

WHAT NOBODY WROTE DOWN

Why AI runs on context, and how to capture it and keep it.

A GITLAI whitepaper · October 2026

Jason Gilbreath · GITLAI, LLC · getintheloop.ai · jason@getintheloop.ai

For operators of PE-backed and owner-led companies, and the firms that back them.

01 · THE GAP IS CONTEXT, NOT THE MODEL

The models are ready. Companies are using them. In KPMG's latest quarterly survey of large U.S. companies, 62% say they are building, deploying, or developing AI agents, up from 53% the quarter before. A quarter of them are already building systems where several agents work together.¹

The results haven't kept up. In PwC's August survey of CEOs, 39% said AI had raised their revenue, cut their costs, or both. That's up from about a third at the end of last year. It still leaves most CEOs without either.²

So where does it go wrong? When it goes wrong, look at what the AI was given to work with.

In a June survey of 101 enterprises, VentureBeat found that 57% had traced a wrong answer from an AI agent to missing or inconsistent business context within the past six months.³ Deloitte's August survey of 501 U.S. leaders piloting agentic AI found that 72% say their organization lacks unified, accessible data. Only 21% say their business processes are ready.⁴

Read those findings together.

The AI isn't the hard part anymore. The hard part is what it knows about your business.

62%	Large U.S. companies building, deploying, or developing AI agents ¹
57%	Enterprises that traced a wrong AI answer to missing or inconsistent business context ³
21%	Leaders piloting agentic AI who say their business processes are ready ⁴

02 · WHAT CONTEXT IS

Context is the knowledge that lives only in your company, written down, curated, and kept current.

Start with the difference between data and context.

Your systems hold data. They record that a price was \$41.50, that an order shipped on Tuesday, that a customer is on net-60 terms.

Your people hold the context. They know why the price was held at \$41.50 when the sheet said \$44. They know which customer gets a call before the invoice goes out. They know which exception is fine and which one needs a manager. They know who to call when the system is wrong.

Data tells AI what happened. Context tells it why, and what to do next.

Most of that context has never been written down. It lives in a few people's heads, in inboxes, in side spreadsheets, in the code someone wrote five years ago and never documented. It shows up in the decision, not in the record.

Knowledge is the raw material. It becomes context when it's written down, curated, and owned. Until then, AI can't use it, and neither can the next person you hire.

03 · WHAT HAPPENS WITHOUT IT

We tested this directly. In July 2026, GITLAI built an AI assistant for a fictional company: its own policies, pay rules, systems, and five fictional employees in their first 90 days. We gave the assistant curated company knowledge and a written job description. Then we built its twin: the same assistant, the same AI model, with no company knowledge at all.⁵

We asked both the same questions.

The assistant with context answered 20 of 20 correctly on the mid-size and large models: dates, dollar amounts, policy rules. On the small model it answered 19 of 20. When it didn't know something, it said so and named the right person to ask.

The twin gave fluent, generic advice. It didn't produce one fact specific to the company, because it had none.

Same AI model. Same questions. The only difference was context.

That's the cheap version of the failure. The twin sounded helpful and was useless. The expensive version is an AI that sounds just as sure of itself and gets your pricing, your terms, or your exceptions wrong. That's what the 57% in the VentureBeat survey describes.³

There's a second cost, and it shows up whether you use AI or not. Mid-market companies carry their operations in a few heads. When one of those people leaves, the context leaves with them. Writing it down is key-person risk reduction you can point to.

04 · HOW TO CAPTURE IT

Capture starts with people, not systems. Your systems will tell you what happened. Only your people can tell you why.

Watch the work as it happens. Interviews get you the version people remember. Recorded work sessions get you what they actually do. In a GitLAI engagement, the people who run a desk narrate their work as they do it: what they're looking at, what they're deciding, and why. Recording is stated up front, in writing, in the engagement agreement and the kickoff materials.

Review it every night, and send it back. Each night, AI reviews the day's recordings and drafts what it saw. That draft goes back to the person who did the work with one question: is this right? Their corrections become the record. Nothing becomes context until the person who knows it has checked it.

Turn decisions into a log, and the log into rules. When the same kind of decision shows up often enough, it becomes a rule. At one company, people in key roles narrated their decisions as they made them. Those recordings became a decision log, and the log became a working rulebook. It works today. It will always be a working draft, because the work keeps changing.

Gather what the systems hold. Map where the data lives and who owns it. Review every code base, so the logic buried in old code becomes visible. Document every workflow. Interview the leaders and the people who own each function.

Package it so someone else can pick it up. The output of capture is a context pack: organized, sourced, and readable by a person or an AI. The test is a handoff. Give the pack to someone who wasn't there, or to a second AI model, and see whether they can pick up the work.

Capture is diagnosis. We use AI to find the problem before deciding whether AI is the solution. Sometimes what capture finds is a rule, a cleaner spreadsheet, or a person who needs backup. That's a complete answer too.

