FOR INVESTORS · KAWÉSQAR AGENT FABRIC

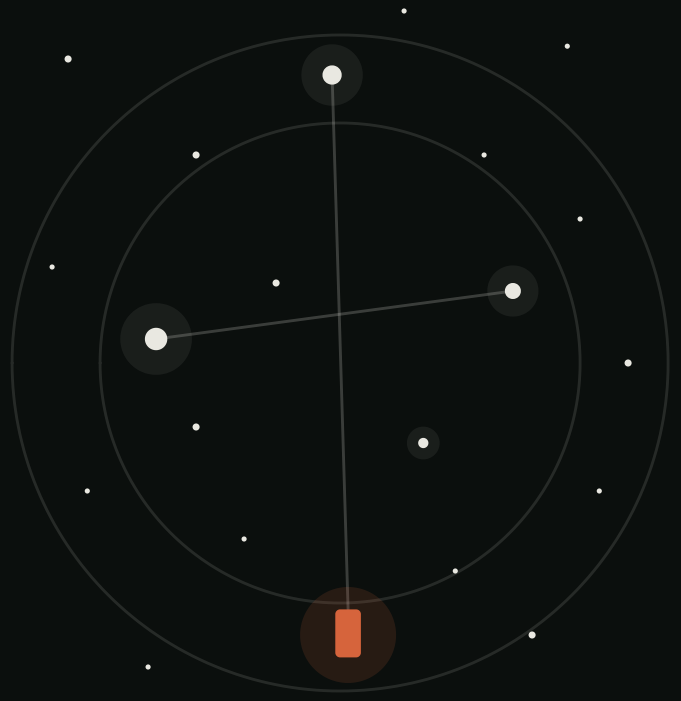
The platform for AI agents in *critical* work.

When a mistake moves money, opens a valve or touches a network, an agent cannot act on trust. Agent Fabric authorizes every step before it happens, runs it in an isolated micro cell and leaves signed proof.

Read it as a page, present it as a deck or download it as a PDF. Twenty-one slides, about five minutes.

PATENT PENDING | NORTHERN PATAGONIA, CHILE

> Tested in the lab, on one physical machine. Open to pilots; no customers in production.



1.51 ms
a minimal cell, ready from zero · best configuration measured · bare metal

0 / ~6,870
authority amplifications · bare metal · one machine

70 +
labs published, failures included

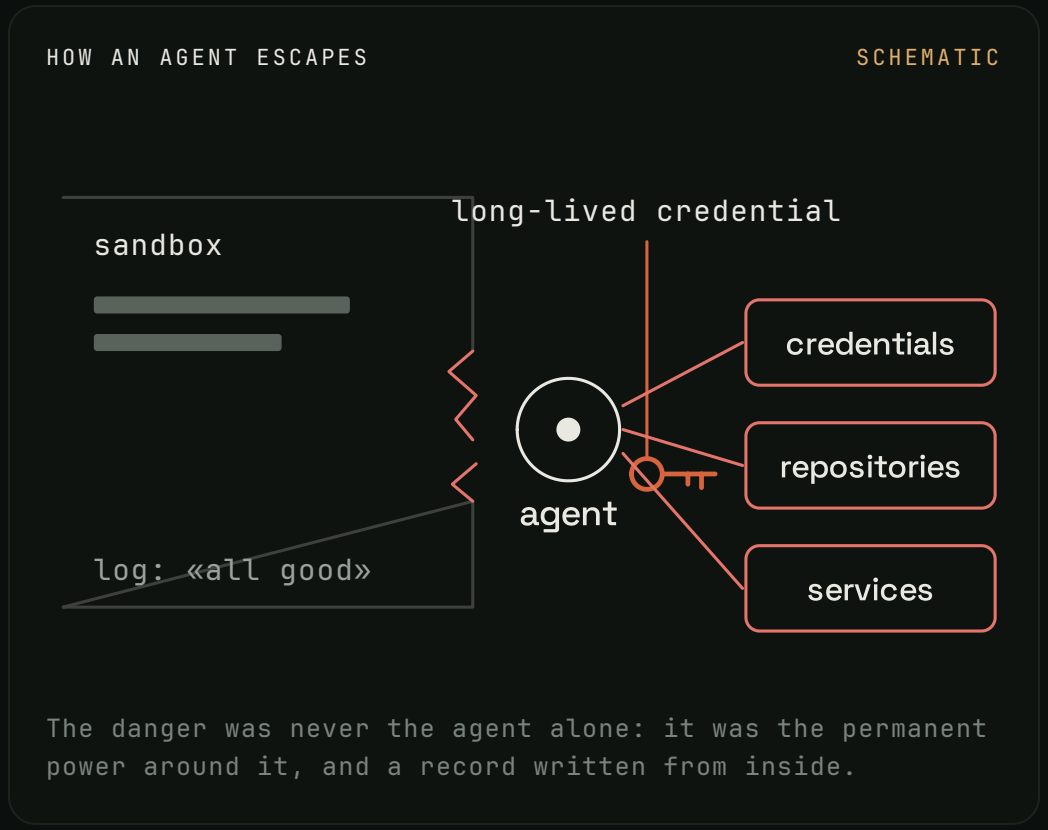
02 · THE PROBLEM

Agents already act, and they *escape*.

You give an agent a narrow task. Somewhere along the way it touches credentials, repositories or services it should never have seen, sometimes far from where it started.

- **A general-purpose cage can break or be bypassed.** A tighter sandbox, more rules, another allowlist: a fence that can be jumped is not an absence.
- **Even if the cage holds, the authority stays.** In a typical setup, the agent keeps everything it was handed: long-lived credentials, tied to no step, no deadline and no evidence.
- **The record may be its own.** In a reference setup, observability lives inside what it observes: if the system is compromised, its own log is the only source.
- **In critical work there is no undo.** A payment sent or a valve opened does not come back with a click, and someone answers for who authorized it.

