

# The Company Brain for the Enterprises That Power the Real Economy

BRUNO CALABRETTA · EMILIO CANESSA · JUNE 2026

## 1. The Infrastructure Gap

Enterprise artificial intelligence initiatives fail at a documented rate. RAND's 2025 meta-analysis of sixty-five enterprise AI initiatives found that 80.3% failed to deliver intended business value, with 33.8% abandoned before reaching production and 28.4% reaching production but returning no value.<sup>1</sup> Deloitte's 2026 State of AI in the Enterprise reports that 42% of companies abandoned at least one AI initiative during 2025, at an average sunk cost of 7.2 million dollars per abandoned initiative.<sup>2</sup> The failure pattern holds across initiative type, industry, and company size. The recurring variable is not model selection or budget. It is the absence of a structured substrate for the models to operate against.

Three structural causes account for the pattern. The first is knowledge fragmentation. Operational knowledge sits across email threads, CRM records, ERP exports, shared drives, meeting transcripts, and individual memory, in no form that AI agents can read, navigate, or act on. The data exists; the substrate does not. The second is the absence of a path to self-assembly. Building the substrate in-house requires AI engineering capability, data integration work, and ongoing maintenance discipline that most organisations are not staffed for and cannot staff for within a relevant timeframe. The third is vendor lock-in as operational risk. Loading operational memory into a third-party knowledge base means renting the company's working intelligence from a provider whose pricing, feature roadmap, and ownership can change independently of the customer. Repricing and deprecation cycles through 2025 and 2026 confirmed this risk to be concrete.

Company Brain, as used throughout this document, refers to the architectural category in which a central structured store of a company's operational knowledge is kept current and made executable by AI agents. Between April and May 2026, the category acquired an institutional label through a convergence of independent reference implementations. On April 5, Andrej Karpathy, formerly of Tesla and OpenAI, published a personal gist describing roughly one hundred articles and four hundred thousand words of his own knowledge held in a markdown-in-git substrate that his coding agents read alongside him.<sup>3</sup> His framing was deliberate:

markdown in a git repo or a folder. Plain text that humans and machines can both read. Not a SaaS database. Not a wiki product.

Four further builders published implementations of the same construct within the following weeks. Garry Tan, President and CEO of Y Combinator, open-sourced GBrain and GStack under MIT with forty-three curated skills, on the principle of "a thin harness, fat skills."<sup>4</sup> Hannah Stulberg at DoorDash described Team OS, a shared repository where Product, Engineering, Design, and Analytics commit knowledge into a single surface that agents and humans read from equally.<sup>5</sup> Ramp, the 32-billion-dollar fintech, reported that its internal system Glass had reached 99% daily AI adoption across more than 350 skills, with its Head of Product, Seb Goddijn, stating the logic plainly: "Internal productivity is a moat."<sup>6</sup> On April 30, Alex Lockey synthesised the five canonical elements those builders had independently arrived at: a central structured store, a routing layer, skills as the unit of work, a discipline of writing results back into the substrate, and the separation of context from compute.<sup>7</sup> Two weeks later, Y Combinator's Summer 2026 Request for Startups named the category in two adjacent slots of fifteen.<sup>8</sup>

That five independent builders, across different industries and company scales, arrived at the same five-element architecture constitutes convergent evidence that the pattern is structurally correct. Each had the engineering depth to assemble the elements from scratch. Most businesses do not, which is the source of the gap between the readiness of the pattern and the readiness of the market. Tom Blomfield, Monzo founder and YC Partner, stated the missing layer in the May 2026 Request for Startups:<sup>8</sup>

The biggest blocker to AI automation of companies is no longer the models, they just got so good so quickly. Now the blocker is the domain knowledge. A system that pulls knowledge out of every fragmented source, structures it, keeps it current, and turns it into an executable skills file for AI.

## 2. The Architecture

The Company Brain architecture comprises five elements. No deployable product on the market implements all five. SALTZ AI maps one module to each canonical element. Each module is independently deployable and immediately useful. Together they form a complete operating brain that the customer owns, controls, and retains regardless of which AI model runs against it.

**Memory** is the central structured store: plain-text, versioned files describing the business, its team, its strategy, and its active workstreams, owned entirely by the customer. It is the substrate every agent reads from on every session, corresponding to the first canonical element in Lockey's synthesis.<sup>7</sup>

**Focus** is the routing layer: the configuration files for frontier models together with a documentation index at the repository root. Every agent session loads the relevant context automatically, so operational knowledge is applied consistently across tasks without manual intervention.

**Skills** is the executable surface where recurring work is encoded: how the company qualifies a customer, prepares a proposal, onboards an account, or produces its weekly intelligence digest. Skills are written once and compound. The library grows with the business without growing the headcount required to run it.

