

BEYOND RETRIEVAL

The Corporate Brain, Again — Why Retrieval Will Not Save Knowledge Management

Why the machine that can finally read the archive still cannot fix what broke knowledge management the first time

Giovanni Leonardi

December 2022



Contents

1. Executive Summary	3
2. The Oldest Ambition in the Building	3
3. The Graveyard We Keep Rebuilding	4
4. What Actually Changed This Winter	5
5. The Case That This Time Is Different	5
6. The Ghosts That Do Not Leave	6
7. Why the Dream Never Dies	7
8. The Brain We Actually Need	7

The Corporate Brain, Again — Why Retrieval Will Not Save Knowledge Management

Transformation · No. 290

Originally written December 2022 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

*“It has done nothing at all to make the knowing honest,
and it is the knowing that was always the hard part.”*

1. Executive Summary

For roughly three decades, organisations have chased the same ambition under a rotating set of names: capture what the enterprise knows, hold it in one place, and make it retrievable the moment someone needs it. Each generation of technology promised the corporate brain, and each left behind a graveyard — abandoned intranets, unread repositories, expertise directories pointing at people who left years ago. In the closing weeks of this year a new candidate has arrived with unusual force: systems that pair semantic retrieval with a language model able to read what it finds and answer in ordinary prose. Weeks after a research chatbot became a household curiosity, the same underlying capability is being wired into internal document stores, and the pattern is being greeted as knowledge management reborn.

This essay argues that the excitement is half right, and that the half it mistakes is the half that always mattered. Retrieval-augmented generation genuinely dissolves one of the two great failures of the old discipline — the retrieval problem, the labour of finding and phrasing what the archive already holds. But knowledge management never failed only at retrieval. It failed at the corpus, because what was written down was partial, stale, and self-contradictory; it failed at the tacit layer, because the most valuable knowledge was never in any document; and it failed at trust, because people will not lean on a source whose reliability they cannot judge. None of those failures is touched by the new architecture, and generation introduces a fourth of its own: the fluent, confident synthesis of material that is wrong, superseded, or contradictory — a failure mode more corrosive than the blank silence it replaces. The corporate brain was never a retrieval problem waiting for a better index. It was, and remains, a problem of curation and trust, and the machine that can finally read the archive has moved that problem without solving it.

2. The Oldest Ambition in the Building

Picture the scene that is playing out, this month, in a handful of organisations bold enough to try it early. A question is typed into a plain box — what is our position on

holding customer records offshore? — and nine seconds later a fluent paragraph comes back, citing three internal documents by name. The people in the room, who between them have spent a decade watching enterprise search return forty blue links and no answer, go quiet. Something has changed, and everyone can feel it.

What has not changed is the ambition behind the box. The wish to build a single place that knows what the organisation knows is one of the oldest in corporate life, older than most of the technology that has been thrown at it. It survives because the pain it addresses is real and permanent: knowledge walks out of the door every Friday, sits trapped in the heads of a few overloaded experts, and has to be painfully rediscovered by each new joiner. Every few years a new tool arrives promising to end that waste, and every few years we believe it. The question worth asking, before we believe it again, is why the previous attempts are buried in the same graveyard — and whether this arrival is different in kind or merely in polish.

3. The Graveyard We Keep Rebuilding

The discipline that called itself knowledge management crystallised in the mid-nineties, built on a genuinely useful distinction between explicit knowledge, which can be written down, and tacit knowledge, which lives in judgement and habit and resists capture. The strategy that followed was seductive in its simplicity: convert the tacit into the explicit, store the explicit, and retrieve it forever. Two decades of intranets, wikis, lessons-learned databases, communities of practice and expertise locators were built on that premise. A wave of organisations even appointed a chief knowledge officer to preside over it.

Most of it failed, and it failed in ways worth naming precisely, because the new tools will inherit whichever of these we do not fix.

