

Silicon Valley just named what we do

Y Combinator named the category: *Company Brain*. Four paths are already under construction — and the question that decides the product's final shape has not been asked out loud.

The accelerator that spotted Airbnb, Stripe and OpenAI before anyone else has just pointed, officially, at the problem this company is built on. The name it gave is “*Company Brain*”. The question nobody is asking is: a brain built how, and to serve whom?

In May 2026, Y Combinator published its Requests for Startups — the problems the accelerator most wants the next founders passing through it to attack. Fourth on the list, partner Tom Blomfield named the category: ***Company Brain***. The official definition: a system that pulls knowledge from fragmented sources, structures that knowledge, keeps it current, and turns it into an executable layer — for people and for AI agents to read and act on.

YC was explicit about what this is **not**. Enterprise search and RAG retrieve text by similarity; they find the document, but they neither prove a relationship between data points nor guarantee that two agents are using the same definition of a metric. A real *Company Brain*, according to the accelerator itself, is “a living, governed model of what the organisation knows” — not a search index with a nice name.

The reason for the interest has a documented price: according to Gartner, poor-quality data costs a typical organisation an average of US\$ 12.9 million a year. Thomas Redman, a researcher the *MIT Sloan Management Review* has published on this subject for years, estimates that bad data consumes between 15% and 25% of most companies' revenue. YC is not betting on a passing trend. It is betting on an invoice that has already arrived.

THE WAYS IT IS ALREADY BEING BUILT

What is interesting about a newly named category is that it reveals, all at once, every bet that was already under way without a common name. Looking at what has already emerged inside YC's own current batch, at least four distinct paths appear — and each deserves to be mapped fairly.

Digital-trace aggregation. Startups such as Cerenovus attack the problem by pulling company files, email, Slack messages, spreadsheets and meeting minutes into a single structured knowledge graph. It is, in essence, the same logic as the mining movement that already exists outside the Valley — read what the organisation has already produced digitally, and organise it.

Structured field capture. Kebra solves another piece of the same problem in another sector: field technicians are terrible at documenting their own work, and that jams the back office. The solution captures what happens on the spot, before the truck leaves, and turns it into a complete note, a warranty claim and a parts order. The phrase the company itself uses is revealing: every job feeds the “*Company Brain*”, so that the best technician's expertise never leaves the company when he does.

External-signal verification. Poth Labs calls itself, literally, “Poth Company Brain” — an agent that looks for truth inside customer feedback, trying to separate real signal from noise in the voice of whoever buys.

System integration with a declared human touch. Bizmark connects systems, inboxes, spreadsheets and documents — and describes, in its own words, also capturing “the knowledge that lives in people's heads”. It is the formulation closest to the problem this text discusses, but the product, by its own public description, starts from system integration inward, with conversation as a complement — not as the central method.

Four companies, four bets, one problem in common, and none of them — on the public material available up to this analysis — places structured conversation with the people who live the work as **the primary mechanism**, with anonymity built by architecture before any recording. All of them start from the trace the organisation has already left; conversation, when it appears, is an accessory.

WHAT NOBODY HAS NAMED YET: WHOSE BRAIN IS IT

There is a question the entire category has not yet asked out loud, and it decides everything about this product's final shape: **does a “company brain” that learns by mining email, Slack and tickets learn about the people who work there, with identity attached to the data, or does it learn about the operation, with identity removed before any record exists?**

The difference is not a technical detail. It is about who the system was designed to protect.

Much has been said, over the past two years, about building “AI-first” companies — automate first, decide later who is left. This company chose the opposite path, and says so without hedging: **here, the human comes first, always — and artificial intelligence exists, among other things, to protect that priority**, not to work around it. Anonymity is not a policy setting; it is a property of the data, decided before any line of code records anything. The rule of backing off in the face of discomfort is not optional courtesy; it is a rule without exception, even when the subject seems important. No individual answer is visible to anyone, not even to the client who commissioned the cycle — because the field that would allow it simply does not exist in the schema.

It is not the only valid way to build a *Company Brain*. The four approaches mapped above solve real parts of the same problem, for different audiences and sectors, and deserve to be recognised for it. But there is a slice of organisational knowledge that only the conversation mechanism, with anonymity from origin, can reach: the decision that never generated an email, the adjustment that never became a ticket, the reason nobody would put their own name under. When an entire category measures success by how much trace it can mine, that slice — usually the most valuable, because it is the most candid — is left out by definition.

WHY TEN, NOT TEN THOUSAND

There is a second question the whole *Company Brain* category is asking the wrong way: everyone asks how to scale the product. Almost nobody asks whether scaling the product is the right way to scale the impact.

Zthex serves at most ten organisations a year. It is not a technical limitation, nor shyness about growth — it is a business model, chosen on purpose, and there is strong historical precedent for it working.

