THE RACE FOR CORPORATE MEMORY

Billions invested in enterprise AI produced no return, layoffs are erasing the knowledge that would make the technology work — and a new category of companies has emerged to capture what no system ever recorded.

Whoever gets there first takes the asset that is hardest to copy in the corporate world.

There is one number that sums up the strangest moment in enterprise artificial intelligence: 95%. That is the share of generative AI pilots which, according to an MIT study published in 2025, produced no measurable impact on the results of the companies that bought them — after an accumulated investment estimated between 30 and 40 billion dollars. The study, run by Project NANDA across 300 public deployments, 52 interviews inside organisations and 153 executive responses, demolished the market's favourite explanation. The difference between the 5% that worked and the rest was not model quality, nor regulation. It was something more prosaic: the tools that failed retained no context about the operation they were dropped into. They shone in the demo and broke on Monday.

MIT's closing recommendation to executives fit in a single line: fund the memory layer — what does not retain context does not scale.

Meanwhile, in the same year they spent billions on tools with no memory, companies were erasing the memory they already had. More than 90,000 jobs were cut in the technology sector in 2026, a growing share attributed to AI restructuring. In February, the outplacement firm Careerminds published a survey of 600 human resources professionals: **two thirds of the companies that laid people off because of AI in the past year say they are already dealing with the consequences.** The consultancy IDC estimates that failures in knowledge transfer cost 31.5 billion dollars a year to Fortune 500 companies alone.

The two numbers tell the same story from opposite angles. The context missing from AI pilots and the knowledge walking out the door in layoffs are, to a large extent, the same thing — and it has had a name since 1966, when the philosopher Michael Polanyi defined it in six words: we know more than we can tell. Tacit knowledge is the judgement applied without being able to explain the criterion, the informal adjustment that holds a deadline together, the never-written reason behind a recurring exception. In March 2026, Berkeley's *California Management Review* raised the concept to a strategy headline: tacit knowledge, the magazine wrote, is the next competitive advantage of the age of AI agents.

Where there is named pain and billions badly spent, a category appears. This one has just been born — and already has three paths competing to reach corporate memory first.

THE MINING PATH

The loudest move came from Munich. In March 2026, Interloom announced a 16.5 million dollar round led by DN Capital, with participation from Bek Ventures and Air Street Capital, for a thesis its founder Fabian Jakobi sums up with one figure: around 70% of operational decisions were never formally documented. The company's bet is to reconstruct those decisions by mining what already exists — support emails, tickets, call transcripts — and turning them into a “context graph” that feeds AI agents.

The announced clients are impressive: Zurich Insurance, where the startup won an internal competition against roughly 2,000 contenders; Volkswagen; the property firm JLL; the logistics company Fiege; and Commerzbank, owner of the category's most cited case — analysing millions of emails against internal documentation, the gap between the documented and the practised process reportedly fell from around 50% to 5%. Jakobi likes to provoke the consulting market: the Zurich underwriter, he told the press when announcing the round, knows how the process works better than any outside consultancy ever will.

THE EXTRACTION PATH

In Texas, Matt Cretzman's Skill Refinery attacks from the opposite side of the table: instead of mining the operation's trail, it extracts method from what specialists have already produced. A 200-page book becomes 30 to 80 structured “skill cards”, delivered inside the tools the professional already uses — ChatGPT, Copilot, Claude — with the specialist monetising their own name under a revenue-share model. It is the category's answer to another fear of 2026: professionals watching AI absorb their craft would rather package and sell their own expertise before it is simply copied.

