Skan AI Raises $63M to Map How Work Actually Happens—Because Enterprise Agents Keep Automating Processes That Don’t Exist

$681.5M went into AI agent startups in the first 15 days of August 2026. The largest cheques didn’t go to agents. They went to the layer that tells agents what the process actually is.

The news

On August 12, 2026, Skan AI announced a $63 million Series C, co-led by Cathay Innovation and Dell Technologies Capital. Citi Ventures, Bloomberg Beta, State Farm Ventures and Wipro Ventures joined. Total funding now sits at roughly $120 million; no valuation was disclosed, per coverage of the round.

The pitch is narrower and stranger than most agent startups: Skan uses computer

vision and activity capture to observe how employees actually move work through enterprise applications — including legacy and regulated systems that don’t expose clean APIs — and builds what it calls a **Context Graph of Work**.

Not a process diagram from a Confluence page. A record of what people actually clicked, in what order, with which exceptions.

The product stack, per Skan’s own description, has three layers:

- **Blueprint** — discovery. Maps which workflows are actually followed versus only documented.
- **Intelligence** — analysis of decisions and exception handling.
- **Agents** — execution, tested against real observed cases rather than idealised specs.

Coverage from devcuration describes “Agent Operating Procedures” (AOPs): living process models encoding rules, exceptions and escalation paths for agentic governance. Skan also states that the resulting Context Graph belongs entirely to the customer enterprise, not to Skan — a claim I’ll return to, because it’s doing more work than it appears.

Why the timing matters more than the size

$63M is not a headline number in 2026. The interesting part is where it sits in the month’s pattern.

Per the Gravity.Fast AI Agent Funding Tracker, AI agent startups took roughly $1.32 billion across 20 disclosed rounds from July through mid-August 2026: ~$643M across 13 rounds in July, then $681.5M across just 7 rounds in the first 15 days of August. Round sizes went up sharply as round count went down. Capital is concentrating.

And look at what the concentration bought:

- **HappyRobot** — $150M Series C at a $1.2B post-money valuation, August 4, led by Prysm Capital, co-led by Eurazeo, with a16z, Base10 and Y Combinator returning. Agents that run phone calls, emails, document workflows and scheduling end-to-end in freight and logistics without human handoff. Total funding ~$200M.

- **CodeRabbit** — $143M Series C, August 12, led by Atomico and Smash Capital.
- **Zenity** — $125M Series C, August 3, led by Norwest, in AI agent security and governance.
- **Skan AI** — $63M Series C, August 12, in process discovery and context.

Three of those four are not agents. Zenity guards agents. Skan tells agents what the process is. CodeRabbit reviews what agents (and humans) write. Only HappyRobot sells the agent itself — and notably, it sells into one vertical with one very concrete job.

> The 2026 agent market is quietly splitting into two businesses: doing the work, and knowing what the work is. The second one is turning out to be harder.

The technical problem Skan is attacking

Enterprise agent deployments fail on context, not capability. That’s the bet.

Here’s the failure mode in concrete terms. A team documents an invoice exception process: four steps, two approvals, one system. An agent is built to that spec. In production, the real process has eleven steps, involves a spreadsheet nobody mentioned, routes around a broken integration via manual re-entry, and has three undocumented escalation paths that senior staff use for edge cases. The agent handles the happy path — which might be 60% of volume — and silently mangles the rest.

Process mining tools have attacked this for a decade by reading event logs from source systems. That works when systems emit good logs. It fails precisely where enterprise work is messiest: desktop applications, terminal emulators, regulated systems that can’t be instrumented, and the human glue between them.

Skan’s approach — computer vision plus activity capture at the workstation layer — is a deliberate answer to that gap. Observe the screen and the interaction, not the database. That gets you the spreadsheet. It gets you the copy-paste between two systems that no integration covers. It gets you the pause where an experienced employee reads something and makes a judgement call.

What “Agent Operating Procedures” actually implies architecturally

The AOP framing is the part senior engineers should pay attention to, and coverage has treated it as marketing.

My take: an AOP is essentially a versioned, machine-readable specification of a process derived from observation rather than from interview. If it works as described — encoding rules, exceptions and escalation paths — it’s the artefact that sits between a process and an agent, and it’s the artefact you can diff, test and audit.

That matters for three reasons:

- **Regression testing.** Skan says agents are “tested against real observed cases.” That’s the only sane way to validate an enterprise agent. You need a corpus of actual historical cases, including ugly ones, not synthetic prompts.
- **Drift detection.** Processes change. If your context model is a living observation feed rather than a one-off consulting deliverable, you can detect when reality diverges from what your agent was built for. Static process documentation cannot do this and never could.
- **Auditability.** When an agent makes a decision a regulator questions, “the LLM decided” is not an answer. “The agent followed AOP v4.2, escalation path 3, which was derived from 1,400 observed cases” is closer to one.

> If you can’t diff your process model, you can’t govern your agents. You can only hope.

The ownership claim

Skan’s statement that the Context Graph belongs entirely to the customer, not to Skan, reads like a legal footnote. It isn’t.

Workstation-level observation of employee activity is one of the most sensitive datasets an enterprise can generate. It touches employee monitoring law, works council agreements in the EU, and — for the regulated systems Skan specifically targets — data residency and privilege questions. State Farm Ventures and Citi

Ventures on the cap table suggest the insurance and banking buyers who care most about this are already in the conversation.

