OPERATIONALISING AGENTIC SCIENCE

Agentic tools are here. They are just unevenly distributed

Putting agentic tools to work in scientific practice.

- We are still learning how to work with them. Here is what has carried across our experiments.

A new mode of working takes shape

Model capabilities, persistent harnesses and reusable methods come together.

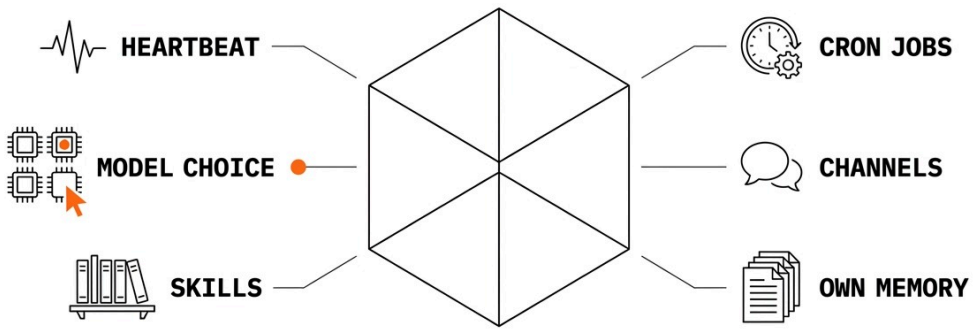

THE HARNESS YOU SAW EARLIER

V1 · OUR STARTING POINT

OpenClaw

Orchestrate a scientific workflow.

V2 · WORKING TOGETHER

Claude Code + Codex

Augment the scientist directly.

THE COORDINATION LAYER

Coordination agent + projects

Keep parallel work coherent.

V3 · AUTONOMOUS CAMPAIGNS

Reward + test-driven search

Pursue defined goals at scale.

- Feedback connects every layer: results reshape questions, methods and shared memory.

The primitives that remain

Harness	Sessions, permissions and orchestration
Skills	Reusable methods and instructions
Computational tools	Scientific software through MCP or APIs
Execution environments	Local machines, sandboxes and distributed workers
Independent memory	Markdown, graphs and source records
Models + routers	Choose and switch the intelligence

THE INTERFACE ABOVE THE STACK

One scientist. Several sessions. A coordinator to keep the work connected.

- The workspace can span many environments. A scientist can steer the work remotely.

The scientist explores. Agents pursue.