05 · WHY IT GOES STALE, AND HOW TO KEEP IT

Context starts going stale the day after you capture it.

A pricing manager leaves. A supplier changes terms. A system gets patched and a workaround stops working. A new customer arrives with exceptions nobody has seen before. None of that updates itself.

The AI doesn't know any of this changed. It keeps answering from what it was given. That's how a system that was right in March is confidently wrong in June.

So context needs an owner and a cadence, like inventory or the books.

Keeping context current means four things, every month:

1. **Re-capture what changed.** Short recorded sessions with the people whose work moved.
2. **Read the logs.** Every AI action leaves a record. The logs show where AI hesitated, guessed, or got corrected. Each of those is a gap in the context.
3. **Retire what's stale.** A rule that no longer matches the work gets updated or removed, with the date and the reason.
4. **Report what changed.** Leaders see what was added, what was retired, and what the AI is now doing differently.

Every part of the context has a named owner inside your company. When the AI doesn't know something, it says so and names that person.

Captured context is a snapshot. Kept context is an asset.

06 · THEN THE JOB DESCRIPTION

Once the context exists, you can decide what AI should do with it.

That's the job description: what AI does, who supervises it, how it's scored, and where it stops. We wrote a full paper on it, *The Job Description Comes First*. In practice, the job description is step two. You can't write a good one until you know how the work is really done.

Software is installed. AI is employed. Employing it well starts with what it will know.

The order matters:

1. **Context first.** Capture how the work is done, and keep it current.
 2. **The job second.** Write down what AI does with that context, and who owns it.
 3. **Proof before production.** Test it against the job description, in your environment, before anyone relies on it.
-

07 · IF YOU HOLD THE PORTFOLIO

Private equity sees this problem at portfolio scale. The barriers KPMG's private equity specialists name are data fragmentation, a shortage of people who can do the work, governance risk, and resistance to change.⁶

Context work fits portfolio work for three reasons.

It's repeatable. The capture method starts with the work, not the software. It runs the same way at a company on Salesforce and a company on spreadsheets. Each company's capture makes the next one faster.

It's measurable. A context pack is a thing you can inspect. Decisions logged, rules written, owners named, gaps closed each month. That turns "AI initiative" into something an investment committee can review.

It reduces the risk you already worry about. Getting operating knowledge out of a few heads and into owned, kept context is key-person risk reduction you can point to in diligence.

For a portfolio, the goal is one playbook instead of ten experiments.

08 · WHAT IT LOOKED LIKE AT A PE-BACKED OPERATING COMPANY

- A working draft of every Phase One deliverable was in the leadership team's hands on the day of kickoff.
 - An eight-person leadership team surfaced eleven roadmap candidates in one facilitated hour.
 - People in key roles narrated their work as they performed it and their decisions as they made them. Those recordings became a decision log and, over time, a working rulebook.
 - Every code base the company runs on was reviewed, and every workflow was documented.
 - The captured context was complete enough for a second AI model to pick up the engagement.
 - The context turned hours of work across a team of three into seconds every week, giving them back the time they needed to increase production.
-

09 · CAPTURE AND KEEP

Capture starts with Phase One. It takes about 30 days at a fixed fee. In that time we:

- Align with the owners and leaders on what matters.
- Capture the work with the people who run it.
- Map your data and your key-person risk.
- Build one working proof in an environment you own.
- Hand over a sequenced roadmap your team can execute, with or without us.

Keep is monthly. We keep your context current, read the logs, tune what's running, and report what changed.

Build comes when the roadmap calls for it. We scope it with you, and bring in a build partner when the work needs one.

One test tells you whether an AI partner built you an asset or a dependency. Imagine the AI platform you use disappears tomorrow. What survives? If the answer is your context, your rulebooks, your job descriptions, and your trained people, you own your capability. The platform is a choice you keep making.

And if the proof does not earn the next phase, you keep everything you paid for.

Start with Phase One. Fixed fee. Weeks, not quarters. Working proof in your environment. You keep everything either way.

Jason Gilbreath · GITLAI, LLC · getintheloop.ai · jason@getintheloop.ai

GITLAI, because your team belongs in the loop.

SOURCES

1. KPMG, AI Quarterly Pulse Survey, Q3 2026, September 24, 2026 (314 U.S. leaders at companies with \$1B+ revenue).
2. PwC, CEO Survey Snapshot, August 4, 2026 (351 CEOs, 59 countries).
3. VentureBeat, VB Pulse survey, June 2026, published July 10, 2026 (101 enterprises with 100+ employees).
4. Deloitte, "AI Agents are Only the Beginning: Deloitte Survey Examines the AI Readiness Gap," August 12, 2026 (501 U.S. senior leaders at organizations piloting agentic AI).
5. GITLAI platform evaluation, July 2026. Platform and figures are documented in GITLAI's evaluation records. All test content was fictional; no client data was used.
6. KPMG specialists, via Private Equity International, "Making the AI roll-out count," September 2026.

© 2026 GITLAI, LLC. All rights reserved.