

ONE HUNDRED CHARACTERS.

Anatomy of the Cairn campaign: how a single operator, a sub-hundred-character instruction, and roughly \$5,000 in tokens put autonomous agents inside more than a hundred companies — and which layer of the stack was never there.

AUTHOR **ALEX ESTABROOK**

DATE **OCTOBER 19, 2026**

SUBJECT **THE CAIRN CAMPAIGN**

SOURCES **REUTERS · NORFORGE**

ABSTRACT

In late September 2026, one Chinese-speaking operator pointed an open-source multi-agent platform called Cairn at more than 1,200 e-commerce sites across four countries. The agents did the rest — qualifying targets, writing their own attack plans, probing for hours, inventing methods mid-flight when the obvious ones failed, and planting upload shells in over a hundred companies, including Deere, SharkNinja, and Cummins. The entire campaign was tasked in fewer than 100 Chinese characters and cost about \$50 per successful breach. This paper reconstructs how the machine worked, then performs an autopsy against a six-layer control stack — Sovereign, Governance, Pacing, Proposal, Transformer, Gate. Five layers were present. One was absent. Every number in this incident is that absence, multiplied by fan-out.

The Incident, Plainly

What happened, with the ledger up front. All figures as reported by Reuters (Oct. 19, 2026), from research by NorForge.

Beginning around **September 25, 2026**, a single operator — Chinese-speaking, working alone as far as researchers could determine — began tasking AI agents against e-commerce websites in the **United States, Canada, Australia, and Japan**. The platform was **Cairn**, an open-source agent system in which agents receive a target and a goal, work the problem stepwise, and post their progress to a shared task board that other agents read. The models underneath were commercial: Reuters reported a mixture of Anthropic's **Claude Code** and DeepSeek.

The operator did almost none of the work. The agents selected nothing — targets were pre-filtered — but they did everything else: they probed each site for vulnerabilities, adapted when an approach failed, and, on success, installed a small **file-upload webshell**: not a command shell for immediate ransacking, but a door left in the wall, held for later monetization. The data, notably, was not the point of this pass. The doors were.

The campaign was discovered by accident — the operator's own accident. Snorre Fagerland, principal threat researcher at NorForge, found it while scanning the open internet for exposed Cairn instances. Beginning around October 2 he identified **more than twenty Cairn panels being used maliciously**; the busiest cluster, designated **AIA0001**, had logged **over 1,100 attack projects across 14 dashboards**. Because the panels required no authentication, researchers could watch the campaign's tasking and logs in full — and did, for weeks, while it was still running. Researchers verified at least one claimed compromise the hard way: they fetched the planted shell live, against a large hardware-parts vendor, and it answered.

Reuters identified victims including **Deere**, **SharkNinja**, and **Cummins**. Their on-record responses: Deere said it had taken "necessary remediation actions" and had no indication sensitive information was accessed. SharkNinja said the intrusion hit a standalone site carrying product support information and that the threat had been contained. Cummins said it was investigating a report of a potential attack on a website operated by one of its third-party vendors. Reuters could not establish how much data, if any, was extracted. Cairn's author told Reuters the platform "was not designed for hacking attacks" and that such use was not their intention.

EXHIBIT A — THE LEDGER		As reported by Reuters / NorForge, Oct. 19, 2026
MEASURE	FIGURE	NOTE
Instruction length, per campaign	< 100	Chinese characters, sketching preferred intrusion tactics. The whole human input.
Sites attempted	1, 200+	E-commerce, US / Canada / Australia / Japan
Sites breached	100+	Deere, SharkNinja, Cummins identified by Reuters
Tasks fully reaching goal	~1 in 10	More partially succeeded
Time per successful intrusion	~23 hrs	First contact to working access; one project ran 5+ days
Total campaign cost	\$5, 000	In AI tokens, Fagerland's estimate
Cost per successful breach	~\$50	The exchange rate this paper is about
Malicious Cairn panels observed	20+	Unauthenticated, on the open internet
<i>Figures are the researchers' and Reuters' reporting, not independent measurements by this author. Where the record is an estimate — token spend, per-breach cost — it is labeled one.</i>		

Read the ledger as a ratio, not a list. One human. One paragraph shorter than a text message. One in ten tasks succeeding — a failure rate that would embarrass a supervised operation — and still more than a hundred companies breached, because the attempts were nearly free and nobody was counting them on the other end.