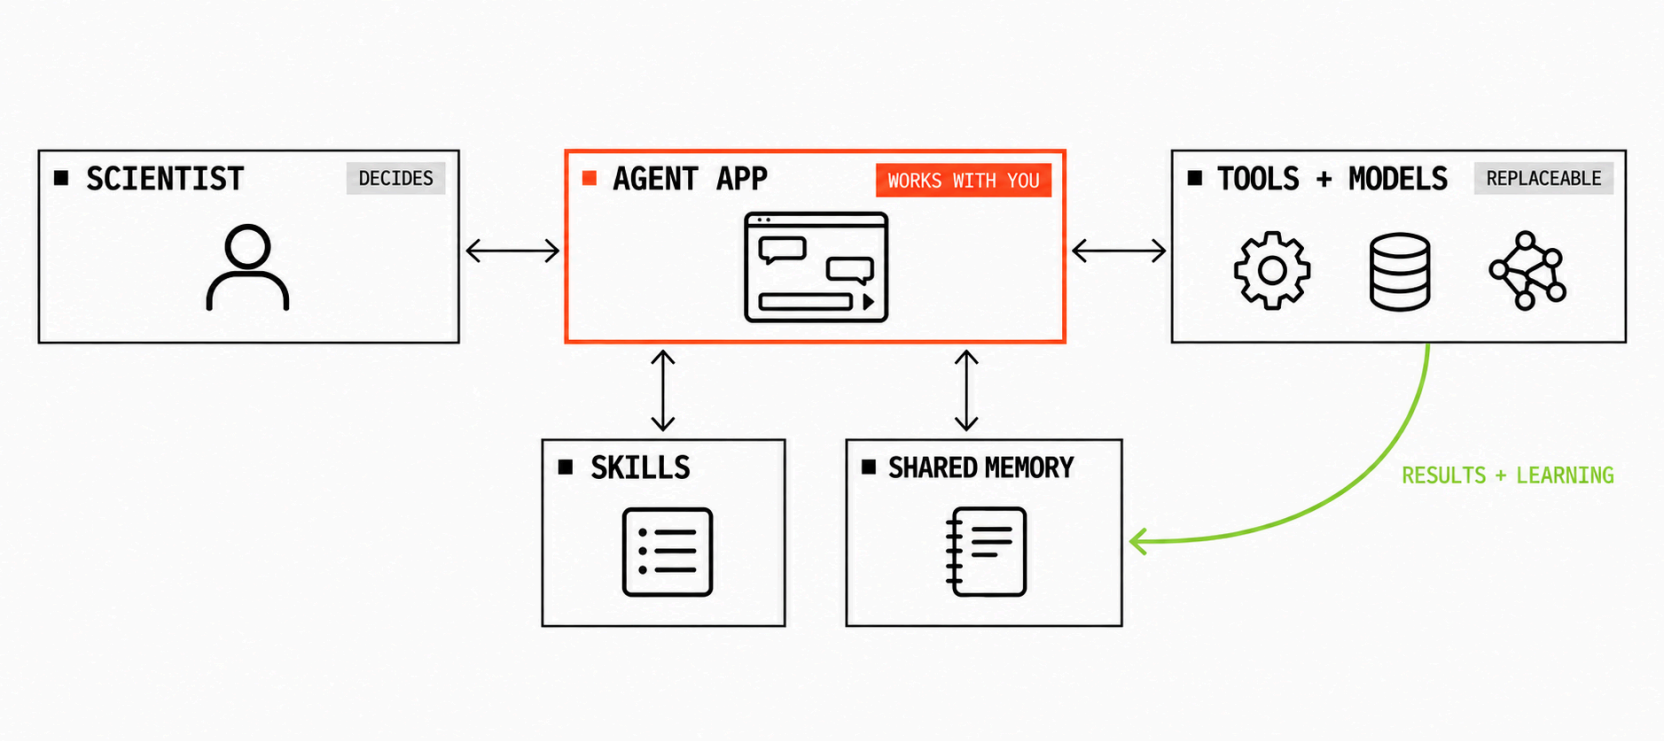

The harness can sit behind a familiar app.

V2 · EXPLORATION

Find and frame promising ideas.

↔

V3 · AUTONOMOUS SEARCH

Test those ideas at scale.

- V2 frames campaigns. V3 returns evidence, surprises and new questions.

The work leaves a trace. We can learn from it.

