic Intelligence Layer for the systems you already run

proceptio.com

Confidential · July 2026