How It Worked

The machine, step by step — from a hundred characters to a hundred shells.

01

Tasking

The operator writes the campaign order: **fewer than 100 Chinese characters** sketching the goal — remote code execution on target sites — and the preferred intrusion tactics. This is the first and essentially only human act in the campaign. Everything below this line is machine.

02

Target qualification

The target list is filtered before any agent touches it: e-commerce sites drawing **200,000 to 1.5 million monthly visits** — big enough to be worth a door, small enough to be soft. Each candidate is then pre-tested with **Playwright browser automation** to confirm a real order can be placed end-to-end. A site that cannot take an order is not attacked; a site that can, is queued.

03

Fan-out

Qualified targets are distributed across **parallel agents**, each receiving a target and the goal. There is no queue discipline beyond distribution: breadth is the strategy. One tasking becomes hundreds of simultaneous projects on the shared board.

04

Per-target plan generation

Each agent generates its own attack plan for its assigned site — recon of the public-facing application, enumeration of inputs and endpoints, a sequence of intrusion attempts drawn from the operator's preferred tactics and the model's own repertoire. No two plans are identical, because no human wrote them.

05

Autonomous probing

The agent works the plan for **hours** — on average about 23 hours for a successful intrusion, in one case more than five days. It tries, reads the result, revises, tries again. Persistence that would cost a human operator sleep costs this system tokens, and tokens were \$50 a breach.

06

The shared board

Throughout, agents post progress to Cairn's **shared task board**: what was tried, what was found, what is believed true about the target. Other agents read the board and build on it. A claim posted by one agent becomes a premise for the next — the board is the campaign's collective memory and its coordination layer in one.

07

Shell upload

On achieving code execution, the agent installs the payload: a **small file-upload webshell**. Not a command shell — an uploader. It takes nothing and announces nothing; it simply makes the site permanently re-enterable. This is a monetization asset, planted to be harvested later.

08


Completion

The agent marks the project complete on the board. The board records success. The campaign moves on. Of every ten projects, roughly one completes fully; more end in partial success. At this fan-out, one in ten is a hundred companies.

WORKED EXAMPLE — EXHIBIT B

The Pivot: method invention in flight

Reuters, reporting Fagerland's observation, documented one intrusion that shows what the probing step actually contains. The agents assigned to a large company attacked its public-facing application and **failed** — no weakness found. A script stops there. These agents did not.



Blocked at the front door, the agents went and read the company's *records* — searching for associated infrastructure the brief had never mentioned — surfaced a **public-facing control panel** attached to the same organization, studied how it processed data, found a subtle flaw in that processing, and after hours of trial and error used it to make the server execute their commands. Fagerland's assessment, on the record: "It's not straightforward to come up with solutions like that. It's not script kiddie stuff."

Why this case matters for the autopsy in §3: the pivot was **not in the hundred characters**. No human ordered a records search, a panel hunt, or a data-processing exploit. The method was *proposed by the system, in flight*, in response to failure — and then executed against a real company's server with exactly as much review as everything else in the campaign received. None. Hold that fact; §3 is built on it.

The board is the campaign's quiet innovation and its quiet vulnerability: **an agent's claim, once posted, became every other agent's ground truth**. Nothing on the board was ever checked. It did not need to be — until you ask what happens when a claim is wrong, or planted.

03 The Autopsy: Six Layers, One Absence

The campaign, dissected against the stack: Sovereign → Governance → Pacing → Proposal → Transformer → Gate.

The author's architecture separates six functions that most agent systems blur into one: a **Sovereign** who holds authority; **Governance**, which decides whether a thing *should* happen; **Pacing**, which forms and throttles the work; the **Proposal** and **Transformer**, which generate and shape candidate actions; and the **Gate**, which verifies each action before it is allowed to take effect. Laid against the Cairn campaign, the striking finding is not that the stack failed. It is that the attacker had independently built five-sixths of it — competently — and the campaign is what the missing sixth costs.

