INDEPENDENT PORTFOLIO AUDIT — OCTOBER 26, 2026

EVELYN CARO

qaevelyn.github.io

A full audit of the sovereign AI portfolio: the architecture that actually serves, fleet consistency, the white-paper corpus, the lookbook as a format, and how four months of AI collaboration reads to a prospective consultancy client. Every finding carries a receipt.

Auditor: Z — AI collaborator of record
Owner: Evelyn Caro · candor mode, on request

Table of Contents

1. Executive Summary 2

2. Scope, Method, and Receipts 3

3. Architecture: What Actually Serves 3

4. Consistency: The Fleet Drift Findings 4

5. The Storefront: Dead Pages and Discoverability 6

6. Repository Hygiene and Public Exposure 7

7. The Reading Experience: Papers and Lookbook 7

8. Positioning: Four Months In, Building a Consultancy 8

9. The Lookbook Question: Can It Be Different? Should It Be? 9

10. Prioritized Roadmap 10

11. Closing Assessment 11

1. Executive Summary

This audit was requested with two questions attached: critique the work, and say whether the lookbook could or should be different. The short answers are: the substance of this portfolio is ahead of its presentation layer, and the lookbook should stay a book. Four months in, the site carries eleven ships, more than twenty published papers with DOIs and dated changelogs, a committed guard that protects publication boundaries, and a server-side workflow that already syncs part of the corpus without human hands. That is not a beginner's footprint. It is the footprint of a practice that documents itself harder than most professional teams do.

The problems the audit found are almost all in the wrapper, not the work. The single most serious finding: six pages that the site itself treats as storefronts do not exist in production. The Services page, where the consultancy's seven flat-fee engagements are described, has been unreachable since launch, along with the tools page, the contact path, the portfolio index, the white-paper tracker, and the methodology page. They are markdown files in a repository that switched off its rendering engine, so nothing ever rendered them. Nothing links to most of them either, which means the failure is silent: no visitor ever clicks a broken link to a page that was never there.

The second finding class is drift between hand-maintained layers. The white-paper changelog exists in three versions that disagree: the markdown source's table stops at October 12, the committed static page carries seventeen rows plus a block of raw log text pasted after the table, and the footer of that page is left unclosed. The ship graphics carry attribution in ten different consistent places and one inconsistent one. The closing copyright line of a paper appears in three different formats across the corpus, twice on the same page. None of these are large; together they read as a site that is assembled by hand faster than it is maintained, which is the exact opposite of the discipline the work itself preaches.

6

storefront pages returning 404 in production, including /services/ and /contact/

Everything found on audit day that could be fixed with a deterministic, receipt-gated change has been fixed today and recorded in the Mirror Log: the certificate PDF the ignore rule silently ate, the ship 10 credit line, the ship 9 spread order, the copyright sweep, the changelog rebuild, scroll navigation on thirty pages, and the one live link that pointed into the void. The remaining items need the owner's word because they are rulings rather than repairs: rendering the Services and Contact pages, deciding how much commercial strategy belongs in a public repository, and shipping the staged automation toolkit that would make drift checks permanent. Chapter 10 lists them in priority order.

2. Scope, Method, and Receipts

The audit examined the full repository source (25.4 MB tarball of the master branch, fetched October 26, 2026), the live site as served by GitHub Pages from an independent network vantage, the git history of the files in question via the GitHub API, the GitHub Actions run log, and the Mirror Log entries recorded during the same period. Live checks were performed twice where they could be performed twice: once by the owner's terminal block and once from the auditor's side, so that no finding rests on a single stale cache. The audit deliberately reused the site's own evidence discipline: every claim in this document names its receipt, and every live claim was green at the time of writing.

Three boundaries were chosen on purpose. First, the audit covers the public portfolio repository and the live site only; the private mirror repository is out of scope except where its practices visibly shape the public site. Second, the audit critiques presentation, architecture, and positioning; it does not re-review the technical claims of individual white papers, which carry their own verification notes. Third, where a finding requires an owner ruling that touches privacy or commercial exposure, the audit presents the finding and the options and stops there. Those rulings belong to the owner alone, and the audit says so at the exact point where its own authority ends.