BOX 1 — THE NUMBERS OF THE RACE

95%
of generative AI pilots with no measurable return *(MIT, 2025)*

US$ 30–40 bn
invested up to the measurement *(MIT)*

2/3
of companies that cut staff because of AI already face the consequences *(Careerminds, Feb 2026)*

US$ 31.5 bn/yr
lost by the Fortune 500 to knowledge transfer failure *(IDC)*

US$ 16.5 m
in the category's largest round *(Interloom, Mar 2026)*

>40%
of agentic projects cancelled by 2027 *(Gartner)*

~130
genuinely agentic vendors, among thousands *(Gartner)*

THE THIRD PATH — AND THE SLICE THE OTHER TWO CANNOT REACH

Both models share a silent dependency: they only work on what, at some point, became a digital artefact. Mining reads emails and tickets that exist; extraction structures what the specialist chose to declare and sign. But the most valuable slice of Jakobi's "70% undocumented" never produced any artefact at all — it was decided in a corridor, agreed on the phone, adjusted by a criterion nobody put into words. You cannot mine what does not exist.

It is in that void that the Brazilian company Zthex operates, with a mechanism that inverts the previous ones: instead of reading trails, it asks — through structured conversation with the people who do the work, in cycles of up to five weeks at a fixed price. The technical difference its founders insist on is the origin of the data: information mined from a ticket is born with an owner — sender, assignee, attribution — and anonymising it afterwards is reconstruction. In conversation, identity is removed before any recording; in the company's systems, they state, the identity field simply does not exist. What was said is then compared against two other sources — what policy states and what the systems record — and every divergence among the three becomes a finding with a traceable origin, delivered as a question to investigate, never as a verdict. At the end, the company is left with a "Divergence Index": the share of its processes where the documented, the recorded and the lived do not match.

The contrast in philosophy is visible even in what companies display. While the category competes over client logos on its websites, Zthex publishes none — by declared policy: whoever hires an investigation into their own operation, its founders argue, is buying discretion, including discretion about having hired at all.

WHO TAKES IT

The category is too young to have a winner, and mature enough to have an investment thesis: DN Capital justified its cheque to Interloom with the sentence that sums up the race — an AI agent is only as good as the expert decisions it can lean on. Gartner adds urgency from the destruction side: more than 40% of agentic AI projects will be cancelled by the end of 2027, and among the thousands of vendors calling themselves agentic, the firm estimates around 130 actually are.

What the three paths have in common is the underlying bet: over the next decade, a company's edge will not sit in the models — which will be commodity, ever better and identical for everyone — but in the proprietary memory each organisation manages to structure before it evaporates through layoffs, retirements and reorganisations. Or, as Zthex's founders chose to carve into their manifesto: everyone buys the same intelligence; the game is building the one you can't.

Z END OF FEATURE

BOX 2 — THREE PATHS TO THE SAME MEMORY

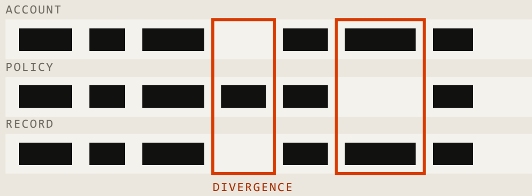

/* the three sources along a process: where they do not coincide, there is a finding. A representation of the method, not of client data. */

Mining

(INTERLOOM)

Reads existing emails, tickets and transcripts. **Strength:** scale over a dense corpus. **Limit:** cannot reach what never became a record; identity is born attached to the data.

Extraction

(SKILL REFINERY)

Structures what specialists have published and declared. **Strength:** speed over ready material. **Limit:** captures the explicit and signed, not the tacit.

Conversation

(ZTHEX)

Asks the people who do the work, anonymises before recording, compares with policy and record. **Strength:** reaches what left no trail. **Limit:** depth over volume — closed cycles, one client at a time.

Editorial content produced by Zthex. Facts about third parties come from public sources cited in the text.

ABOUT ZTHEX

Zthex structures the layer of knowledge no system shows. It talks to the people who do the work, removes identity before recording, and compares what was said with what policy states and what the systems record. Every divergence becomes a finding with a traceable origin.

PRESS

Press kit, images and additional reporting:
zthex.com/en/imprensa
info@zthex.com

ZTHEX

ZTHEX.COM
point your camera