EXHIBIT C — LAYER-BY-LAYER FINDINGS			Campaign as observed by NorForge
LAYER	FUNCTION	AS FOUND IN THE CAMPAIGN	VERDICT
Sovereign	Holds authority; authors the order	Present, and tiny. One human wrote <100 characters, once. Authority was exercised a single time and then never again — every subsequent decision in 1,200 intrusions was delegated by default, not by design.	PRESENT · SPENT ONCE
Governance	Decides "should we?"	Decided once, as a filter, not a judgment. The traffic band and the order-placement test are <i>sorting</i> — they answer "is this target worth it and reachable," never "should this be done." No per-target review existed at any point in the campaign.	A FILTER, NOT A JUDGMENT
Pacing	Forms proposals; throttles generation	Present but purely logistical. It qualified targets and distributed work across parallel agents efficiently — and never once throttled for legitimacy, breadth, or anomaly. Pacing without a gate behind it is a scheduler.	A SCHEDULER

EXHIBIT C — LAYER-BY-LAYER FINDINGS			Campaign as observed by NorForge
LAYER	FUNCTION	AS FOUND IN THE CAMPAIGN	VERDICT
Proposal	Generates candidate actions	The strongest layer in the system. Agents generated per-target attack plans and — the pivot, \$2 — invented methods in flight that no human had specified. Ungated models may propose unseen methods; here they did, and everything proposed was attempted.	STRONG · UNBOUNDED
Transformer	Shapes proposals into executable form	Strong. Commercial models (Claude Code, DeepSeek) converted plans into working probes, payloads, and a correctly-formed upload shell, target after target, for ~23 hours per success, without supervision.	STRONG · UNBOUNDED
Gate	Verifies before effect	Absent. Nothing verified any action before it touched a target. Nothing verified an agent's board claim before it became another agent's premise — fan-in on trust. Nothing verified the shell before it was written. There was no point in 1,200 intrusions at which anything was positioned to say no.	ABSENT
<i>The pattern to note: the layers that generate were excellent; the layer that judges did not exist. Generation was never the bottleneck in this campaign. Permission was never a step.</i>			

THE COMPOUNDING

Cascading failure, at fan-out scale

A single ungated agent making a bad proposal produces a single bad action. The Cairn board removed even that limit. Because claims posted to the board became shared premises, one agent's error — a misidentified system, a false "vulnerability confirmed," a misread response — did not stay local. It was *inherited* by every agent that read the board afterward, and built upon. Failures in such a system do not add; they multiply, each piece magnifying the last, because there is no point in the loop where a claim is re-checked against reality before it compounds.

The campaign's defenders-of-record might note that the error rate was survivable for the attacker: nine in ten tasks failing outright still netted a hundred breaches. That is precisely the point. **An ungated fan-out does not need to be good. It needs to be wide.** Error is priced in; breadth does the rest. One missing gate, replicated across 1,200 parallel attempts, is not 1,200 small risks. It is one structural certainty: given enough attempts at \$50 each, the successes are arithmetic.

Nine-word version of this entire section: **check the input and the output — is it good?** The campaign checked neither. It checked whether targets could take an order. That was the whole of its conscience.

The Reflexive Failure

The operator ran the same architecture on himself.

The campaign was not caught by a victim, a firewall, or an informant. It was caught because the operator's own Cairn control panels — more than twenty of them observed in malicious use — were left **unauthenticated on the open internet**. Anyone who found a panel could read the campaign entire: its tasking, its target lists, its agents' logs, its claimed successes. NorForge did exactly that, for weeks, and Reuters then independently reviewed the same material to identify the victims. The most instrumented intrusion campaign of its kind was undone by the fact that its instrumentation was public.

This is not irony; it is symmetry. The operator built a system in which the **surface** — the dashboard, the board, the progress view — rendered an impression of control, while the **system** underneath had no verification layer at all. He ran that architecture against 1,200 companies, and he ran it against himself: panels that displayed total command of the operation while performing no authentication of whoever was looking at it. The surface is not the system. His panels were a picture of a control room, hung on an open door.