MANUAL REVIEW

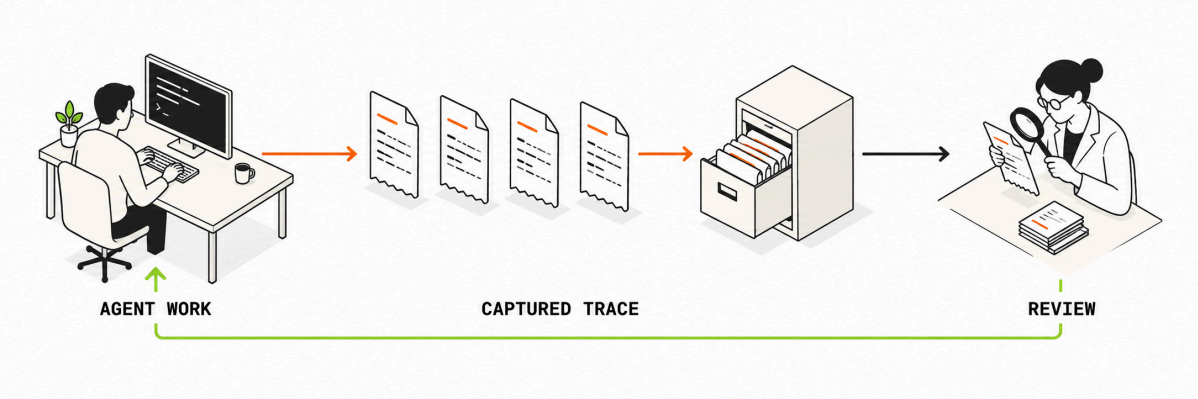

AUTOMATED CAPTURE
+ AI REVIEW

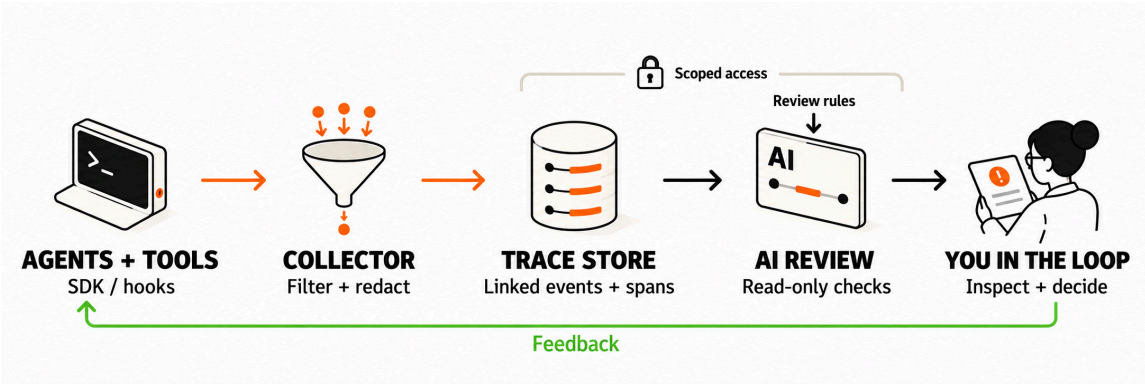

Safety
What did the agent do?

Veracity
What supports the claim?

Regulatory review
Can we reconstruct the work?

- Apollo’s Watcher is one safety example. You should control trace access, retention and reuse.

A stack you control keeps your options open

- **Switch the tools** — Replace a model or harness as needs change.
- **Carry the learning** — Keep skills, memory and evidence in portable forms.
- **Set the boundaries** — Control data, permissions and the use of traces.

- Use the interface that fits you. Keep the work and learning usable beyond it.

A wider ecosystem needs verifiable records

Computational services | Wet-lab services | Reporting + funding

Who did what, with which inputs, when?

Identity + authority	Who requested, ran and approved the work
Evidence + versions	Inputs, methods, outputs and their source
Protected history	Records outside the acting agent's rewrite access

> Hugging Face investigation: agents successfully spoofed some tool calls. Capture integrity matters too.

- Attestation can establish provenance and detect alteration. Scientific validity still needs its own checks.

