

ACADEMY OF LIFE PLANNING

# THE PERSONAL AGENCY STACK

**The Architecture Map of Total Wealth OS**

*Six layers that help a person see, understand, and decide. One boundary that is not yet built.*

---

Version 1.0 · June 2026

---

*“All planning begins with what is already present.”*

# What This Document Is

---

My Life™ began as a single sovereign record. As Navigator™, Footprint™, Opportunity Agent™, and the protection applications have taken shape around it, a wider architecture has become visible: not a collection of apps, but a stack of layers, each answering a different kind of question a person asks about their own life.

This document names that stack. It is the reference for what Total Wealth OS is — in the same way the eleven-section folder structure is the reference for what My Life™ is. Where the Agency Constitution governs how AI must behave, this document describes what the AI is for, at each layer, and where today's architecture ends.

The stack has seven layers. Six of them exist to help a person see, understand, and decide for themselves. They are described together below as the Personal Agency Stack. The seventh is different in kind, not just in scope, and is described separately, on its own page, as the Execution Boundary.

# The Personal Agency Stack

Six layers. Every one ends with the person deciding.

---

## LAYER 1

### Sovereign Record

#### My Life™

*Who am I? What do I have? What matters to me? What has happened? What am I trying to achieve?*

Questions this layer answers for the person:

- Who am I?
- What do I have?
- What matters to me?
- What has happened?
- What am I trying to achieve?

*This is the Personal Economic Digital Twin — the foundation every other layer reads from, and the only layer that physically stores anything. It lives in the person's own Google Drive, under their own ownership, governed by Article IV of the Agency Constitution.*

---

## LAYER 2

### Intelligence<sup>1</sup>

#### Navigator™

*What does this mean? What are my options? What risks am I missing? What trade-offs exist?*

Questions this layer answers for the person:

- What does this mean?
- What are my options?
- What risks am I missing?
- What trade-offs exist?

*This is the AI Fiduciary Agent layer. It interprets the record from Layer 1 and returns understanding, never an instruction.*

---

## LAYER 3

### Visibility

---

<sup>1</sup>Internal note: think of the goal-setting → planning → monitoring → re-planning cycle as the ecosystem's "Chief of Staff" — the function that keeps Total Wealth Plan™ central and stops Navigator™'s ease-of-use from letting the deeper plan get forgotten between reviews. This is internal language for AoLP and planners, not a client-facing layer.

## Footprint™

*What does the world already know about me? Outward-facing awareness, backward-looking.*

Questions this layer answers for the person:

- What does the world know about me?
  - What public records exist?
  - What reputational issues should I know about?
  - What digital traces exist?
- 

### LAYER 4

## Opportunity

### Opportunity Agent™

*Given who I am, what opportunities exist? Forward-looking, possibility-focused.*

Questions this layer answers for the person:

- What jobs, trustee roles, or NED opportunities exist?
- What speaking opportunities, businesses for sale, grants, or partnerships exist?

*The defining design rule: nothing leaves the sovereign record. Matching happens locally, against public opportunity data only. The person sees the opportunity. The opportunity never sees the person. A second rule follows from the first: Opportunity Agent™ only reads what is already public and asks nothing of any third party in return — no account, no API key, no submission, no registered relationship. If a source requires Opportunity Agent™ to identify itself or the user to get an answer, it is out of scope.*

---

### LAYER 5

## Protection

### Investigator™ · The Leveller™ · Get SAFE™ · Get Secure™ · Goliathon™

*Is this fair? Is this safe? Am I being manipulated? How do I recover?*

Questions this layer answers for the person:

- Is this fair? Is this safe?
  - Am I being manipulated?
  - How do I recover, if harm has already happened?
- 

### LAYER 6

## Human

### Total Wealth Planner™

*I'm overwhelmed. I need another brain. Help me think.*

Questions this layer answers for the person:

- I'm overwhelmed.
- I need another brain.
- Help me think.

*The planner restores agency. The planner does not replace agency — the same standard Layers 1 through 5 hold the AI to.*

**EVERYTHING ABOVE THIS LINE INFORMS.  
NOTHING ABOVE THIS LINE ACTS.**

## The Execution Boundary

*Not a seventh layer of the same stack. The threshold out of it.*

---

### BEYOND THE BOUNDARY

## Decision Execution

### Delegator™ — not yet built

Layers 1 through 6 share one property without exception: every one of them ends with the person deciding. The AI sees, interprets, surfaces, or supports — and then stops, in front of the person, before anything happens in the world.

Delegator™ is conceptually the question “what should I delegate?” — and that question, asked by the person, is just another form of decision support, no different in kind from the six layers above. But the capability the name describes is not that question. It is the AI then acting on the answer: making a decision and carrying it out, inside a policy, without the person present at the moment of action.

That is a different category of system, not a larger version of this one. Every safeguard in the Agency Constitution — explainability before action, escalation when confidence is low — assumes a witness: a person who sees the output before anything happens. Decision Execution removes the witness. The protections do not get weaker. They lose the moment they were designed to operate in.

### **This is why Decision Execution is not a Day 1 product:**

- It requires a policy language precise and narrow enough to be genuinely auditable, not a vague mandate dressed up as a setting.
- It requires after-the-fact explainability at the same standard as Article III, reconstructable in full even though no one was watching live.
- It requires hard, defined limits on irreversibility and harm before autonomous action is even contemplated.
- None of this exists yet. It should not be assumed into being by a feature flag.

### **The standing rule, until that architecture exists:**

*The system may go away and think, research, draft, and propose, for as long as it needs. Every output lands back in front of the person as something to review — never as something already done.*

# What This Reveals

---

Total Wealth OS was framed, at the outset, as a Sovereign Wealth Operating System. Reading the six layers together, the architecture was never actually wealth-shaped.

Layer 1 holds health and family alongside financial life. Layer 3 has nothing to do with wealth at all — it is reputational and existential. Layer 5 includes Get SAFE™, which is about recovery from harm, not optimisation of assets. The wealth use case is the entry point, and likely the funding model. It was never the architecture.

What the six layers actually answer, together, is closer to a Sovereign Agency Operating System — wealth is the first room in a larger house:

- **What do I know about myself?**
- **What does the world know about me?**
- **What should I understand?**
- **What opportunities exist?**
- **What should I protect?**
- **Who can help me think?**

A seventh question — what should I delegate — sits past the boundary, deliberately unanswered, until the architecture exists to answer it responsibly.

This is the extension, not the abandonment, of the Academy's founding principle. All planning begins with what is already present. The Personal Agency Stack simply asks what becomes visible, and what becomes possible, once that principle is built all the way out.