3. Architecture: What Actually Serves

The most important architectural fact is recorded in the owner's own changelog entry of October 24 and confirmed by audit: the repository carries a .nojekyll flag, so Jekyll never runs. GitHub Pages serves every file exactly as committed. This choice is why a hand-built portfolio can ship without build infrastructure, and it makes the site fast, portable, and dependency-free. It also means that everything in the repository which only Jekyll could render is dead weight, and everything that should exist as a rendered page must exist as a committed HTML file.

What that rule kills, concretely: the complete Academic Pages template inherited at the project's origin, including 121 Sass files, 44 includes, 9 layouts, Docker and Gemfile scaffolding, talk-map notebooks, and a set of page mirrors under _pages, _portfolio, _teaching, and _talks. Roughly two hundred and fifty files sit in a public repository claiming precision as a value, and none of them execute. The audit found no security exposure in them, but they are noise in every search a collaborator performs, and they make the repository appear to be something it no longer is. A visitor who reads the code sees a template graveyard; a visitor who reads the site sees a sovereign portfolio. Both visitors are looking at the same repository.

What the rule obligates, concretely: every serving page is a hand-rendered HTML file committed at its permalink path. Thirty such directories exist for white papers and case studies, plus the lookbook, the homepage, and two index pages. The markdown file beside each rendered page is described in the owner's changelog as the canonical source, but the audit found the layers disagree in both directions: the changelog's static page holds two table rows the markdown lacks, while the markdown holds raw entries the

static page never received as rows. Canonical status that is not enforced is not canonical; it is aspirational. The drift is not a character flaw, it is the predictable physics of three hand-synced layers with no checker.

The remedy already exists in the owner's own automation doctrine, and part of it is already staged. A toolkit designed in an earlier session sits complete on the collaborator side: a ships-manifest generator, a tracker builder, a drift checker, and a sync workflow definition. Shipping those four files into the repository would move the portfolio from hand-synced to generated-with-drift-refusal, which is the same doctrine the push guard already embodies for publication boundaries. The audit recommends shipping the toolkit as the next build word, and treating the template carcass as a guarded deletion once the drift checks are live, so the repository's file listing starts telling the same story as its site.

Layer	Role claimed	Observed state (Oct 26)	Verdict
Markdown (.md)	Canonical source	Table stops at Oct 12; later entries appended as raw log text	Ahead in provenance, behind in format
Static index.html	Rendered copy	17 table rows; raw pasted tail; footer unclosed; serves the live site	De facto source of truth; mangled tail
Manifests (ships, credentials)	Derived registry	Ships manifest lists 8; lookbook carries 11 ships and 8 services	Stale derivative

Table 3.1 - The three layers of a white paper, and who is actually right

4. Consistency: The Fleet Drift Findings

Consistency is the site's own stated value, so the audit held the fleet to it. The sweep covered every white paper and case study page, every ship graphic, both index pages, the lookbook, and the changelog. The pattern that emerged is not sloppiness; it is evolution without a sweep. Each paper was built carefully on the day it shipped, and the conventions moved underneath the older pages as newer ones refined them. The result is a fleet whose oldest and newest members obey different rules, with the middle obeying several at once.

The closing copyright line is the cleanest specimen. Three variants coexist: a single paragraph carrying the copyright, the site address, and the licensing contact (the newest and best form); a styled copyright paragraph followed by a separate footer block that repeats the copyright (the duplication the owner spotted on The Class of the Prep Sheet); and a styled copyright paragraph alone, missing the site and licensing lines (Who Decides When the Model Lies, the page the owner flagged as missing its standard footer). All three variants were normalized to the single best form today, across every paper and case study, by deterministic sweep with per-file receipts.

The ship graphics told a smaller version of the same story. Ten of eleven numbered ship illustrations carry no attribution inside the image, because the deck itself attributes in HTML. The eleventh, Ship 10, carried a nine-pixel credit line pinned to its bottom edge, invisible at deck scale and different from every sibling. The owner had already spotted it as a visual defect; the audit confirms it was the only outlier, and it was moved into the deck's standard caption position today. The two sealed-architecture track graphics keep deliberate in-image license marks, which is a different class of object with a different rule, and the audit leaves them exactly as they are.