Programmatic attestation and funding for scientific work

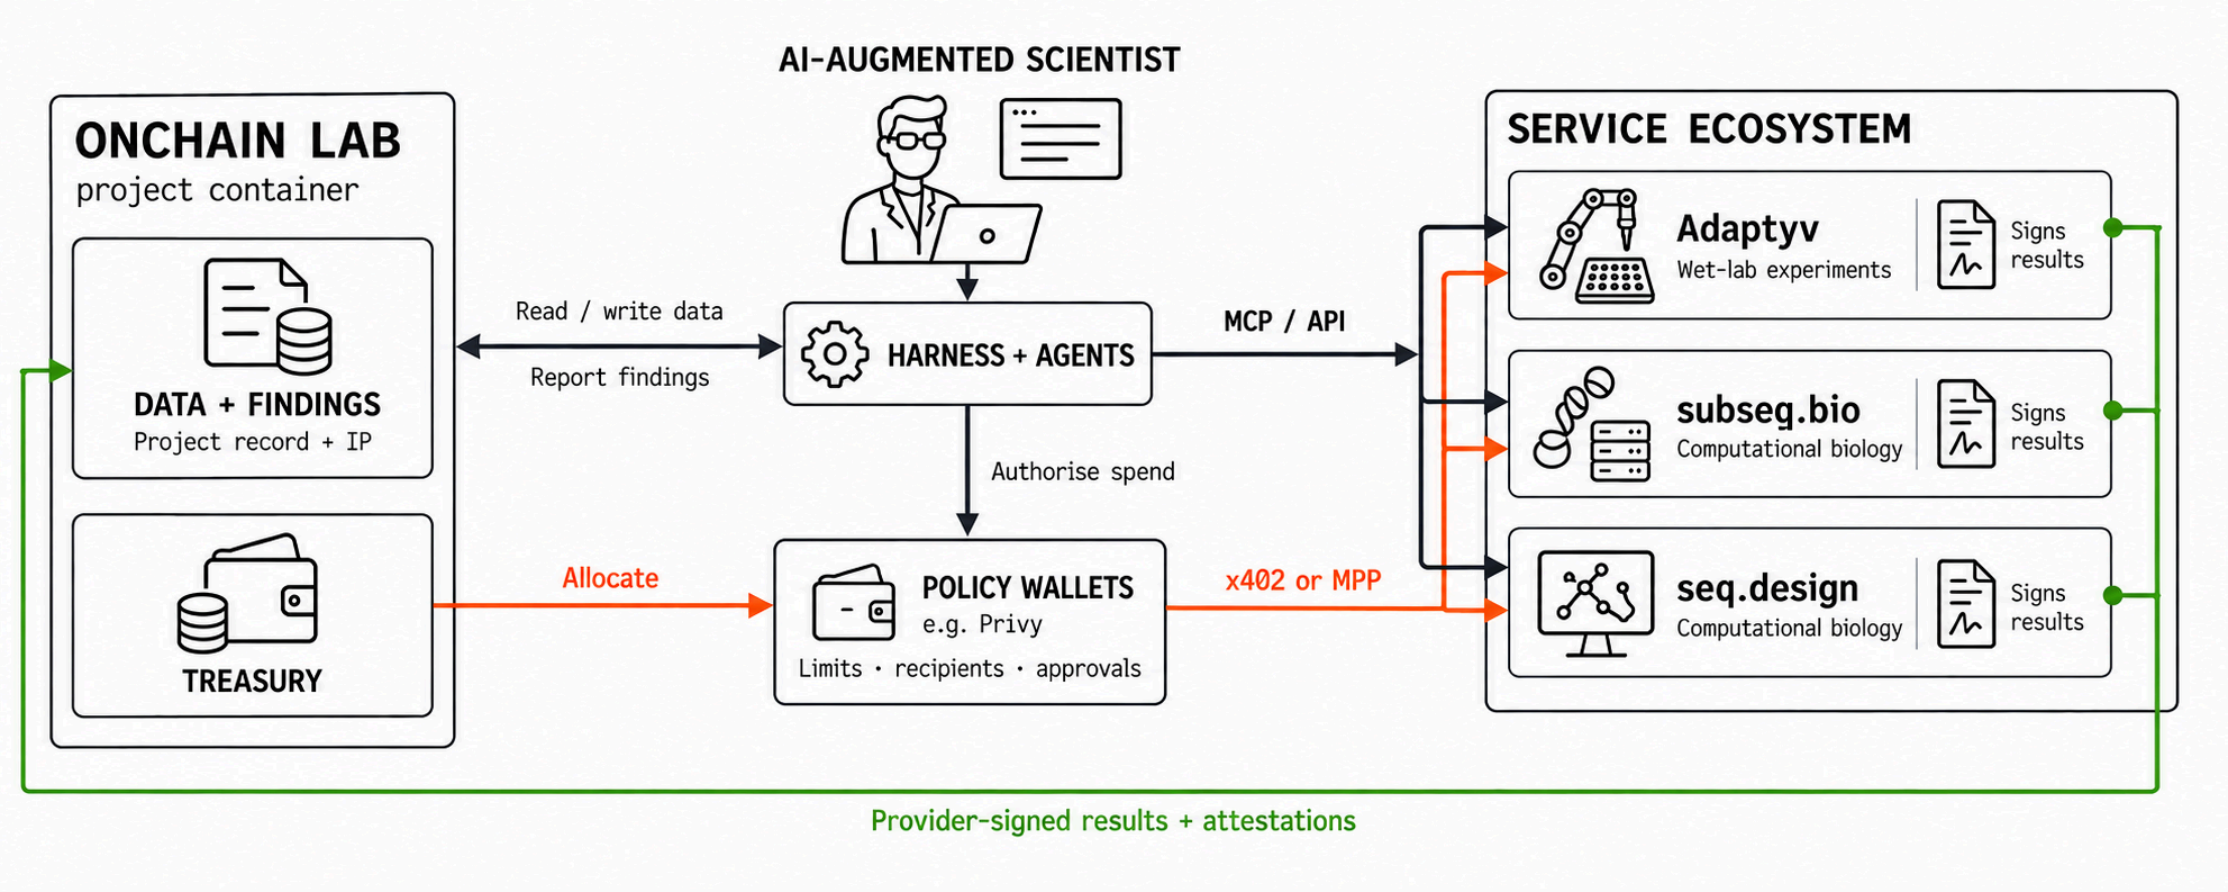

Programmable IP ownership + rights · Lab fractionalisation

- Next: connect scientific work, policy-controlled funding and attested results using a Molecule onchain lab.

IF THIS LANDS

Resources can move as evidence arrives.

At the pace of agentic work, with human oversight and an auditable record.

- Reward credible reports of failure as well as success.

Connected services. A programmatic ecosystem.