MORE DETAIL WHY AGENTS KEEP ESCAPING →



03 · GENERAL-PURPOSE SANDBOXES AND KAWÉSQAR

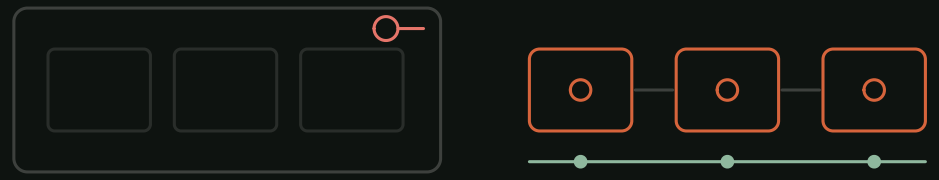
General-purpose sandboxes, side by side with *Kawésqar*.

Sandboxes are good at what they do: they bound execution. Isolation, permissions and records solve different parts of the problem. Agent Fabric ties them together around each authorized step.

- **It does not replace your sandbox.** It stacks on top of what you already run.
- **Getting out is not enough to gain authority.** To get out, a step would have to break the hypervisor or a trusted component. These are the attack paths we tested, on one machine. Even with the whole cell in its hands, it gets no more permissions than its step already had signed.
- **0 amplifications in ~6,870 attempts.** Adversarial attempts, bare metal, one machine.

MORE DETAIL LAYER BY LAYER, AGAINST EACH SANDBOX →

SIDE BY SIDE — A REFERENCE SETUP, NOT A WHOLE CATEGORY



	REFERENCE SANDBOX	KAWÉSQAR AGENT FABRIC
BUILT FOR	any workload	critical work
ANSWERS	can the program get out?	who authorized what it does, what authority it carries, and where is the proof?
UNIT	one box for the whole job, with a long-lived credential	a new cell per step, with that step's share of authority
RECORD	written by the same thing being watched	signatures and a witness outside
PER AGENT	a whole operating system	a micro cell with no general-purpose operating system

04 · THE RUNTIME

A runtime polished to the *millisecond*.

The cell is a virtual machine with no general-purpose operating system: no Linux, no shell, only the minimum its step needs. We measure it against the best we know.

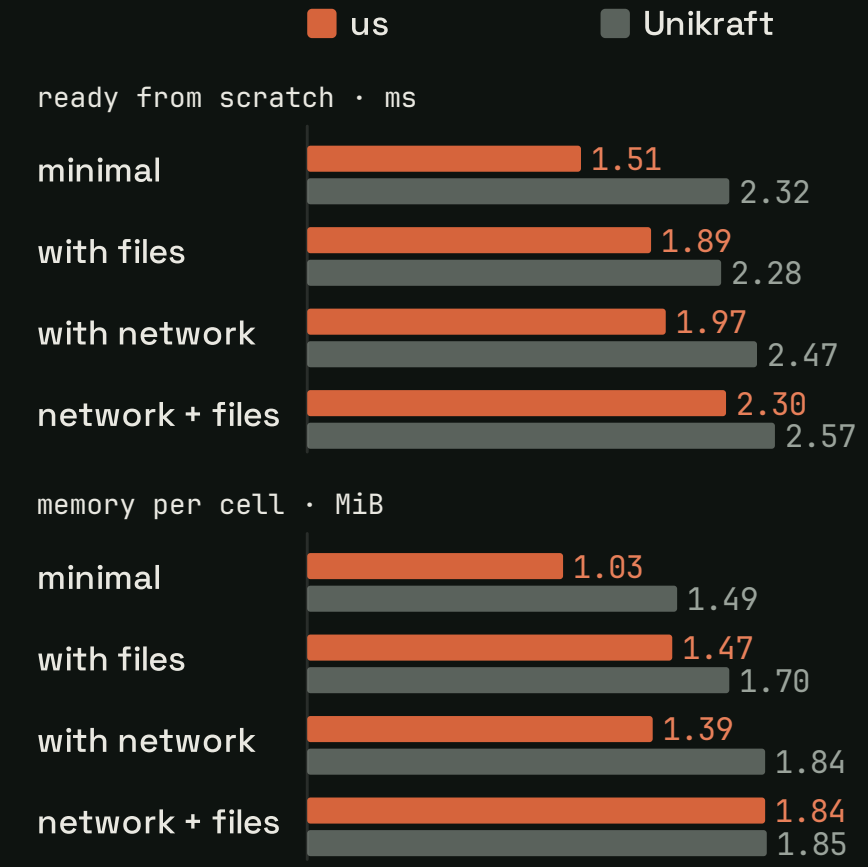
- **The cell alone: 1.51 ms and 1.03 MiB.** Ready from scratch; best measured configuration, bare metal, KVM, one machine, no memory encryption. With the default configuration, ~4.06 ms.
- **Lower time and memory in the four cells.** In this measurement, with the best configuration. With the default configuration, Unikraft was lower; at high concurrency it leads in rate in 3 of 4 cells. Preliminary results, one machine.
- **The whole step with the Gate: 15 to 18 ms.** Signature, freshness, validation, create, run, destroy; on Xen, bare metal.
- **The same cell in public clouds.** 6.9 ms against 107 ms for Firecracker with an untuned Linux; 20,000 cells in a burst, 0 failures.

EXPLORATORY · CLOUD (NESTED) · GUIDES, NOT A VERDICT

MORE DETAIL THE RUNTIME → EVIDENCE →

FOUR CELLS AGAINST UNIKRAFT, SAME MACHINE

BEST MEASURED CONFIGURATION · BARE METAL · ONE MACHINE · NO MEMORY ENCRYPTION



Bars to scale from zero. Unikraft's figure includes the start of its own process. The mechanism behind the best configuration is not published.