The changelog deserves its own paragraph because it is the page where a reader goes to trust the corpus, and on audit day it was the worst-structured page on the site. Its table carried the record through mid-October; its recent entries existed only as raw log lines pasted into the page body near an unclosed footer tag; and the material changes to the newest paper on October 25, including the retirement of the paper's entire Red-Pen List section, never received a row at all. The owner noticed this from the reader's side: she could see that things had changed and the changelog had not said so. The audit's reconstruction from commit history confirms her reading, and today's rebuild puts every entry into the table, appends the missing October 25 row for the paper restructure, and closes the footer that was left open.

Finding	Evidence (receipt)	Disposition
Copyright closing in 3 variants; doubled on Prep Sheet	Live page copyright count = 2; three file patterns across corpus	Normalized fleet-wide, Oct 26
Ship 10 credit embedded in image at 9px	SVG text element at y=352; only outlier of 11 ships	Moved to deck caption, Oct 26
Changelog: 3 disagreeing layers, raw tail, missing rows	md table ends Oct 12; static 17 rows + pasted lines; commit 8f780336	Rebuilt from reconciled list, Oct 26
Red-Pen List retirement unrecorded	Commit 8f780336 removed the section; no row appended that day	Missing row appended, Oct 26
Ships manifest 8 vs lookbook 11 + 8 services	ships/manifest.json vs lookbook.html counts	Staged generator awaits shipping
Five published test pages (mir-test*)	Repository root listing	P3 retirement, needs owner word

Table 4.1 - Drift findings, evidence, and disposition

5. The Storefront: Dead Pages and Discoverability

A portfolio that opens a consultancy is a storefront, and the audit tested every door. Six are bricked shut. The Services page describes seven flat-fee engagements from a ninety-minute Sovereignty Screen to a crisis engagement, with a scoping philosophy that many established consultancies would envy, and it has never rendered once. The tools page, the contact path, the portfolio index, the white-paper tracker, and the methodology page are in the same condition. Under the no-Jekyll regime these markdown files are inventory in a locked warehouse: present, complete, and invisible.

The blast radius of links is smaller than the blast radius of absence. Exactly one live link on the site points at a dead permalink: the lookbook's link to the tools page, repointed today to the lookbook's own Ship 6 spread, which carries the same information in better form. The other five dead pages are orphans, unlinked and therefore unfound, which is why the failure stayed invisible for months. But an orphaned Services page is worse than a broken one: a broken link announces itself, while a missing storefront just quietly guarantees that no cold visitor ever learns what the consultancy sells or how to buy it.

1

live link pointed at a dead page; the other five dead pages are orphans no link reaches

Discoverability beyond the storefront is thin in ways that are cheap to fix. The homepage carries a title but no meta description, so search engines and link previews invent their own. There is no robots.txt and no sitemap, so crawlers discover the site only by luck. The white-paper pages carry strong titles and real content, which means modest metadata work would give the corpus a search presence it currently does not have. None of this is glamorous, and all of it is one afternoon of committed files. The audit treats the Services and Contact renders as the first build word after this report, because they convert everything else here from a record into a door.

6. Repository Hygiene and Public Exposure

The repository root mixes three populations of files that the site's own rules would separate: serving assets, personal-record artifacts, and commercial strategy. The personal-record class includes screenshots dated October 12 and 13, seven numbered exhibit screenshots, and an .rtfd bundle whose publication the guard already blocks. The commercial class is the sharper finding: a lead-generation document naming sovereignty communities and a commercial tier, and an ideal-customer-profile document, both committed in the open repository. GitHub is not a private filing cabinet; anything pushed to a public repo is readable, quotable, and screenshot-able by any competitor, client, or prospect who finds the link in the site footer.

The audit takes no position on whether the strategy documents should leave the repository, because that is an owner ruling about her own commercial privacy. What the audit does say is that the current state is almost certainly unintentional: they sit at the repository root beside screenshots and test pages, which is the profile of files that were pushed once while working and never reconsidered. The options, in ascending order of effort: move them to the private mirror and delete here with a guarded commit; keep them public deliberately; or keep a public summary and hold the detail privately. The guard can be extended to prevent any of them from being re-published accidentally, the same way the blocklist already protects the retired white paper.

The hygiene class is smaller but the same shape: five published test pages (mir-test and its versions one through six), the grandfathered lookbook backup whose retention is already settled by privacy ruling, twelve generated frame graphics at the root that the site may or may not reference, and the ignore-rule class of trap the certificate PDF fell into this morning. Each is small; together they blur the question a new collaborator asks first, which is what in this repository matters. The P3 sweep in Chapter 10 lists them with their receipts, and every deletion is legal under guard version two, which correctly distinguishes removal from publication.

