Tandem

WHITE PAPER

Same company, different answers

Company AI tools speak for the business every day. This paper looks at how they decide what's true about a company, why that usually goes ungoverned, and what it takes to change that.

Matthew Snyder, Founder and CEO, Tandem · November 2026

Summary

Most companies now give their AI tools access to the company's documents and call the result a company brain. It answers questions, drafts proposals and writes in the company's name. But nobody can say what it believes about the company, because it doesn't hold a view. It searches the pile each time and works from whatever it finds: current or out of date, approved or never approved, right or wrong.

The platforms behind these tools govern who can see which files. They don't govern which statements are true, current and approved. That second job belongs to the company, and today it mostly goes undone.

This paper describes the gap, what a governed layer needs, and how Tandem provides one: a single, readable record of what the company has decided is true and how it wants to be described, used by every tool the team works in, and changed only when a named person approves it.

How a company brain works today

The setup looks different from company to company. It might be a business account for ChatGPT or Claude with shared projects, Copilot reading SharePoint and email, an enterprise search tool, or a knowledge base the company built itself. Underneath, most work the same way.

1. Someone asks a question or makes a request.
2. The system searches the documents it's allowed to reach and pulls back the passages that seem to match.
3. The AI writes an answer from those passages.
4. The conversation ends, and nothing about the company is kept.

This is a good design for finding things. It lets a new hire ask about a project from three years ago and get a useful answer. But it means the company brain has no settled understanding of the company. It has a search result, assembled fresh every time, from whatever was in the folders.

Why more documents don’t make a truer answer

A company's files are a record of everything it has ever written, not a record of what it currently stands behind. Connect the AI to more of them and it gets better at finding things, but not better at knowing which of them is right.

- Old and new sit side by side. Last year's capabilities statement and this year's are both in the folder, and the search has no way to know one replaced the other.
- Drafts count as much as finals. A rough proposal, a working document, or a claim written once and never approved carries the same weight as the version that went out.
- The answer depends on the question. Two people asking slightly different questions can get two different descriptions of the same company, because different passages matched.
- There's nothing to look at. An answer may cite its sources, but there's no place to read the company as the AI understands it, and nothing to correct when it's wrong.

None of this is anyone's fault. It's what happens when a search over a growing library is asked to speak for the company. The more the company writes, the more versions of itself the AI can find.

Access is governed. Accuracy isn’t.

Enterprise AI platforms take governance seriously, and they should. Admins decide which files each person can reach, which connections are allowed, how long data is kept, and whether it can be used to train models. Those controls answer an important question: who can see this?

They don't answer a different one: is this true, current and approved for us to say? A platform can't know that a pricing page was retired, that a claim never cleared legal, or that the company now describes itself differently since the rebrand. Only the company knows those things, and in most companies that knowledge lives in a few people's heads.

What enterprise AI governs	What usually goes ungoverned
Who can open which files	Which version of a statement is current
Which tools can connect	Which claims are approved, and which are off limits
How long data is kept	How the company wants to be described
Whether data trains a model	What the AI has concluded about the company

The result is a company brain that is secure and still unreliable about the company itself.

What a governed layer needs

If the documents are the record, the company also needs an approved version: a small, deliberate set of what it has decided is true and how it wants to be described. For that to work with AI tools, it needs five things.

1. One place. The approved version lives in a single source the company controls, not scattered across folders and prompts.
2. A status on everything. Each fact and claim is marked approved, needing review, or prohibited, so a tool knows what it can say and what it can't.

3. Something you can read. People can see what the AI has been given about the company, in plain language, and change it.
4. A gate for new knowledge. Anyone can suggest something, but it only becomes part of the approved version when a named person signs off.
5. The same answer everywhere. Every tool the team uses draws from the same source, so a change made once reaches all of them.

How Tandem does it

Tandem is built around that approved version. We build it from what the company already has (its brand guide, templates, decks, website and approved language), and the company's team reviews and refines it with us.

What a governed layer needs	How it works in Tandem
One place	The guide, assets, templates, approved claims and rules live together in one place the company controls.
A status on everything	Every claim carries a status: approved, needs substantiation, needs legal review, or prohibited.
Something you can read	A brand understanding page shows what Tandem has concluded about the company. People can read it, edit it, and lock the parts that must not change.
A gate for new knowledge	Anyone can suggest a new fact or a better line from wherever they work. It waits in a queue until someone the company chooses approves it.
The same answer everywhere	Claude, ChatGPT and other AI assistants, plus Slack, Canva and other tools, connect to the same source. A rule changed once applies on the next request.

Tandem also hands over the real files, not descriptions of them. When someone asks for a deck, the AI starts from the company's actual template and approved logo instead of recreating them. And before anything goes out, a draft can be checked against the approved claims, with the exact phrases at issue and suggested wording.

Alongside your knowledge base, not instead of it

A company that has built a knowledge base or connected its files to an AI platform should keep it. It's the best way to find what the company has done: past projects, old decisions, who worked on what.

Tandem does a different job. The knowledge base answers “what have we ever written?” Tandem answers “what have we decided to say?” Used together, a team can pull the history of a project from one and describe it in approved, current language from the other.

Tandem also keeps its scope narrow on purpose. It holds brand and marketing material: the guide, files, templates, approved claims and rules. Project data, client figures and confidential records stay in the company's own systems and tools.

What it doesn’t do

A governed layer makes the approved version visible and available everywhere. It doesn't make anything automatic, and we'd rather say so plainly.

- It needs people. The brand understanding is drafted from the company's materials, and it's only as good as the people who review and lock it. The value is that it can be seen and corrected, not that it's right on its own.
- It steers the tools. It doesn't control them. A connected AI tool can still mix Tandem's guidance with other material in the same conversation, or drift from it. The compliance check and a person's review still matter.
- It supports review. It doesn't replace it. Legal, compliance and regulatory review stay with the company's own people. Drafts simply reach them cleaner.

Questions to ask of any company brain

Whatever a company uses today, these questions show whether it governs what the AI says about the company, or only who can see the files.

1. Can we read what the AI understands about us, in one place?
2. Can we correct it, and lock what must not change?
3. When two documents disagree, which one wins, and who decided?

4. Is every claim marked as approved or not, and does the AI respect the difference?
5. Who approves new knowledge before it's used?
6. Does every tool our team uses get the same version?
7. When we rebrand or retire a message, how long until every tool knows?

If the answer to most of these is “we'd have to check the folders,” the company brain is a search engine with good manners. That's useful. But it isn't governed.

Closing

AI tools will keep changing, and teams will keep adopting new ones. The question that won't change is who decides how the company is described. Today, in most companies, the answer is whatever the search finds.

Tandem puts that decision back with the company: one approved version, readable, correctable and the same in every tool the team uses. Your documents remain the record. Tandem is what you've decided to stand behind.

Build boldly. Stay the course.

Matthew Snyder, Founder and CEO, Tandem · builtintandem.com