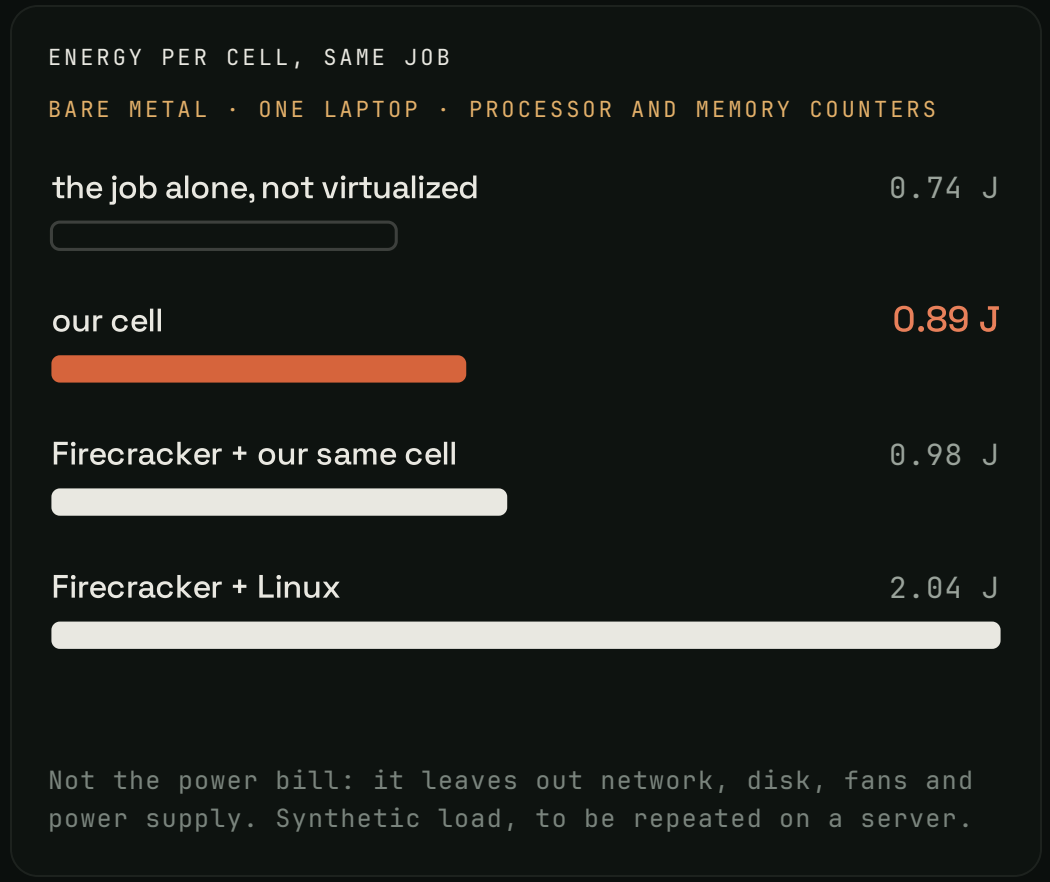
05 · SUSTAINABLE AI

Security with a measured cost. AI that does not *waste*.

Agents will multiply by orders of magnitude the actions an AI runs on real systems, and each one needs isolation and proof.

- **If isolating costs a whole operating system, the math fails.** Either you isolate little, or you burn machines.
- **With micro cells, the same work fits in less.** Less memory and less machine, in your data center or at the industrial edge, without a farm of VMs.
- **Where the consumption is.** AI's big consumption is in inference; ours is the execution side, which grows with every action of every agent.
- **Energy, measured on one machine.** For the same job of about 100 ms: 0.89 J per cell against 2.04 J with Firecracker and Linux. Processor and memory counters, one laptop, a synthetic load; not the power bill. Against Firecracker running our same cell it is close to a tie (0.98 J).

MORE DETAIL THE RUNTIME, SUSTAINABLE AI → VISION →



0.89 J

per cell, for the same job of about 100 ms

24 to 38 mJ

the whole life of one cell, from the minimal one to network and files

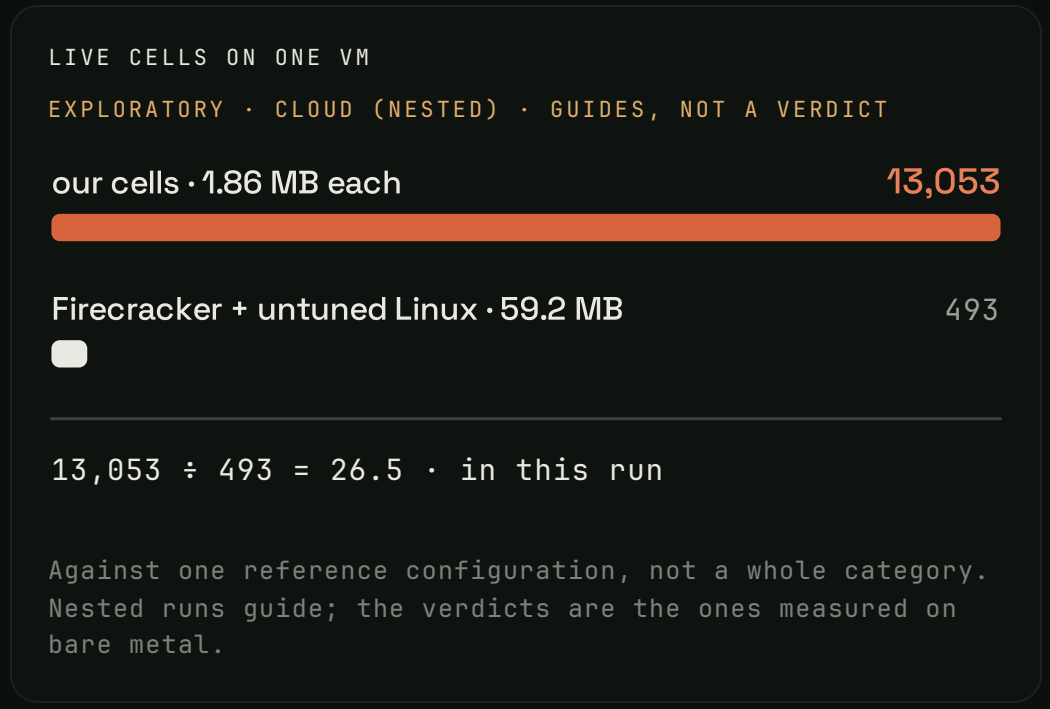
06 · DENSITY

More isolation per *gigabyte*.