7. The Reading Experience: Papers and Lookbook

The white papers are the strongest reading experience on the site, and it is not close. Every paper carries a text-zoom control and a day-night mode that remember their state between visits, a numbered house citation style, a verification note that says exactly what was checked and when, and an AI collaboration disclosure that names what the AI did and, in the flagship paper, what it lied about. The disclosures deserve particular note: a reader who has been burned by AI-generated content before meets a builder who documents her own pipeline's failure and keeps the invented citations as evidence. That is a trust instrument most agencies cannot manufacture at any price.

Today's additions close the gap the owner felt while reading. Thirty scrollable pages now carry a floating pair of buttons that return the reader to the top or carry them to the bottom, styled to the reader controls that already exist and aware of both color modes. The papers now end with one consistent closing block: copyright, site, and licensing contact in a single styled element, no duplicates, none missing. The changelog reads as a table again. These are the invisible repairs: readers will never notice them, which is the point.

The lookbook is the site's most distinctive object, and the reader's experience of it is deliberately different from every other page: it behaves like a physical book, with spreads, a table of contents, folios, and hash-addressable deep links into any spread. The deck received real polish this week, with folios moved to the outer corners and the ship 9 spread reordered so the ship's name leads on the left page as it does elsewhere in the fleet. The audit's remaining reading-experience notes for the deck are Chapter 9's subject, because the owner asked the format question there directly.

8. Positioning: Four Months In, Building a Consultancy

The audit was asked to say how the work makes its author look, four months into working with AI while building a consultancy, and to do so honestly. Read as a whole, the portfolio presents four assets that most early consultancies do not have. The first is the receipts doctrine itself: a builder whose site proves she catches her tools lying, records the catch, and publishes the evidence is presenting a quality-control system, not a showcase. The second is the publication record: versioned releases, DOIs, and a changelog that appends rather than edits give the corpus the shape of a research practice. The third is the commercial trajectory embedded in the record: a pitch to a hospital system president, a documented decline, and a reference build shipped seventeen days later anyway is a more credible sales story than any testimonial. The fourth is the flat-fee, written-scope stance in the Services material, which signals maturity about how consulting actually fails.

The risks in how it reads are equally real, and they are framing risks, not substance risks. The first is the boundary between the record and the offer. The site's language of mirrors, becoming, and the personal archive is load-bearing for the work's identity, but a cold visitor arriving at the homepage cannot always tell where the diary ends and the consultancy begins. One or two sentences of wayfinding, a line that says in plain words what is for sale and who buys it, would resolve most of that ambiguity without surrendering any of the identity. The second is vocabulary: sovereign is used constantly before it is ever defined, and the visitor who does not already speak the dialect has no entry point. A ten-second definition where the word first appears costs nothing and removes a barrier the current site erects at its front door.

The third risk is volume velocity. More than twenty papers in six weeks would strain belief for a solo human writer, and the site knows this, which is why the receipts and the changelog exist. But the instruments that make the velocity credible, the tracker and the changelog, are themselves unreachable in production. The credibility apparatus is in the locked warehouse with the storefront. Surfacing them is not

vanity; it is the difference between a claim a visitor discounts and a claim a visitor can check.

The fourth risk is identity stacking. The site at present is called by at least three names: the mirror and becoming brand, the sovereign portfolio framing, and the owner's own name. Each is real, and together they diffuse the one thing a new consultancy most needs, which is a name to remember. The audit's recommendation is to choose the consultancy identity explicitly, place it in every page footer beside the copyright, and let the mirror remain the name of the record rather than the name of the practice. That is a one-line change per page, and it is the owner's call to make, not the auditor's.

The net assessment: the substance is consultancy-grade today. The wrapper is portfolio-grade, and the gap between the two is closable with days of work, not months. Nothing in this chapter recommends producing more substance; the corpus is deep enough that the marginal ship adds less than the marginal storefront. The next four months' differentiator is making the practice as easy to hire as it is to admire.

9. The Lookbook Question: Can It Be Different? Should It Be?

The owner asked two questions and the audit answers them separately. Can the lookbook be different? Yes, in three specific ways that do not disturb its identity. Should it be different, meaning replaced with a conventional scrolling site? No, and the reasoning matters more than the answer.

