AUGELA
The Human AI Interface

THE AI PLAYBOOK FOR BUSINESSES

Capturing your company's AI knowledge.

Ground your agents, govern what they do, and compound what they learn, so nothing useful evaporates when the tab closes.



CONTENT

· Why AI feels useful but never adds up
· The three-step loop
· The two things worth capturing
· The four tutor outcomes
· A 3-hour rollout plan

WHY THIS MATTERS NOW

Right now, your company has no memory of what its AI knows.

Your team is already using AI. The problem is that almost nothing it produces survives the conversation it happened in. A colleague works out the right way to explain your pricing to a sceptical client, someone finds a workaround that saves an hour a week. Then the tab closes, and all of it is gone. Next week, someone else starts from scratch.

- **Corrections stop being repeat work** – the source gets fixed, not just the answer in one chat.
- **Discoveries become company property** – not something one person happens to remember.
- **Your knowledge base sharpens as you use it** – instead of quietly going stale.

THE TRANSFORMATION

From "Every useful thing your AI produces dying in a chat window."

To "A loop that turns it into permanent company knowledge."

BEFORE AND AFTER

Four everyday failures. One better way.

What changes once the loop is actually running.

OLD · THE SAME WRONG ANSWER
Someone spots a wrong answer and moves on. Next day a colleague gets the same wrong answer.

NEW · FIX THE CAUSE
The first person marks the wrong answer for a review, a reviewer fixes the problem at the source.

OLD · EVAPORATED INSIGHTS
Someone works out a genuinely brilliant explanation in a conversation. There is nowhere for it to go, so it dies there.

NEW · ELEVATE THE INSIGHT
The conversation can be distilled into insight review. A tutor elevates it and adds to company knowledge.

OLD · NO IDEA WHAT WAS SAID
A client asks what your AI told them last quarter. You don't know.

NEW · A VERIFIABLE RECORD
A tamper-evident audit trail of every AI decision with reasoning.

OLD · KNOWLEDGE IN ONE HEAD
What the company knows lives with the oldest employee or documents nobody has opened since onboarding.

NEW · GROUNDED AGENTS
Agents answer from your connected sources, so the answer reflects your reality instead of a plausible guess.

THE LOOP

Three steps. One compounding system.

Each step makes the next one work better. And the third feeds back into the first. That is what makes this compound rather than accumulate.

1

Ground your agents

Stop your AI guessing. Connect the places your knowledge already lives, and give the agent an identity your colleagues will trust.

- Connectors
- Brand voice
- Fallback behaviour

2

Govern what they do

Make AI something you can approve rather than tolerate, because you can see and control what happens.

- Roles
- Guardrails
- Watch words
- Audit trail

3

Compound what they learn

Turn every correction and every discovery into permanent knowledge. Reviewed by a human, placed at the right level of your organisation.

- Response review
- Insight elevation
- Escalation

↑ STEP 3 BECOMES KNOWLEDGE FOR STEP 1

STEP 1 – GROUND YOUR AGENTS

Goal: stop your AI guessing

An assistant that answers from your documents, your policies and your customer history is a colleague. Connect places where your knowledge already lives.

CONNECT	FLOW	UNLOCK
Google Drive / OneDrive	Docs, decks, policies	Answers grounded in the documents you already maintain
Slack / Gmail	Threads, decisions, action items	Informal decisions become searchable instead of scrolling away
HubSpot CRM	Contacts, deals, history	Customer-aware answers without anyone looking a record up
GitHub	Issues, code, project status	Technical context available to people who don't read code
MCP server / upload	Anything else you run	A path for systems that don't have a native connector

TURN "AI" TO "YOUR AI"

Define your **agent's voice**, **specialization**, and strict **fallback rules** – because a bad guess is worse than a referral. Start small by deploying one agent with a single connected folder for one team, as company-level knowledge flows down automatically and allows early setups to evolve.