Unconfirmed: the research data doesn’t specify how observation is scoped, redacted, or consented to. Anyone evaluating this needs to ask, in detail, before the pilot.

The contrarian take

My take #1: “the context layer is the real moat” is half right, and the half that’s wrong is expensive.

The bull case writes itself: models commoditise, agents commoditise, but a proprietary map of how your specific enterprise actually works is durable and non-transferable. Every agent vendor becomes a customer of the context layer.

The problem: Skan explicitly does not want to own the graph. It says the graph belongs to the customer. That’s the right ethical and commercial answer for enterprise sales, and it is also a direct argument against the moat thesis. If the artefact is the customer’s, the vendor’s defensibility rests on the capture pipeline and the tooling — which is real engineering, but it’s engineering that Microsoft, ServiceNow and the incumbent process-mining vendors can also fund.

My take #2: the third layer is the risk.

Blueprint and Intelligence are observation and analysis. Layer three is Agents — Skan executing work itself. That puts it in competition with the agent vendors that would otherwise be its natural distribution channel. Discovery tools that also sell automation have a long history of getting stuck: too automation-focused to be a neutral system of record, too discovery-focused to beat specialists on execution.

My take #3: HappyRobot’s valuation is the more honest signal in this dataset.

$1.2B post-money for agents doing phone calls, emails, documents and scheduling in freight, end-to-end, no human handoff. Narrow domain, measurable output, an industry where the incumbent workflow is literally people on the phone. That’s a $150M round on demonstrable substitution, not on a platform narrative.

Compare: Skan’s round is $63M on the thesis that everyone else’s agents need better ground truth. Both bets can pay. But one has a clearer path from product to invoice.

My take #4: what nobody in the coverage flagged.

Neither Skan AI nor HappyRobot disclosed ARR, customer counts, or call-volume metrics in their funding materials. A $1.2B valuation and a “Context Graph of Work” were both announced with no operating numbers attached.

That is normal for Series C in this market, and it should still change how you read the round. What you’re seeing priced is investor conviction about a category, not disclosed traction. Treat these announcements as a map of where smart capital thinks the bottleneck is — which is genuinely useful — and not as evidence that either product works at scale.

> Seven rounds, $681.5M, fifteen days, and not one disclosed revenue figure. Funding size is a signal about the category, not about the company.

What to actually do about this

If you’re a CTO with an agent programme that’s underperforming, the useful takeaway isn’t “buy Skan.” It’s that the diagnostic they’re selling is one you can partly run yourself.

Before your next agent build:

- **Assume your process documentation is wrong.** Not incomplete — wrong. Pick your highest-volume workflow and reconstruct it from logs, tickets and screen recordings of three experienced operators. Compare to the documented version. The gap is your agent’s failure rate ceiling.
- **Build the exception corpus first.** Collect 100+ real historical cases, weighted toward the ugly ones. That’s your test set. If you can’t assemble it, you’re not ready to deploy an agent, regardless of vendor.
- **Instrument for drift.** Whatever your process model is, it needs a version number and a mechanism that fires when production behaviour diverges. Most

agent deployments have no such mechanism, which is why they degrade silently.
- **Separate the system of record from the executor.** Whether you buy Skan or build internally, keep the process model as an artefact you own, in a format you can export. Do not let it live only inside a vendor’s runtime.

Vendors worth tracking in this layer: Skan (discovery/context), Zenity (agent security and governance — $125M led by Norwest is a serious vote for that problem being real), CodeRabbit (review of machine-generated output). The pattern across all three: the money is going into control planes.

Diligence questions for any context-layer vendor: How is workstation observation scoped and redacted? What’s the export format for the process model? What happens to the graph if the contract ends? How does drift detection work, mechanically? Who signs off on employee monitoring in your jurisdiction? None of these are answered by the research available on Skan, and all of them will decide whether a pilot survives legal review.

Where this goes in 6–12 months

I expect the “process model as a first-class artefact” pattern to get a name and a format fight. Right now Skan calls it a Context Graph and AOPs; others will call it something else. Whoever gets to a portable, inspectable specification that agent frameworks can consume directly wins more than whoever has the best capture technology.

I expect at least one major incumbent — process mining, ITSM, or hyperscaler — to acquire or clone in this space within twelve months. The strategic logic is too clean: whoever owns the process model owns the sequencing of every agent deployment behind it. Cathay Innovation and Dell Technologies Capital co-leading a $63M round in a company with ~$120M total raised is a positioning move as much as a growth investment.

I expect the July-to-mid-August pattern to hold: fewer rounds, larger cheques, concentrated in control and governance rather than agent capability. $1.32B across 20 rounds in six weeks, with round sizes rising as count falls, is the shape of a market where investors have decided which layer is scarce.

I expect the first serious enterprise agent failure post-mortem to blame context,

not the model — and for that post-mortem to be the best marketing this category ever gets.

What I don’t expect, and would flag as the genuine open question: whether observation-derived process models stay accurate long enough to be worth the capture cost. Enterprise processes change constantly. A context graph that needs continuous re-observation is a subscription business; one that decays faster than it can be refreshed is a consulting engagement wearing SaaS clothing. The research data doesn’t tell us which Skan has built, and no disclosed retention or ARR figures exist to infer from.

The agent era’s real bottleneck isn’t reasoning capability — it’s that almost no enterprise can produce a truthful, testable, machine-readable description of its own work, and $681.5M in fifteen days says investors have finally noticed.