In 1945, the Bell Telephone research laboratory gathered a handful of scientists under the direction of William Shockley — among them John Bardeen and Walter Brattain. The group never tried to become an electronics factory at scale. In December 1947, it built the transistor — and let the entire world scale on top of it. The three shared the 1956 Nobel Prize in Physics for the discovery. Almost every electronic device on the planet today descends from that small room, which never needed to grow in order to change everything.

In 1970, Xerox opened a research centre in Palo Alto — PARC. Out of it came the graphical interface with windows and icons, the mouse and the Ethernet network, all running on the Alto computer as early as 1973. In December 1979, Steve Jobs visited the laboratory — part of an arrangement in which Xerox bought pre-IPO Apple shares in exchange for the access — and saw there what would become the visual foundation of the Macintosh. Xerox did not scale those ideas as a mass-market product; Apple did. The story is not as simple as “Xerox did not realise what it had” — it is more accurate to say Xerox chose not to bet its own core business on it, and let the idea multiply through the hands of whoever did.

In agronomy, the example is even more to the point. Norman Borlaug spent decades working with a small number of selected wheat-breeding programmes — first in Mexico, then extending the method to India and Pakistan at moments of real and imminent famine. He never tried to plant the world's wheat himself. He chose where hunger was most urgent, handed the method to local governments and research institutes, and let those partners scale the harvest. The “Green Revolution” that resulted is credited, according to the Nobel Committee itself, with having helped avert mass famine for hundreds of millions of people — some estimates cite more than a billion. Borlaug received the Nobel Peace Prize in 1970 — not for having grown the world's wheat, but for having chosen, with extreme care, who would receive the method first.

The contemporary example closest to the Zthex model is in the largest-scale philanthropy on the planet. The Gates Foundation does not deliver vaccines door to door. It helped found and fund Gavi, the Vaccine Alliance — a partnership that works through a relatively small number of very-high-leverage partners: national governments, the World Health Organization, UNICEF. Since 2000, that model has immunised more than 1.2 billion children and averted more than 20.6 million future deaths, according to Gavi's own figures — without the Gates Foundation having grown, as an organisation, in proportion to the impact generated. The size of the foundation is not the size of the result. The size of who it chooses is.

The pattern in all four cases is identical, and it is the same one that governs Zthex: **the small organisation does not try to become large. It chooses, with rigorous criteria, who is already large enough that the depth of a single well-done piece of work multiplies on its own afterwards.**

WHAT THIS MEANS IN PRACTICE, FOR ZTHEX

Ten clients a year is not a ceiling on ambition. It is the same arithmetic as Bell Labs, PARC, Borlaug and the Gates Foundation, applied to organisational knowledge: a multinational that goes through a Zthex cycle does not keep the result to itself. The knowledge captured shapes how that company trains new hires, how it decides its next investment, how it treats thousands of employees and millions of customers who will never know the name of the company that helped unlock that process.

Zthex is not trying to be heard through its own scale. It is trying to be heard through the scale of whoever it chooses to serve. Each selected client is not a sale — it is a leverage point. That is why choosing who enters a year's calendar is the company's most strategic decision, more important than any product decision.

// none of the four examples above grew in order to change the world. What grew was whoever they chose to let grow with what they created.

THE INVITATION, AGAIN

Silicon Valley has just validated, with the institutional weight only it carries, that the problem this company has worked on since before it had that name is real, urgent, and expensive enough to justify fifteen of the most coveted bets of a generation of founders.

Zthex chose to build that organisational brain by the path that begins in conversation, not in the trace — and by the rule that the artificial intelligence taking part in it exists to protect whoever speaks, never to replace them. It is one bet among several valid ones. It is the one we make, convinced that a company brain that cannot keep a secret for whoever trusted it is not a brain — it is just another database with a nice name.

None of the stories above grew in order to change the world. What grew was whoever they chose to let grow with what they created. That is the bet behind the limit of ten — it is not the size of Zthex that matters. It is the size of whoever Zthex helps to see its own truth first.

Z END OF ARTICLE

BOX – WHY DEPTH, NOT VOLUME

Bell Labs, 1947	a handful of scientists invent the transistor; they do not manufacture electronics at scale. The world scales on top of the invention.
Xerox PARC, 1970s	a small research group invents the graphical interface; Apple and Microsoft scale the product.
Norman Borlaug, 1960s–70s	a wheat method handed to a few countries in crisis; local partners scale the harvest. Nobel Peace Prize, 1970.
Gates Foundation / Gavi, since 2000	1.2 billion children immunised and 20.6 million deaths averted since 2000 (<i>Gavi</i>); partners scale the vaccination.
Zthex, today	up to ten cycles a year; each selected client scales the result to thousands of people who will never know this company's name.

// map built from public material published by the companies cited, with no judgement about which approach is “better” – each solves a real piece of the problem.

Editorial content produced by Zthex. Facts about third parties come from public material published by the companies and organisations 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

Signed republication and additional reporting:

zthex.com/en/imprensa

info@zthex.com



ZTHEX.COM
point your
camera