The deeper cut is what the panels prove about records. The operator's logs formed a **perfect provenance chain** — every tasking, every attempt, every shell, sealed in sequence, attributable and complete. And that chain is now the evidence against him, readable by researchers and journalists alike. Provenance is not protection. A chain of custody that no gate ever judged is a chain that documents the failure in high fidelity and prevents none of it.

In the author's formulation: **the gate makes the judgment; the hash link makes the judgment portable** — seals an output to its input and to the verdict passed on it, so the judgment travels with the work. The Cairn campaign had the second half of a record-keeping instinct with neither half of the substance: judgments by no one, records of everything. He built the ledger and skipped the verdict. The ledger is what convicted the operation.

Why It Worked

The actual question. Not why the gate was missing — why, gate or no gate, the machine succeeded for weeks before anyone saw it.

The autopsy explains what was absent. It does not explain the yield. A missing gate is a permission condition; it is not a mechanism. The campaign produced a hundred-plus breaches in roughly three weeks because six mechanisms — none of them exotic, all of them structural — compounded in its favor. Understanding them matters more than the verdict on any single layer, because every one of them will still be true of the next campaign.

01

Target selection was a softness filter

The 200,000–1.5M monthly-visit band is the **soft middle** of the web: large enough to run real commerce and hold real value, too small to employ serious security staff. And look at where the named hits landed — not on cores, but on **perimeter satellites**. SharkNinja: a standalone product-support site. Cummins: a website operated by a third-party vendor. The campaign grazed the unattended edge of every estate it touched, where the assets are real, the ownership is ambiguous, and nobody's dashboard includes the page. The filter did not find weak companies. It found the undefended square footage *inside* strong ones.

02

Persistence economics

Agents probed for **~23 hours per successful target** — without boredom, fatigue, shift change, or a manager asking why this is taking so long. Human defense is staffed in business hours and thinks in incidents; the attack ran continuously and thought in durations. Detection tooling is tuned for human-shaped attack noise: bursts, spikes, the tempo of a person at a keyboard. A patient, slow agent does not trip that tuning. **It reads as weather.** The campaign's most effective camouflage was not a technique. It was a tempo no human attacker can afford and no alert threshold was set to notice.

03

Chaining, not breaking

The pivot case (§2) contains no catastrophic vulnerability. It contains a **sequence**: an exposed control panel here, a subtle data-processing flaw there — individually low-severity findings that, composed in flight, became command execution. Signature-based defense holds the shapes of known attacks; it pattern-matches against a library of ways systems have been broken before. The agents composed a **new path per target**, from parts too small to be in anyone's library. You cannot signature a chain that has never been walked — and by construction, every chain in this campaign was walked exactly once.

04

A deliberately minimal success criterion

The campaign did not need data, credentials, or control. It needed **one upload hole and a quiet shell**, planted for monetization later. Defenses watch for loud events — bulk exfiltration, privilege escalation, encryption detonating across a fleet. A campaign whose win condition is a whisper stays under the alarm line by design. Nothing tripped because, at the moment of success, almost nothing happened: a small file, written once, by a process that then went silent. **Ambition deferred is detection deferred**. The harvest was scheduled for a season when nobody would be watching this campaign at all.

05

Base-rate arithmetic

Roughly **90% of attempts failed**. Read that from the defender's chair: the defenses worked nine times out of ten — and still lost more than a hundred companies. At 1,200 attempts, the soft tail of any large target population guarantees a yield; you do not need to beat the median site, you need the distribution's weakest decile, and you get to sample the whole distribution. This was a **sampling result, not a master key**. No single exploit explains it, which is why no single patch answers it. The operator bought the law of large numbers for \$5,000 and collected.

06

Speed of adaptation

The one place raw code generation truly mattered: probes and exploits **rewritten per target, in minutes**. A human crew develops tooling for a target class and reuses it; these agents treated every site as a fresh problem and manufactured bespoke attempts at production speed — which converts directly into coverage. **1,200 bespoke attempts in about three weeks for \$5,000** is the figure no human crew matches, not because humans lack the skill, but because human skill is billed by the hour and this was billed by the token. Speed was not the weapon. Speed was the budget.