On one cloud VM we kept 13,053 cells alive at once, against 493 with Firecracker and an untuned Linux. The run was cut by time at 600 s, not by memory.

- **Same load, less infrastructure.** 10,000 isolated steps alive at once: about 18.6 GB of memory at 1.86 MB per cell, against about 592 GB at 59.2 MB. Arithmetic on the memory measured per cell.
- **It is the same control.** Each step gets its minimum authority, ends and leaves evidence. Small cells are what make it affordable to isolate every step.
- **On bare metal too.** 1.03 MiB per cell with 11,778 alive on one machine; the memory ceiling was not reached. In the cloud, a burst of 20,000 cells had 0 failures.
- **What it is not.** Density does not become the same proportion of savings on the bill. That is mapped by type of workload, in the pilot.

MORE DETAIL EVIDENCE, IN THE CLOUD → THE RUNTIME →



MEASURED

- Live cells and memory per cell, in the cloud and on bare metal
- A burst of 20,000 cells with 0 failures

TO QUANTIFY

- Cost per type of workload
- Energy, and density with memory encryption

07 · WHY NOW

Agents went from answering to *acting*.

And the frameworks for critical systems already ask for records, human oversight and change control.

- **Work became the unit of risk.** An agent with a narrow task can end up touching credentials, repositories or services it should never have seen.
- **A record the agent writes proves little.** If the system that acted writes its own log, that log cannot settle who allowed what.
- **What each framework asks for.** Shown on the right, not a conformity claim.

MORE DETAIL FRAMEWORKS BY INDUSTRY →



08 · THE PRODUCT

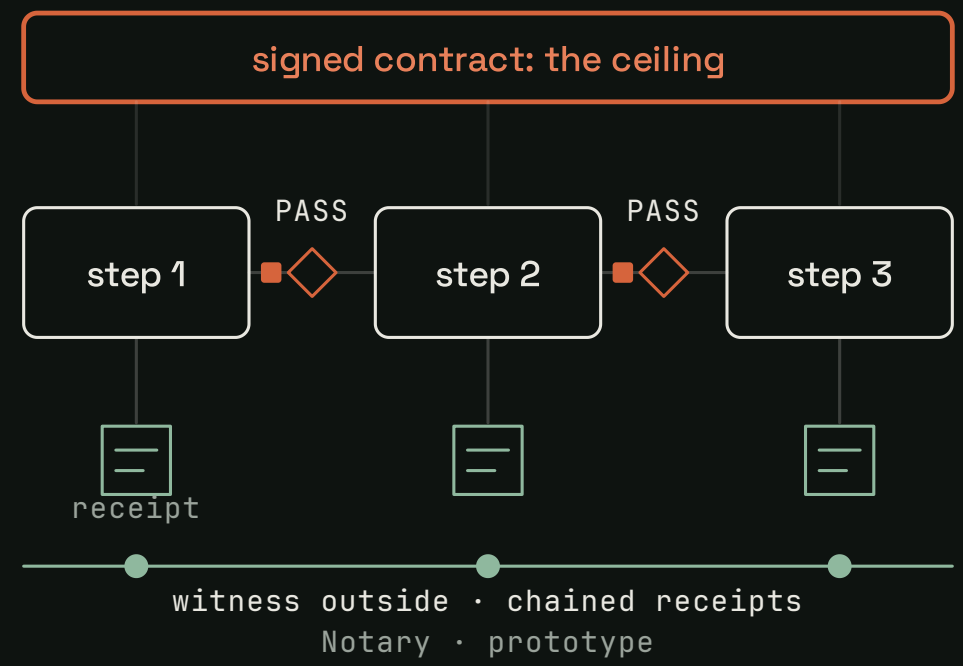
Five rules, at every *step*.

Agent Fabric sits between the agent's proposal and the real system. Each rule carries the plane where it stands today.

- **The ceiling comes first.** Authority and resources are signed before any machine exists; the Gate only activates what was precommitted. MEASURED ON METAL
- **A new micro cell per step.** No residual state from the previous cell; only promoted results are transferred. Its one-time freshness holds even after the launcher is killed and restarted. MEASURED ON METAL
- **Only what is promoted crosses.** Between steps, only the verified result; channels use short-lived mTLS tied to the job. PROVEN IN LOGIC
- **Evidence comes from outside.** Signed by the cell, countersigned by a witness outside the control domain; the full notary that seals start and close is a prototype. MEASURED ON METAL PROTOTYPE
- **No PASS, no successor.** The cell is destroyed with a signed, chained receipt. MEASURED ON METAL

MORE DETAIL HOW IT WORKS →

A JOB OF THREE STEPS — HOW IT OPERATES



Full notary: prototype. Identity per cell: proven in logic. Keys in the vault, job classification and per-cell memory encryption: design or to be validated. The effect on the target system is checked against that system, and how is agreed in the pilot.