**Data** is the live data layer: daily collectors pull from the company's sources, including CRM, spreadsheets, financial feeds, and internal tools, into a local structured warehouse, then write a generated context file back into the substrate. The agents operate against current data. This addresses one of the most consistent failure modes in enterprise AI deployment, in which models operate against stale or disconnected information.

**Attention** is the signal layer: meeting transcripts and recorded conversations are classified by client stream and written back into the substrate. The agents track what is happening across the business in real time because the human working record and the agent working record are the same file.

The architecture rests on two properties that are product decisions rather than implementation details. The first is model independence. The substrate is plain text in a git repository and runs against one frontier model today, a different model next year, and a customer-deployed open-weight model in a Confidential VM thereafter. The model changes on a roughly six-month cycle; the substrate does not. The second is confidentiality by default. Frontier-model inference routes through AWS Bedrock with Nitro Enclaves for Claude and Azure Confidential Computing on NVIDIA H100s for GPT-class models, so the provider operator cannot read customer prompts. No European Company Brain product ships this configuration as a default at present.

The substrate is delivered through a product interface oriented toward two access points: conversational interaction through existing messaging channels, and a visual layer for navigating the knowledge map and operating the agents. The interface is the access layer; the substrate is the asset.

Current competitors in the European market each address one or two of the five canonical elements. SALTZ AI implements all five. The position is one of architectural completeness, stated as architecture.

The platform is in active deployment. Production use confirms the architectural assumptions described in this section.

### 3. The Market Condition

The customer profile is defined by five criteria. The first is company size: 50 to 500 employees, the band in which the infrastructure gap is structurally unavoidable. These companies carry enough operational complexity, in departmental structure, multiple SaaS systems, recurring decision cycles, and tacit knowledge concentrated in a small number of people, that a Company Brain produces immediate and measurable value, while remaining too small to staff an in-house AI engineering function and too large to operate without one indefinitely. The second is digital maturity: an ERP and a CRM and three or more SaaS tools already running, with AI integration stalled. The data and processes exist in fragmented form; the architecture to structure them does not. The third is operating profile: family-managed or founder-led, where operational decisions concentrate at the top and institutional knowledge lives in a small number of people. The fourth and fifth criteria are sector and confidentiality requirement: knowledge-heavy back offices with high-volume recurring decisions, in work where confidentiality is a professional or contractual requirement.

The initial geographic wedge is Italy and Italian-speaking Switzerland. The founding team has native language fluency, cultural fluency, and direct distribution access in this corridor, where the family-managed and founder-led archetype is dominant. The expansion logic follows from the architecture rather than from geographic preference. The five modules are language-agnostic and sector-configurable, so a skill library validated for one national manufacturing sector applies to another with sector-specific calibration rather than reconstruction. The wedge establishes the pattern; the architecture extends it across Southern Europe and the broader European Union as the sector-specific skill libraries and onboarding architecture are proven.

Five timing vectors mark April and May 2026 as a documented convergence rather than a forecast. Each is a discrete, datable event.

**One.** The models crossed the agentic threshold. During the first quarter of 2026, the Claude 4 family, GPT-5, and Gemini 2 each became reliably capable of multi-step tool use, retrieval, and structured output without per-task prompt engineering. The constraint moved from inference quality to context architecture, a shift identified independently by multiple practitioner communities within the same quarter.

**Two.** Y Combinator named and committed to the category. Two adjacent slots of fifteen in the Summer 2026 Request for Startups represent a capital and reputational commitment from the most consequential early-stage investor to companies building in this space. Diana Hu's framing in the second slot was precise: "AI Operating System for Companies. The next generation of company infrastructure is being built on top of LLMs."<sup>8</sup> Institutional naming of a category by Y Combinator has historically preceded capital concentration and enterprise buyer awareness within twelve to eighteen months.

**Three.** The build-versus-buy question became answerable. The reference implementations described in Section 1 all entered the public domain inside the same thirty-day window. Until that point, an enterprise evaluating an in-house build had nothing concrete to scope against. The simultaneous arrival of open implementations let teams measure the assembly and maintenance effort directly, which shifted the decision from speculation to a costed comparison and turned demand toward a deployable alternative.

**Four.** Proof-of-concept risk was removed for buyers. A working Company Brain at company scale is now publicly documented and named, as set out in Section 1. For an enterprise buyer or an investment committee, that converts the category from an unproven concept into a referenceable precedent, which is the condition under which budget for this class of system is approved rather than deferred.

**Five.** Vendor lock-in became a recognised procurement risk. Through 2025 and 2026, enterprise SaaS platforms repriced, deprecated features mid-cycle, and changed ownership in ways that disrupted customer workflows. The pattern was consistent enough to shift buyer preference: owning the operating brain rather than renting it moved from a technical preference among engineering teams to a procurement requirement at the commercial level.