1. The capture tax. Writing knowledge down is work, and the person who holds the knowledge is almost never the person who benefits from recording it. So it does not get recorded, or it gets recorded thinly, in the last rushed hour of a project when the budget is already spent.
2. Content rot. What was captured aged. A repository of forty thousand documents accumulated over fifteen years is not an asset; without relentless curation it is a sediment, in which the current policy and the policy it replaced sit side by side with equal authority and no visible date of death.
3. The tacit residue. The most valuable things the organisation knew — why a particular integration always failed, which supplier could be trusted under pressure, how a deal really got done — were exactly the things that never made it into a document, because they were too contextual, too political, or too obvious to their holder to seem worth writing.
4. The trust deficit. Even when the answer was in there, people rang a colleague instead. A trusted human source came with provenance: you knew who they were, how current they were, and how much to discount them. The repository offered a paragraph with none of that, and so it was not believed.

Enterprise search, when it came, was sold as the cure and treated only the first symptom. It made the pile marginally easier to rummage through; it did nothing about the pile being stale, incomplete, and unjudgeable. The lesson, paid for many times over, was that

retrieval was never the binding constraint. The archive's problem was not that we could not find things in it. It was that too much of what we found was not worth finding.

4. What Actually Changed This Winter

It is important to be exact about what the new approach does, because its genuine advance is easy to overstate and easy to dismiss, and both errors lead somewhere false.

The mechanism is, at heart, a pipeline. Documents are broken into passages and each passage is converted into a numerical fingerprint — an embedding — that captures something of its meaning rather than its mere words, using the kind of cheap embedding endpoints that have only this quarter become widely available. A question is fingerprinted the same way, the passages whose fingerprints sit closest to it are pulled back, and those passages are handed to a language model with an instruction of the form: answer this using only what follows. The model then does the thing that is new — it reads the retrieved material and composes a direct, phrased answer, rather than leaving the human to read forty documents and synthesise for themselves.

Set against the graveyard, two things are genuinely different. The first is that the burden of phrasing has shifted from the reader to the machine; the nine-second paragraph is real, and it is not nothing. The second is subtler and more important: the system reads a corpus that already exists rather than demanding that people write a new one. For the first time, the capture tax is not levied at the point of use. The documents an organisation already has — the policies, the tickets, the transcripts, the wiki nobody loved — can be pointed at directly.

The old dream asked people to pour their knowledge into a system before it could be useful. The new architecture reads what is already there. That single inversion is why this feels different — and why it is tempting to conclude that the discipline's oldest failure has finally been engineered away.

That temptation deserves the strongest possible hearing before it is answered.

5. The Case That This Time Is Different

The optimist's argument runs like this, and it is not a straw man. Knowledge management's failures, on this view, were overwhelmingly failures of the interface. The knowledge was in the building; it was simply locked behind bad search, rigid taxonomies, and the impossible demand that busy people file their expertise into someone else's schema. Remove those frictions — let anyone ask anything in their own words and receive a composed answer drawn from the real corpus — and adoption follows, because for the first time the system is easier than ringing a colleague. And adoption was always the true bottleneck; every previous tool died of neglect, not of some deep conceptual flaw. The capture tax, the optimist adds, falls away precisely because we no longer need people to write new summaries: the model summarises on demand from whatever raw material

exists. Even the tacit layer erodes at the edges, because informal traces — the long message thread, the meeting transcript, the half-finished note — are now legible to a system that can read prose, where before only formally filed documents counted.

This is a serious case. It correctly identifies that friction killed the previous generation, and it correctly sees that reading an existing corpus is categorically better than commissioning a new one. If the whole problem were retrieval and phrasing, the optimist would simply be right, and this essay would be a celebration. The difficulty is that the whole problem was never retrieval and phrasing.

6. The Ghosts That Do Not Leave

Return to the four failures and ask which of them the new architecture actually touches. It resolves the first — the phrasing and finding — handsomely. It does nothing for the other three, and it worsens the situation in a way the old tools never could.

The corpus problem is untouched, and arguably made more dangerous. Consider the forty-thousand-document store again, and suppose an honest audit finds that roughly a third of it is superseded, a further tenth is duplicated across shared drives, and a long tail of the remainder was written by people no longer at the organisation and never reviewed since. Enterprise search, faced with that store, returned the superseded policy as one link among many, and a human reader often noticed the smell of it — an old template, a defunct department in the footer, a date. The generative system retrieves the same superseded passage, finds it a perfect semantic match for the question, and renders it into a crisp, confident, present-tense answer with the old document's date nowhere in view. The staleness that was merely present in the old world is now laundered into authority.