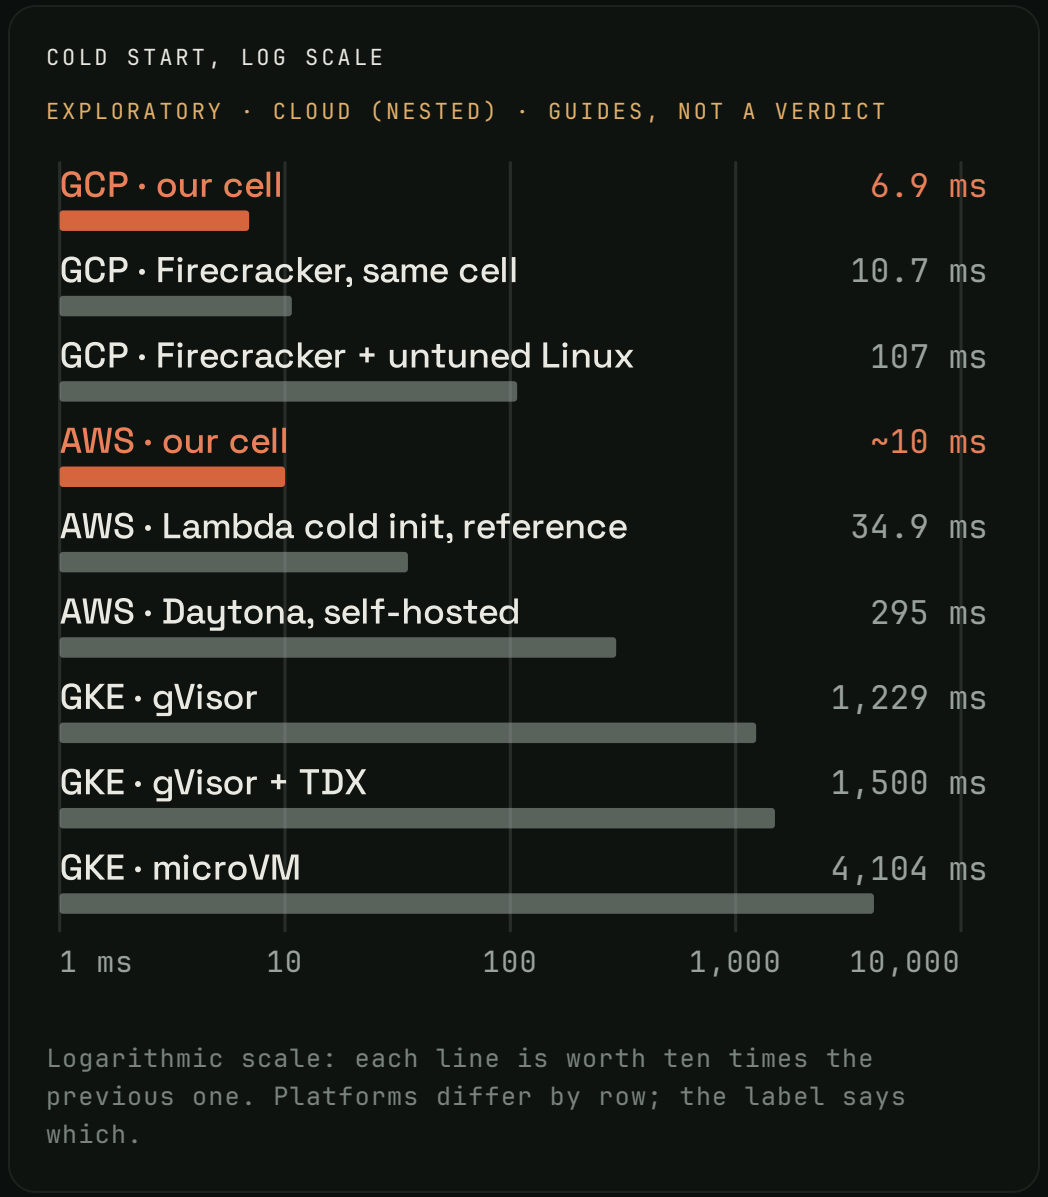
09 · MEASURED AGAINST THE KNOWN · CLOUD

In the cloud, the same *cell*.

Exploratory · cloud (nested) · guides, not a verdict. Only the cell's start was measured there, with no Gate or witness; performance and isolation verdicts are the bare-metal ones.

- **Against Firecracker.** In GCP the cell starts cold in 6.9 ms; Firecracker with an untuned Linux, 107 ms. With the same cell, Firecracker reaches 10.7 ms.
- **Managed runtimes guide, they are not a duel.** In GKE «ready» is defined another way. Lambda's cold init is a reference. Self-hosted Daytona runs with Privileged=true and runc.
- **Density in GCP.** 13,053 live cells at 1.86 MB each, against 493 Firecracker + Linux at 59.2 MB; 20,000 in a burst, 0 failures.

MORE DETAIL EVIDENCE, IN THE CLOUD →



10 · THE EVIDENCE

Measured on metal, failures *included*.

Every result carries the plane where it ran. Gaps are listed as clearly as results. On metal, one machine:

- **In the tests, an order with an invalid or repeated signature got no cell.** Written from the control domain itself, it is rejected for its signature. A signature shows who signed, not that the content is true.
- **What runs is what was verified.** One byte changed after verification: 10 of 10 rejected, and the cell never runs.
- **0 amplifications in ~6,870 attempts.** In the tests, authority did not grow under adversarial attempts. More than 70 labs published, with the ones that failed.
- **What we do not claim.** No memory encryption, one machine; with defaults Unikraft leads; no outside team has reproduced the core verdicts yet.

MORE DETAIL EVIDENCE → LAB NOTEBOOK →



0 / ~6,870
authority amplifications · bare metal

70 +
labs, failures included

11 · FOR WHOM

Six industries, each with its *critical* action.

We start with one concrete action: the one that weighs most if it goes wrong.

BANKING
Verify authorization before releasing a payment

ENERGY AND WATER
Review the setting before authorizing the change

INDUSTRY, OT AND IOT
Authorize a maintenance order from its evidence

DEFENSE AND GOVERNMENT
Record who authorized each change

TELECOM
Verify the prefixes after an approved BGP change

HEALTHCARE AND PHARMA
Authorize the batch disposition before recording it

MORE DETAIL SEE THE SIX CASES → THE BANKING FLOW →

BANKING · SUPPLIER PAYMENT — DEMONSTRATION

- The agent proposes the instruction
  agent
- Approvals and conditions
  payment owner and treasury sign
- Gate verification
  contract and clean screening
- One cell issues one instruction
  cell
- The hub submits it · no undo
  payment hub
- Settlement confirms the result
  RTGS
- Final receipt for audit
  receipt · full notary: prototype

The critical part on your hardware. The rest, in the *cloud*.

The same cell runs on your hardware and in public clouds. What touches a secret or an effect stays; what only computes can go out.