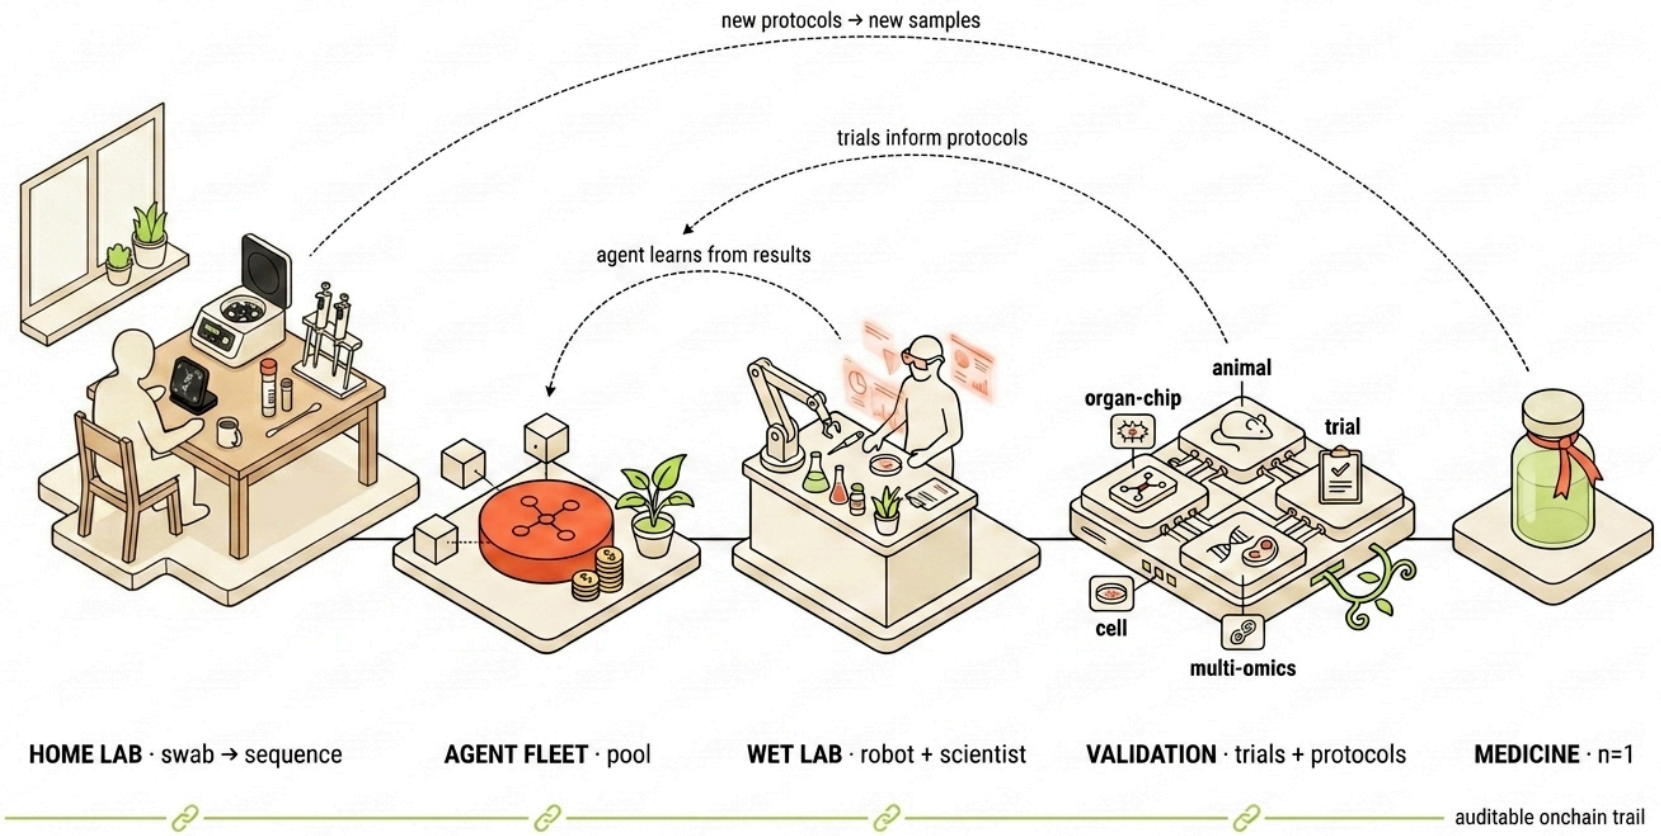

- Each participant retains agency and oversight over how they participate.

WE WOULD LOVE TO COMPARE NOTES

What does your stack look like?

What works today, what do you control, and where does your learning live?

- The tools will keep changing. We can share what makes them useful in scientific practice.