The tacit residue is equally untouched. If the reason a particular integration always failed lived only in the head of an engineer who left in the spring, it is not in the ticket system for the model to read, and no amount of fluent retrieval will conjure it. The system can only be as wise as the corpus is, and the corpus is precisely as incomplete as it ever was. What has changed is that the gaps are now invisible: an empty archive returns a confident answer built from whatever adjacent material it could find, where before it at least had the honesty to return nothing.

And trust — the deepest of the failures — is inverted rather than earned. The old repository was distrusted and therefore underused, which was wasteful but safe. The new system is fluent and therefore over-trusted, which is the more expensive error. A composed paragraph in confident prose carries an authority its sources have not necessarily earned, and it strips away exactly the provenance a human source supplied — who said this, how current are they, how much should I discount it.

The failure	Enterprise search 1.0	Retrieval-augmented generation
Finding and phrasing	Partly solved	Solved well

The failure	Enterprise search 1.0	Retrieval-augmented generation
Stale, contradictory corpus	Exposed, human often notices	Laundered into confident prose
Tacit knowledge never written down	Absent, visibly	Absent, invisibly
Trust and provenance	Under-trusted, so underused	Over-trusted, so over-relied-upon

This is the fourth failure, the one that belongs to the new era alone. The old systems failed by silence; they gave you nothing, and you knew it. The new systems can fail by fluency — synthesising the superseded, the contradictory, and the simply absent into a single smooth answer that betrays no seam. There is a fresh and half-understood worry, too, that a system which answers from whatever text it reads can be steered by that text, so that a stray instruction buried in a document becomes an instruction to the machine. We do not yet know how large that risk is. We know only that confident wrongness, delivered at the speed and polish of these systems, is a more dangerous thing to put in front of a busy decision-maker than an empty results page ever was.

7. Why the Dream Never Dies

If the ambition keeps failing, why does it keep returning? Because the underlying loss is real and the seduction of the single answer is powerful. Organisations do haemorrhage knowledge; experts are a genuine bottleneck; onboarding is genuinely slow and wasteful. Any tool that promises to end that pain will always find an eager buyer, and the more fluent the promise, the more eager the buyer. There is also an institutional force at work: a fluent demonstration is extraordinarily easy to sell upward. A single impressive answer in a boardroom does more to unlock a budget than a year of honest talk about corpus hygiene ever could, and so the discipline is pulled, again, toward the demo and away from the drudgery that actually determines whether it works.

That is the gap between transformation intent and transformation reality in its purest form. The intent is a brain that knows what the organisation knows. The reality that will determine success or failure is unglamorous and almost entirely non-technical: someone must decide what belongs in the corpus and what must be retired from it; someone must mark which documents are current and which are dead; someone must own the answer to how do we know this is right. The technology has genuinely changed. The work has not.

8. The Brain We Actually Need

The honest conclusion is neither the celebration nor the dismissal, but something harder to sell: the new architecture is a real advance on one axis and a real hazard on three others, and which of those dominates in any given organisation is decided entirely by the quality and governance of the underlying corpus — the very thing the excitement invites

us to ignore.

We are being offered a better way to read the archive at the precise moment we are most tempted to stop tending it. The organisations that will get value from this are not the ones that wire a model to their document store fastest. They are the ones that treat the model as a reason to finally do the unglamorous work the old discipline always required and rarely got: curating a corpus worth reading, retiring what is dead, capturing the tacit knowledge that still lives only in people, and — above all — building the provenance and the habits of scepticism that let a fluent answer be trusted the right amount, which is to say, not completely.

“The machine can now read everything we have written down. That is exactly why it has never mattered more what we chose to write, what we let rot, and what we still keep only in our heads.”

The corporate brain, if it is ever built, will not be built by retrieval or by generation. It will be built by the patient, thankless, deeply human work of deciding what is worth knowing and keeping it true. The technology has, at last, made the reading effortless. It has done nothing at all to make the knowing honest, and it is the knowing that was always the hard part.