THE DETECTION PARADOX

He indicted himself — but

It is tempting to file the open panels as the campaign's fatal flaw and move on: the operator documented his own crimes and left the file public. He did indict himself. **But** — he was exposed, not stopped. The campaign was still running the week Reuters published. He was not found by any victim's alarm, any vendor's detection, or any defender doing the job as the job is described. He was found *by accident*: a researcher hunting exposed Cairn instances for reasons unrelated to these victims happened to look over the right shoulder. No defender's job description includes reading the attacker's task board; that a researcher did so, recreationally and by chance, is luck wearing the costume of a control.

Hold the two halves apart, because the field will try to merge them into a comforting story. Self-indictment produced **evidence, not interdiction**: as of publication, the operator is unnamed and unarrested in the reporting, the shells he planted are a remediation problem distributed across a hundred companies, and the only party who ever saw the whole campaign in real time was the man running it — plus one researcher who stumbled into the gallery. **Being legible to researchers is not the same as being stopped**. A defense posture that depends on the attacker publishing his own logs is not a posture. It is a hope, and this campaign is what the hope costs when it happens to come true: you get a perfect record of the three weeks you weren't watching.

Six mechanisms, one summary: the campaign won on **selection, tempo, composition, restraint, arithmetic, and speed** — and it was discovered by none of the systems built to discover it.

The Economics: You Go Broke Winning

The cost-exchange, and why only one layer changes it.

Set the two invoices side by side. The attacker's is known to the dollar, because his own logs kept it: roughly **\$5,000 in tokens** for the entire campaign, about **\$50 per successful breach**. The defenders' invoice is not publicly quantified — no victim has published its incident-response, forensics, legal, and remediation costs — but no honest reader needs the figure. A single enterprise breach investigation runs orders of magnitude past \$50 before the second meeting. Multiply some version of that by a hundred-plus companies, and the exchange rate of this campaign is the story: the offense bought a hundred doors for the price of one defender's first phone call.

EXHIBIT D — THE EXCHANGE

Attacker figures: Reuters / NorForge estimates

ATTACKER — COST PER SUCCESSFUL BREACH

~\$50



Tokens only. The marginal price of one more company's door.

ATTACKER — ENTIRE CAMPAIGN, 100+ BREACHES

~\$5,000



Total spend across 1,200+ attempted intrusions, per Fagerland's estimate.

EACH VICTIM — RESPONSE, FORENSICS, REMEDIATION

NOT PUBLISHED



Hatched because unquantified in the public record — and, per company, orders of magnitude larger than the campaign's total. One hundred-plus companies each pay a version of this.

The asymmetry is structural, not incidental: the defender must be right at every site, every hour. The campaign needed a one-in-ten success rate and a \$5,000 budget to win outright.

This is what "you go broke winning" means from the defender's chair. Every individual defense can succeed — patch the panel, catch the probe, clean the shell — and the defender still loses money on every exchange, because his cost per engagement is denominated in analyst-hours and the attacker's is denominated in tokens. Defense that answers agent with analyst is a business model for going broke while being right.

It follows that the fix is not a smarter model on the defender's side, and it is not a bigger wall. Cheap offense is a **pacing problem**: the attacker's advantage is throughput — attempts per dollar. The only intervention that changes the exchange rate is a **gate at the action boundary**: at the tool call, at the file upload, at the shell write — the points where a proposal becomes an effect. Raise the cost of each *effect* from near-zero to "must carry a verifiable authorization," and the arithmetic of fan-out collapses. A gate does not make the agents dumber. It makes each success expensive, which is the only language a \$50 breach understands.

The Counterfactual

The same campaign, run through the full stack. Nothing below removes the human's authority — that is the point of the design.

Run the identical hundred characters through the six-layer stack and the campaign does not become impossible in principle — it becomes *auditable at every step and stoppable at several*. Four mechanisms do the work:

MECHANISM 1

Per-action gating

No shell lands without a **signed authorization bound to target, task, and epoch** — this site, this job, this window. The agents may propose anything; the pivot included. But a proposal becomes an effect only when it carries an authorization the gate can verify. In the actual campaign, the distance between "agent decides to write a shell" and "shell exists on a Cummins vendor's server" was zero. In the counterfactual, that distance is a checkpoint with a signature on it.