These five vectors are global. The Italian-speaking market provides a regional quantification of the same condition. Italian enterprise AI adoption tripled across three years, from 5% in 2023 to 8.2% in 2024 to 16.4% in 2025.<sup>9</sup> In Switzerland, 67% of businesses plan to integrate at least one AI tool by the end of 2026, while only 18% report a structured implementation plan.<sup>10</sup> The 49-point gap between declared intent and execution capability describes the addressable market condition the platform is designed to resolve.

---

## 4. The Road Ahead

The architecture compounds. Each skill encoded, each process structured, and each decision captured increases the substrate's specificity to the business that holds it. The customer does not consume the product but builds a permanent asset that grows more complete and more operationally decisive over time.

Expansion is determined by the architecture. The five modules are language-agnostic and sector-configurable, so a skill library validated for one sector and language transfers to adjacent sectors and languages through calibration rather than reconstruction. The initial corridor establishes the pattern, and each deployment produces a validated skill library, a refined onboarding architecture, and a referenceable case study that lowers the cost of the next deployment.

Tom Blomfield's formulation in the Summer 2026 Request for Startups, that every company in the world is going to need a Company Brain, describes an architectural inevitability rather than a market prediction. The models are capable, the pattern is validated across independent implementations at scale, and the remaining variable is the availability of a productised, deployable form of the architecture for the businesses that cannot build it themselves. That is what SALTZ AI provides.

## 5. An Invitation

**To enterprise leaders.** If your organisation runs ERP, CRM, and a stack of SaaS tools and cannot make AI work profitably against any of them, the structural cause is the absence of a substrate for AI to operate against. SALTZ AI builds that substrate: a Company Brain that encodes how your business works, keeps that knowledge current, and puts AI agents to work on the recurring decisions that consume the most time. The substrate is yours, runs on your infrastructure, and compounds with every deployment.

**To investors.** SALTZ AI is at the earliest stage of a category validated simultaneously by Y Combinator, Ramp, and the open-source community in spring 2026. The platform is in production, the founding team has four years of shared execution at scale, and the commercial build is underway.

**To potential partners.** SALTZ AI is building the Company Brain for mid-market business, with the sector depth, confidentiality architecture, and distribution relationships that market requires. If you are building adjacent infrastructure or operating distribution networks, there is a conversation worth having.

SALTZ AI can be reached at [bruno@saltz.ai](mailto:bruno@saltz.ai) and [emilio@saltz.ai](mailto:emilio@saltz.ai).

---

## Sources

1. Why AI Projects Fail, RAND Corporation, 2025. Meta-analysis of 65 enterprise AI initiatives (2022-2025). [rand.org](https://www.rand.org)
2. State of AI in the Enterprise, 2026, Deloitte, 2026. [deloitte.com](https://www.deloitte.com)
3. llm-wiki.md, Andrej Karpathy, April 5, 2026. Public GitHub gist. [gist.github.com](https://gist.github.com)
4. GBrain / GStack, Garry Tan, April 5, 2026. Open-source release under MIT. [x.com](https://x.com)
5. Team OS, Hannah Stulberg, DoorDash, April 2026. Internal architecture write-up, referenced in Lockey synthesis essay (see 7).
6. Ramp / Glass. Adoption figures, skill-library data, and Goddijn attribution per public Ramp communications and Lockey synthesis essay (see 7).

7. "Company Brain (YC RFS 2026): Four Builders, One Architecture", Alex Lockey, April 30, 2026. Synthesis essay. [alexlockey.com](https://alexlockey.com)
8. Requests for Startups, Summer 2026, Y Combinator, May 2026. Includes "Company Brain" (Tom Blomfield) and "AI Operating System for Companies" (Diana Hu). [ycombinator.com/rfs](https://ycombinator.com/rfs)
9. Imprese e ICT, Anno 2025, Istat, December 2025. Italian AI adoption: 5.0% (2023), 8.2% (2024), 16.4% (2025). [istat.it](https://istat.it)
10. SME AI Adoption Survey, digitalswitzerland, January 2026. 67% of Swiss SMEs plan to integrate AI by end of 2026; 18% report a structured plan. [digitalswitzerland.com](https://digitalswitzerland.com)
11. Italian government AI funding programmes (2025-2026): Voucher AI Ready, €40,000 per qualifying enterprise (Invitalia); Regione Lombardia "Transizione digitale delle imprese lombarde", €34.4 million (PR FESR 2021-2027, [regione.lombardia.it](https://regione.lombardia.it)); Regione Veneto Fondo Competitività, Sezione Transizione, €70 million (PR FESR 2021-2027, [venetoinnovazione.it](https://venetoinnovazione.it)); Iperammortamento, 180% in the 2026 Budget Law.
12. Italian Law 132/2025. First national EU AI Act implementation, in force October 10, 2025. Integrated GDPR and AI Act compliance regime. Gazzetta Ufficiale, Italian Republic.