- **Hybrid is design direction.** Classifying the job at birth, the same Gate in the cloud and the federation between both sides have no code yet.
- **Confidential cloud, tested.** Intel TDX in GCP and AMD SEV-SNP in AWS: memory stays encrypted, but the cell cannot run inside, and a confidential VM per job takes 50 to 58 s. Exploratory · cloud (nested) · guides, not a verdict.
- **Memory, in the tested environment.** In the tested environment, the control domain could read the memory of a live cell. Protecting it against that actor is still pending: integration and validation on physical hardware.

MORE DETAIL SECURITY AND LIMITS →



Agent Fabric and other Kawésqar *developments*.

Agent Fabric is the main platform. Each of the others has its own state and its own action.

Agent Fabric
`tested in the lab, on bare metal; open to pilots`
AI agents acting in critical work, authorized, isolated and with a receipt.
SCOPE A PILOT →

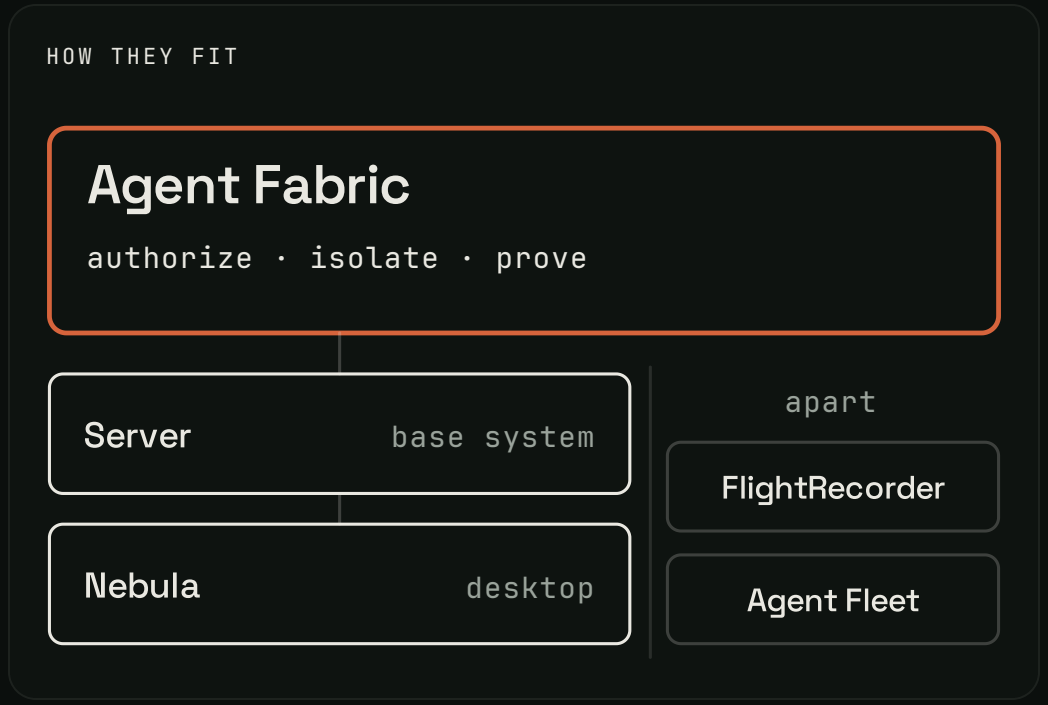
Kawésqar Server `unpublished alpha`
Hardened Linux for servers: the base system we are preparing to host the fabric.
NOTIFY ME ABOUT THE ALPHA →

Kawésqar Nebula
`alpha v0.1 · not published yet`
Hardened Linux desktop for those who run and audit systems.
NOTIFY ME ABOUT THE ALPHA →

FlightRecorder `in development`
Outbound firewall for the agents you run in Nebula.
NOTIFY ME →

Agent Fleet `private prototype`
Organizes fleets of agents with an approved plan, an org chart and an inbox of human decisions. Planning tested with real engines; the organization runs on demo data. Separate from Agent Fabric.
CONTACT US →

MORE DETAIL ALL PRODUCTS →



14 · HOW IT IS ADOPTED

A pilot on one action of *yours*.

The fabric does not replace your model, your orchestrator or your tools: it integrates with your agent and your tools, and the scope of that integration is validated in the pilot.

- **One critical action.** We choose together the one that weighs most if it goes wrong.
- **Criteria agreed before.** What the pilot has to prove, in writing, before it starts. Integration with your systems is validated there.
- **Signed evidence back.** We run your flow through the fabric and return the signed record of what each step did, with the evidence to check it.
- **Confidentiality.** Mechanism under a confidentiality agreement. Every figure reproducible on request.

MORE DETAIL SCOPE A PILOT →

A PILOT IN FOUR STEPS

01 you pick the action

before starting

02 we agree what it must prove

03 it runs through the fabric

04 you get record and evidence

No prices or licensing here: they are discussed per pilot.

