

LAB NOTEBOOK · dated



THE LAB NOTEBOOK

Everything is an agent

Through 2025 we built a platform where every component is an agent with a standard contract, and a system (SelfGrow) where agents create the skills they lack instead of failing. The agent contract became the foundation of everything that followed. Honest edge: self-growth works for small increments and still needs a human for the big architectural calls.

The dated R&D archive



Research – measured, partial

Everything is an agent

Through 2025 we built a platform where every component is an agent with a standard contract, and a system (SelfGrow) where agents create the skills they lack instead of failing. The agent contract became the foundation of everything that followed. Honest edge: self-growth works for small increments and still needs a human for the big architectural calls.



01 The entry, dated

02 What it was

03 What we built

04 What we learned — including the honest negatives

05 Where it went / status

06 Sources, glossary & the honesty ledger

Every claim carries its tier · nothing faked

PREPARED

20 May 2025

VERIFIED BY

Jon Ross — Founder, Vocabotics —
15 years building safety-critical
systems

HOW TO READ IT

Drafted with AI, verified by a
named human. Sources at the back.

01 The entry, dated



The entry, dated

May 2025

entry date

Agents

category

~1-2 years ahead

lead over the world



Public

access

What the world had at the time

Through 2024–25 'agents' mostly meant a single LLM in a loop with a fixed toolbox. A standardised agent contract where every part of the system — even the gateway and the marketplace — is itself an agent, and where agents write the tools they are missing, was well ahead of common practice.

An agent that hits a task it can't do usually fails. Ours wrote the skill it was missing and carried on — and every part of the platform, down to the gateway, was itself an agent.

The costly signal

The measured result

Every number carries its own honesty tier — the tiering is the credibility.

v1→v3

SelfGrow iterations, CLI to live dashboard

 RESEARCH

5-part

standard agent contract
(blueprint/code/settings/types/GUI)

 RESEARCH

13 in 48h

network architecture versions, AI-paced

 RESEARCH

The conceptual breakthrough of 2025 was small to state and large to build: **make every component an agent, give agents a standard shape, and let them write the tools they don't have.** Most of what Vocabotics is grew out of that sentence.

02 What it was



Two ideas, developed together:

The agent contract. Instead of ad-hoc scripts, every agent had the same five-part structure — a `blueprint.json`, the code, `settings.json`, `types.ts`, and a `gui.tsx`. Standardising the shape made agents **composable and replaceable**: you could swap one for another, generate new ones on demand, and reason about the whole system uniformly.

Self-growth. An agent that meets a task outside its abilities normally fails. **SelfGrow** instead had the agent *create the missing skill on the fly* and continue — learning by extending itself rather than stopping.

03 What we built



SelfGrow v1 → v3: three iterations, from a CLI to a real-time Socket.IO web dashboard that visualised skills as the agent acquired them.

Vocabotics Agents: the platform that made the five-part contract concrete, with a CLI, tab-completion, and dynamic agent generation.

Vocabotics Network v1 → v13: thirteen architectural versions in two days, each testing a different boundary — agent registry, identity, networking, even an LLM-based proof-of-work experiment.

Aigen / Grow: pushed "everything is an agent" to its logical extreme — the API gateway and the marketplace were *also* agents.

 RESEARCH — MEASURED, PARTIAL

The measured signal here is less a single number than a **rate**: thirteen coherent architectural iterations in forty-eight hours is a pace that is simply not available without AI-assisted development, and it let us explore the agent design space empirically instead of on paper.

04 What we learned — including the honest negatives

Self-growth has a ceiling. It works well for *incremental* skills — the agent needs a small new tool and writes it. It struggles with **architectural pivots**: deciding to restructure the whole system is still a human judgement call. An agent can grow a limb; it can't yet decide it should have been a different animal.

"Everything is an agent" is not free. Making the gateway an agent, and the marketplace an agent, is elegant — and sometimes it is just **indirection** where a plain function would have been clearer and faster. The abstraction earns its keep in some places and taxes you in others, and knowing which is the skill.

Fast iteration leaves debris. Thirteen versions in two days is a superpower for exploration and a liability for maintenance — you end up with near-duplicates that later have to be pruned.

The debris became a discipline

The near-duplicate sprawl this style produces is exactly what forced the later "cull" discipline — you cannot converge what you refuse to prune. The speed was real; so was the cost, and we account for both.

05 Where it went / status

The **agent contract survived everything** and became foundational. The self-growth idea matured into the lab's later work on a governed fleet, and the Director → Manager → Worker → Tool hierarchy (a separate, protected report) is the production form of it. This was the year "everything is an agent" stopped being a slogan and started being an architecture.

What is still open — kept visible

The honest edges of this lane, shown next to the wins. None of these is a footnote; each is a claim we have not yet earned — and marking it is what earns the right to be believed on the rest.

- ☐ Self-growth is reliable for incremental skills; it struggles with architectural pivots, which still need human judgement.
- ☐ 'Everything is an agent' adds indirection where a plain function would do — the abstraction helps in some places and taxes you in others.
- ☐ Rapid iteration (13 versions in 48 hours) explores the space fast but leaves a trail of near-duplicates that later have to be culled.

06 Sources, glossary & the ledger

Where every claim comes from — and what the plain words mean.

Sources

- 01 **Vocabotics project audit — SelfGrow v1–v3, Vocabotics Agents, Network v1–v13, Grow (2025) —**
Vocabotics internal project history (as of May–June 2025)

THE HONESTY LEDGER

- Shipped — measured, gate-clean.
- 🔬 Research — measured, partial; the gap is shown.
- 🔭 Vision — designed, honest; not yet built.

The bigger claim goes up a tier, never down. We show the eval beside the demo — the tiering is the credibility.

See the agent hierarchy in production

The Director → Manager → Worker → Tool pattern that turned this instinct into a production orchestration system (protected report).

[/lab/production-agent-hierarchy](#)

Nothing faked.

A real, dated lab report from our own R&D. Measured, not asserted; the honest negatives are kept on the page next to the wins — because a lab that hides its failures forfeits the right to be believed on its results.