STEP 2 – GOVERN WHAT YOUR AGENTS DO

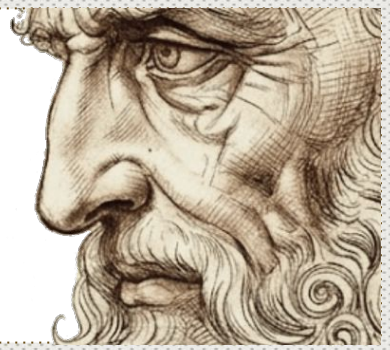
Goal: make AI you approve, not tolerate

Governance sounds like a brake. In practice it is what lets you say yes to more and better AI, because you see and control what is going on.

LAYER	SET UP	WHAT YOU GET
Roles	Admins configure · reviewers assess answers · tutors decide what becomes shared knowledge · everyone else just uses it.	Automated AI provisioning. Company AI knowledge management.
Guardrails and Watch words	Rules the agent must always follow; flags on sensitive terms, competitor names, anything you want seen.	Precise AI answers, guarded in passive (flag and log) or active (warn, hold, or block before a response) mode.
Audit trail	Every AI decision and admin action, hash-chained, with reasoning, confidence and cost.	Always on AI audit with exports to CSV or JSON with chain verification, tied to your SIEM.

AUDIT RULE OF THUMB

The audit trail is the way to turn "We have the visibility into our AI" claim into something you can hand to a client, an auditor, or a regulator. Have it before you need it, not when someone asks.



STEP 3 – COMPOUND WHAT YOUR AGENTS LEARN

Goal: make corrections and discoveries permanent

This is the step a few companies do, and it is where the compounding actually happens. Two distinct things are worth capturing, and they need different paths.

PATH 1 · CORRECTIONS

A user or system flags a wrong answer. A reviewer diagnoses the cause – the source document, configuration, or phrasing – and fixes that.

Reactive: something went wrong, and now it stops going wrong for everyone.

PATH 2 · DISCOVERIES

Anyone, in any conversation, can elevate an insight. One click distils the exchange into a short titled card, which the finder can refine before submitting.

Proactive: it captures anything valuable.

Then a tutor decides, you can get at least four outcomes

Elevate
Becomes real retrievable knowledge, chunked and embedded, so future answers draw on it.

Coach
Returned with guidance. Explicitly not a rejection. It's how people learn to spot what's worth keeping.

Reconfigure
The real fix is a setting, not a knowledge gap links straight to the agent's configuration.

Reject
With a reason, recorded.

The review gate is the feature, not the friction. Nothing becomes company knowledge unless a human elevates it. Automatic learning means automatic propagation of mistakes.



WHERE THE KNOWLEDGE LANDS

An insight goes up one level, not into a pile.

This is the part that matters. An elevated insight doesn't land in a generic knowledge dump. It goes one level up from where it was found, and the target is computed automatically, so the person submitting never has to guess who needs it.

AGENT	where the insight was found
GROUP	the team that shares the same context
COMPANY	everyone benefits – the escalation ceiling

If something turns out to be more broadly useful than it first appeared, it can be escalated further up the same ladder, carrying its endorsement trail, so you can always see where a piece of company knowledge came from and who vouched for it.

Near-duplicate submissions are caught automatically, so the knowledge base sharpens instead of bloating. And because every elevated insight becomes grounding knowledge, the next answer is better, the next correction is about something more subtle, and the next insight is more valuable.

Why this is hard to copy by your competitors. They can copy a feature list. They cannot copy months of your company's curated corrections and discoveries.

ROLLOUT

Your 3-hour plan.

Hour 1

Ground one agent. Pick the team with the most repetitive questions. Deploy AI agent via template from Augela, or create your own. Connect knowledge source, set the brand voice and the behaviour.

Hour 2

Turn on governance. Assign one reviewer/tutor. Add your watch words. Decide whether governance runs passively or actively. Optionally, confirm the audit trail is capturing what you expect.

Hour 3

Close the loop once, deliberately. Have someone flag a wrong answer and watch a reviewer fix the source. Then have someone elevate a real insight and walk it through to Elevate.

THE END STATE

Your AI stops repeating the same mistakes, your team's best thinking stops evaporating, and the longer you run it, the better your AI and your team gets.

Start a 7-day free trial →

· No credit card required
· Flat monthly price for your team
· Compliance and GDPR by design.