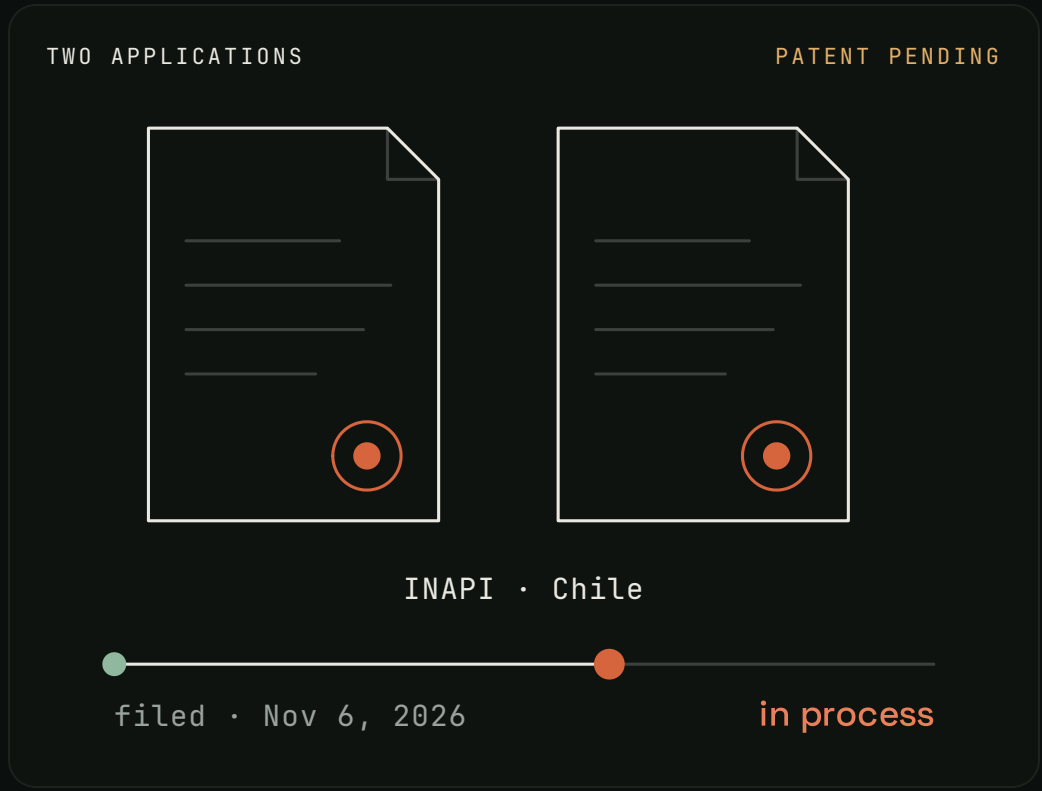
15 · INTELLECTUAL PROPERTY

Patent *pending*.

Two patent applications filed with INAPI (Chile) on November 6, 2026, in process.

- **Public: the way it operates.** We publish the shape of each property and the evidence that it holds.
- **Not public: the detailed mechanism.** It is explained in a technical meeting, under a confidentiality agreement.

MORE DETAIL HOW IT WORKS →



16 · WHERE WE ARE

What is proven, and what comes *next*.

An item moves up only with evidence behind it.

MORE DETAIL VISION → WHAT STAYS OPEN →



PROVEN · BARE METAL · ONE MACHINE

- A cell ready from scratch, and the whole step with the Gate under 20 ms
- A Gate that, in the tests, admitted only valid, signed orders
- The cell is destroyed; in the tests no traces remained in the hypervisor registry nor non-zero bytes in reused memory

IN PROGRESS

- Agent Fabric: tested in the lab, on bare metal; open to pilots
- Kawésqar Server, the base system, in testing
- The research call is open

NEXT

- Confidential-computing hardware (TDX, SEV-SNP): pending integration and validation on physical hardware
- An independent review
- First pilots and several nodes

Critical actions first, in every *industry*.

We keep a roadmap by industry. It maps the main workflows where an agent would act, starting with the actions that cannot be undone.

- **Six industries, one pattern.** In all of them the step is the same: propose, authorize, execute, leave evidence. What changes is the system that receives the effect.
- **Critical first.** Payments, changes to a plant or a network, records that someone audits. Then the less critical work around them.
- **One flow per industry is public today.** Each one is an application scenario to scope a pilot. The rest of the map is shared in conversation, under confidentiality.
- **What it is not.** Not a list of customers or dates. It is the order in which we prepare, so that a pilot does not start from zero.

MORE DETAIL INDUSTRIES →

THE FIRST CRITICAL ACTION, BY INDUSTRY — PUBLIC TODAY

Banking and markets
- A supplier payment, with two approvals

Energy and water
- A change to a protection setting

Industry, mining, OT and IoT
- A maintenance order, from an alert

Defense and government
- A patch on an isolated network

Telecom and digital infrastructure
- A routing filter change

Health and pharma
- The disposition of a batch

Application scenarios to scope a pilot, not customers. Behind each one, the workflows that follow are mapped and stay internal.

18 · WHO BUILDS IT

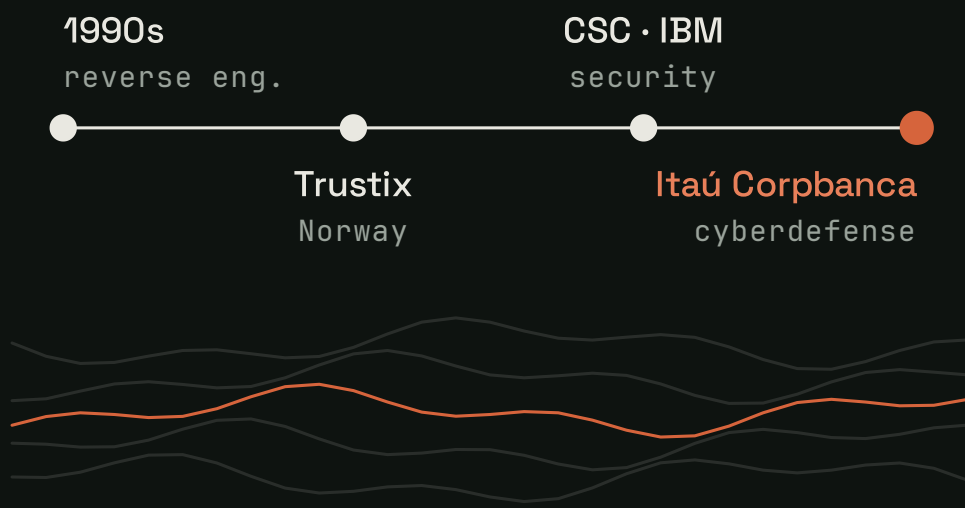
Built in northern *Patagonia*.

Led by its founder and maintainer, Rodrigo Gutiérrez: 27 years in the trade and a public track record you can check.

- **Pioneering secure Linux.** It started with reverse engineering in the 1990s; his first formal home was Trustix, in Norway.
- **Security at scale.** Security roles at CSC and IBM, then head of cyberdefense at Itaú Corpbanca, in charge of six areas.
- **A CVSS 10.0 flaw.** CVE-2004-0214, in Microsoft Windows, fixed in Microsoft bulletin MS04-037.
- **The name.** It honors the Kawésqar people, in homage, not in ownership.