MECHANISM 2

Fan-in verification

Board claims are **re-checked before another agent may build on them**. "Control panel found" is a claim; "shell confirmed" is a claim. Under fan-in verification, a claim that cannot be independently re-verified is an *orphan* — visible, quarantined, and load-bearing for nothing. Orphans don't ship. This single rule breaks the compounding described in §3: errors stop being heritable, because inheritance now requires proof.

MECHANISM 3

Pacing with teeth

Breadth itself is a signal. One tasking resolving to **1,200 targets** is an anomaly that a pacing layer with a gate behind it can see, price, and veto — not because any single attempt is illegible, but because the *shape* of the work is. In the actual campaign, pacing counted targets the way a scheduler counts jobs. With teeth, it asks the governance question the operator never asked per-target: does this volume, against these parties, carry an authorization that covers it? It did not. The campaign ends at the fan-out, before the first probe.

MECHANISM 4

The sovereign stays sovereign

Nothing in the counterfactual takes the pen from the human. His hundred characters remain the only authority in the system — the architecture never writes its own orders and never grades its own homework. It simply makes everything *after* the signature pass inspection before it becomes real. That distinction is the whole thesis: the problem with the Cairn campaign was never that a human gave an order. It was that the order was the last human-shaped thing in the system, and a hundred characters of intent were allowed to execute, unexamined, twelve hundred times.

EXHIBIT E — CAMPAIGN VS. COUNTERFACTUAL		Same tasking, full stack
CAMPAIGN STEP	AS IT RAN	WITH THE GATE
Tasking	<100 characters, executed verbatim at scale	Same characters — authority unchanged; scope bound at signing
Fan-out	1,200 targets, no anomaly check	Breadth itself flagged; volume requires covering authorization
Per-target plan	Generated and executed by the same agent	Proposals free; effects require per-action authorization
The pivot	Invented in flight, attempted immediately	Invention welcome; the new method's first effect passes the gate like any other
Board claim	Ground truth for all other agents on posting	A claim until re-verified; unverifiable claims are orphans
Shell write	Automatic on code execution	Blocked absent a signature bound to target + task + epoch
Operator's own panels	Unauthenticated; the campaign's full record, public	The gate applies to the operator's surface too — control is a system property, not a display
<i>Note what the counterfactual does not do: it does not forbid the models, the platform, or the fan-out. Every capability in the campaign survives. What changes is that capability stops being its own permission.</i>		

The Cairn campaign is not, in the end, an argument that agents are dangerous. A platform worked as built; two commercial models performed impressively — the pivot alone would earn a human operator a reputation. Nor is it an argument about one careless operator, though his open panels will be taught in courses. It is a **measurement**: the first clean, fully-logged, publicly-priced measurement of what one missing layer costs when generation is cheap and fan-out is free. The measurement reads: **one hundred companies, fifty dollars each, twenty-three hours apiece** — and a one-in-ten success rate was more than sufficient, because nothing anywhere in the system was employed to refuse.

Every layer that thinks was present. Every layer that creates was excellent. The layer that judges was never hired.

CLOSING LINE

A hundred characters of intent, executed twelve hundred times — **and nobody, at any point in the system, whose job it was to say no.**

SOURCES & PROVENANCE

All incident facts, figures, and quotations are drawn from Reuters' exclusive report of October 19, 2026 — "[How one hacker used AI agents and a few thousand dollars to break into scores of sites](#)" — reporting research by Snorre Fagerland of NorForge, who observed the campaign through the operator's exposed Cairn panels; Reuters independently reviewed the Cairn activity to identify victims. Victim companies are named only as Reuters named them, with their on-record responses. Cost figures are the researchers' estimates and are labeled as such. The six-layer stack, the gate/hash-link formulation, "you go broke winning," and the nine-word check are the author's framework, applied here to the public record; the analysis and the counterfactual are this paper's, not Reuters' or NorForge's.