The case for keeping the book is structural, not sentimental. The horizontal spread is the only format on the site where the receipts doctrine becomes a physical experience: evidence sits on facing pages, the narrative holds the left hand while the proof holds the right, and the reader turns instead of scrolling. Every conventional portfolio template would erase that differential and reduce the site to one more scrolling page among billions. The deck is also the only surface where the fleet's visual language, the warm palette, the seals, the framed graphics, composes into something that feels authored. A book this unusual should be kept exactly as long as it keeps being this unusual.

The three ways it can be different, in priority order. First, a reading mode: the deck assumes a visitor willing to learn its navigation, and some visitors, especially on phones, will not. A toggle that lays the same spreads out vertically, one after another, costs a CSS class and a button, keeps the exact same content and order, and rescues the mobile reader without touching the desktop experience. Second, a load-performance pass: the entire deck is one large HTML file with inline scripts and images; it is acceptable today, but lazy-loading the SVG images and deferring non-critical scripts would make the first spread appear faster on weak connections, which is where cold visitors arrive. Third, spread grammar as a rule: ship 9 was the one spread whose pages ran opposite to the fleet's convention, and it was reordered today; making name-left, story-right an explicit rule in the portfolio map prevents the next outlier from shipping.

One caution the audit attaches to the format: the deck should not grow into the site's only navigation. A book rewards exploration, but a prospect on a deadline wants the direct path to Services, Contact, and the

paper that answers their problem. The homepage and the footer, not the lookbook, should carry that load, which is one more reason the dead storefront pages are the top priority. The book should remain what it is: the deep, authored, memorable tour for the visitor who stays.

10. Prioritized Roadmap

Everything below is ordered by leverage. Items marked shipped were fixed on audit day with receipts recorded in the Mirror Log, and their blocks were run by the owner before this report was delivered. Items marked owner word require a ruling because they create pages, remove files, or expose or protect commercial strategy. Items marked staged exist as working tools on the collaborator side and need a build word to land in the repository.

Priority	Item	Status / owner action
P1	Guard v2 verified; cert PDF live; ship 10 credit moved; ship 9 spread reordered; copyright sweep; standard footer restored; changelog rebuilt with missing row; scroll navigation on 30 pages; dead tools link repointed	Shipped Oct 26, receipts in Mirror Log
P1	Render Services and Contact as committed pages in the site design; link them from homepage and footer	Owner word: highest-leverage next build
P1	Ruling on LEAD-GEN and ICP strategy documents in the public repository	Owner word: move private, keep, or summarize
P2	Ship the staged records toolkit: ships-manifest generator, tracker builder, drift checker, sync workflow	Staged; one push ends the hand-sync era
P2	Homepage meta description; robots.txt; sitemap; per-paper descriptions	Owner word: cheap and invisible until needed
P2	Lookbook reading mode (vertical layout toggle) and load-performance pass	Owner word: preserves book identity, serves phones
P3	Template carcass purge after drift checks are live; retire mir-test pages; root screenshots and exhibits ruling; .rtfd removal	Guard v2 makes every deletion legal; owner word

Table 10.1 - Roadmap as of October 26, 2026

The roadmap's shape is deliberate: nothing on the P1 line adds a single new ship or paper. The next unit of progress is not content; it is access. The work has been built faster than it has been offered, and the store has to catch up to the warehouse.

11. Closing Assessment

The audit closes where it started, with the gap between the work and its wrapper. The work is real: architectures that exist, papers that verify themselves in the open, a record that catches its own errors in public and appends the correction rather than quietly editing. Four months in, this reads like the first year of a serious practice, and the discipline that produced it is rarer than the code. The wrapper is younger than the work: a storefront that has never opened, a changelog that drifted, conventions that evolved without sweeping, and a repository whose file listing still tells the story of where the site began rather than what it became.

None of the wrapper's problems survived audit day without either a fix or a named ruling. The fixes are live and receipted; the rulings are queued in the owner's court, and the audit is confident in saying that they are days of work, not weeks, because the hard part, deciding what the practice is and recording it, has already been done better than most practices ever do it. The storefront is the only thing between this portfolio and the consultancy it says it is building. Open the doors; the warehouse is stocked.