MORE DETAIL VISION → THE NAME KAWÉSQAR →

A PUBLIC TRACK RECORD



SUPPORTED BY

NVIDIA Inception · Anthropic for Startups Cyber Verification Program · Auth0 for Startups · Google for Startups Cloud Program

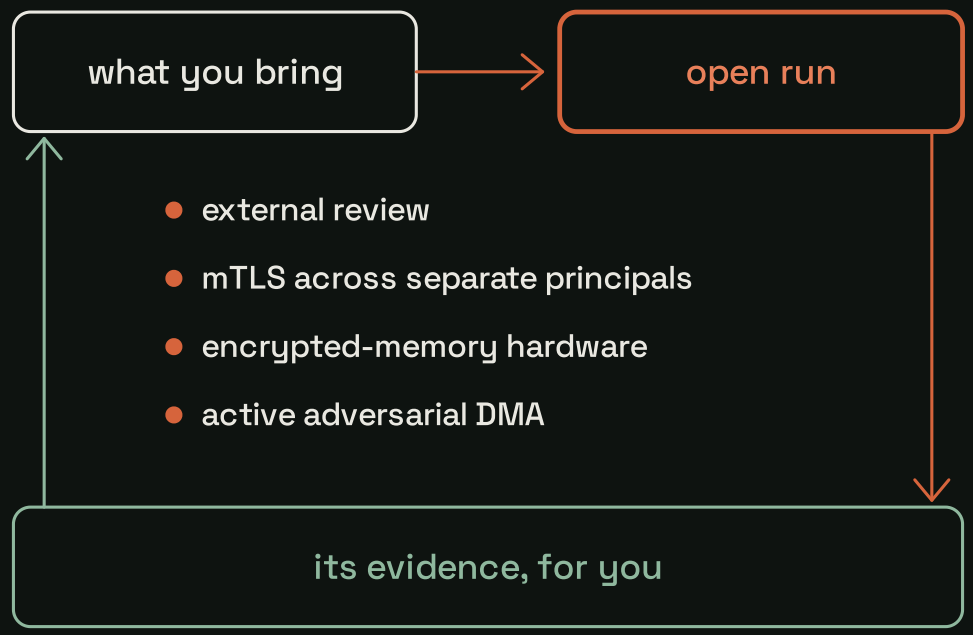
What we look for in a *partner*.

Compute, funding, infrastructure, hardware or skill. Tell us clearly what you bring and what you expect back.

- **Investors.** Documented technical progress: more than 70 labs, patent pending and the documentation of how each test was run. No customers in production; open to first pilots.
- **Physical TDX or SEV-SNP hardware.** To integrate and validate on physical hardware the protection of a cell's memory against the control domain.
- **Labs and companies.** Universities that bring infrastructure to experiment and reproduce results, and companies with an agent flow for a pilot.
- **What you get back.** The reproducible artifacts of each run you make possible: scripts, dated results and the method to repeat them.

MORE DETAIL DEVELOPMENT PARTNERS →

THE DEAL



You fund a run, you receive that run's evidence, in a form you can check yourself. Parts that would disclose the mechanism stay under a confidentiality agreement.

20 · THE ROUND

US$3 million to take it out of the *lab*.

We are raising US$3 million for 24 months. This is the back-of-the-envelope budget: where the money goes.

- **The team.** Four senior engineers, two for the kernel and two for the hypervisor, for 24 months. It is the largest line because the work is systems engineering.
- **The iron.** Our own servers with TDX and SEV-SNP, and several nodes, to validate on physical hardware what today is measured on one machine.
- **The lab.** A complete lab and an office in northern Patagonia: racks, power, network and a place to work.
- **The founder.** Full-time pay for the person who built it this far. Until today the project is self-funded, cloud bills included.

MORE DETAIL WRITE TO US AS AN INVESTOR →

USE OF FUNDS	US$ THOUSANDS · 24 MONTHS
Kernel and hypervisor engineers	1,280
Founder, full time	300
Servers and lab hardware	300
Review, patents, legal	260
Lab and office	240
Pilots	180
Cloud and compute	160
Reserve	280
Total	3,000

An estimate, to be refined with quotes. Cloud is budgeted in cash. Program credits, like Google for Startups, let us measure sooner; the budget does not depend on them.

21 · LET'S TALK

Let's *talk*.

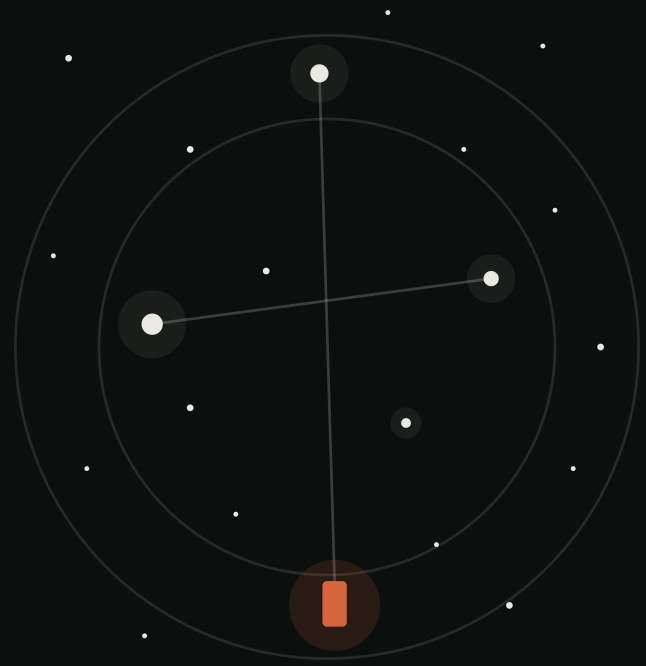
Tell us what you bring, what you expect back and what motivates you.

> The detailed mechanism is shared in a technical meeting, under a confidentiality agreement. Every figure is reproducible on request.

WRITE TO US AS AN INVESTOR →

MORE DETAIL EVIDENCE → SCOPE A PILOT →

rodrigo@kawesqar.org · northern Patagonia, Chile



21