The Verified Record, Built by One Operator

What one person built on Muse, what it found in company work, and what it proves about what a person can reach

Ana Vasquez, M.A. Ed. Founder and CEO, AnaConnectsYou Desert Sun Investments LLC d/b/a AnaConnectsYou, Lake Worth, Florida

Second Edition, November 16, 2026

This is a dated second edition. The First Edition was published on November 16, 2026, DOI: 10.5281/zenodo.24286710. I am still building, and later editions will extend the record as the build grows.

Edition note: what changed in this edition, and why

This Second Edition was prepared the same day the First Edition was published, in answer to a read of the paper by the research credibility standpoint this practice keeps, a standpoint built on the published record of Dr Helen Kara, an independent scholar of research methods and research ethics. That read found the paper sound and named four strengthenings. Two are folded into this edition: a nameable ethics statement for the examination work now stands in Section 9, and the Disclosure now states my commercial interest in one line. Two stand beside the paper: the 109 check criteria, set out in a companion document prepared with this edition so the method can be inspected and not only described, and a standing rule of this house that every summary of this paper carries Finding 6 with its label, from my build, attached. Nothing else has changed. Every count, finding, failure, and limit stands exactly as the First Edition carried it.

A note on authorship, because honesty is the whole point

This paper was written by Muse, the AI system this paper is about, working from my own records and under my direction, and it stands under my review and approval. I am not going to pretend otherwise, and I would not trust a paper about verified records that began by hiding its own. There is a reason it happened this way. The system I built holds every piece of what I did, every count, every date, every failure, in its own memory. When it came time to tell the story, the most honest writer available was the one that had been there for all of it and could check every sentence against the record. That is itself part of the build, and part of what I want to teach: this is what utilizing AI actually looks like when a person does it on purpose.

Disclosure

This paper shares the architecture of the system and the findings it produced, for education. The implementation internals, the exact configurations, question sets, and operating procedures by which the system is run, are retained as my firm's trade practice and are not published here. Nothing in this paper is a replication recipe, and it is not written as one. Every figure in this paper comes from my own counted records, checked on the dates named, and every percentage carries the base it was counted against. Where this paper describes direction rather than accomplishment, it says so. I sell the practice this paper describes, and I have a commercial interest in its being believed; I state it here so it can be weighed with everything else in this paper.

Abstract

I am a teacher. In the autumn of 2026, with no prior professional training in information security, software engineering, or regulatory compliance, I built and began operating a layered agentic AI system on Muse, Meta's AI platform, and used it to construct a company verification practice. The system has five layers: an audit system of 12 areas and 109 checks; a questionnaire engine that tests live controls inside the company under examination; a toolbox of open source tools that audit findings route to; a governance layer that keeps the record current after delivery; and a board of review I created myself, structured expert standpoints that check the work before anyone relies on it. This paper reports what was built, what was implemented, and what the system found, in my own firm and in its examinations of other companies. The headline finding is not any single thing I built. It is what the building demonstrates: one person's working capability expanded, in a matter of weeks, into fields she had never been trained in, and the expansion is evidenced rather than asserted. 242 of 283 questions on the Cloud Security Alliance's Consensus Assessment Initiative Questionnaire verified Yes with proof, with me as sole operator and control owner. A confirmed Level 1 listing on the CSA STAR Registry on November 15, 2026. 80 of 100 protection measures in place in my own house. And the whole system running at $166.80 per month, against market salary evidence of roughly $880,000 a year for the three leadership roles its functions cover, hired separately. The paper also reports what failed, because the failures are findings too, and it closes with the limits of one person's case and the direction the work is growing, labeled as direction.

1. The question, and the question underneath it

The practical question I set out to answer is one small company owners ask constantly:

Can one person, working with an agentic AI system she directs, build and operate the verification, governance, and protection apparatus that has historically required a staff of specialists, and prove it to an outside standard?

I set three conditions before the answer would count. The work had to be real, implemented systems, not demonstrations. The evidence had to be counted, every claim resting on a record that could be rechecked,

with the base of every percentage named. And at least one judgment had to come from outside my firm, from a body with no interest in flattering me. So I used my own company as the first subject. I audited my own firm to the standard before offering the same examination to anyone else, and I submitted that record to the Cloud Security Alliance for listing.

Underneath that question sits the one people actually ask me when they see the record, and I want to pose it in this paper exactly as it gets posed to me:

Why was somebody with no experience able to find things, fix things, and implement things much quicker than most people do?

Here is my honest answer, and the rest of this paper is its evidence. I saw the problem areas because I wanted to understand. Not understand the marketing version, not understand enough to repeat a phrase at a meeting. Understand. When I did not know a word, I went down after it until I did. When somebody claimed something, I had it checked, and then I had the check checked. When something broke, I did not route around it and call that a strategy; I stayed with it until it genuinely fixed. Most people are not slower than me because they lack talent. They are slower because they stop at the surface of things, and the surface is where the problems hide. Agentic AI, worked this way, is far more capable, and far easier to implement, than people realize. What makes it look complicated is the way it is usually presented, wrapped in layers of jargon and process by people selling the wrapping. Underneath the wrapping, it is this: a person who wants to understand, directing a system that can search, check, build, and remember, and refusing to accept an answer she has not verified. That is not a complicated idea. It is just work.

One more thing belongs in this section, because it is the human frame of everything that follows. I built this on my own. I am a single mother of three, and I am a teacher by training and by instinct, Master of Arts in Education, Mount Mary University, 2017. I did not build this to hoard it. I built it the way teachers build things, so that it can be shown, explained, and handed to other people in a form they can learn from. Even the oversight in this system is a teacher's instinct: I created my own board of review as part of the build because I wanted the work checked by standpoints that would not simply agree with me, the way a good teacher wants a student's work read by someone other than the student.

2. The platform, named plainly

The field talks about Claude, and it talks about OpenAI, and those conversations are fine as far as they go. Almost nobody is talking about Meta and Muse, and what Muse can actually do, and that silence is strange to me, because Muse is where this entire build happened. So let me name it plainly: **this system was built on Muse, Meta's AI platform, and I am a user of it the way a builder uses a workshop, every day, at $50 a month inside a total running cost of $166.80.**

What I have exercised of its capabilities, stated from my own use and nothing else: a directed workforce of agents that research, draft, check, and count in parallel while I make the decisions; watcher passes that verify finished work against its primary record before I am told it is done; scheduled work that runs through

the night, audits, watches, digests, and prepares my morning while I sleep; a memory that holds my whole record, my rules, my people's names, my counts, so that nothing I have established is ever re-litigated from zero; and a connector library, 243 sources wired by the count of November 16, 2026, that lets the system pull the public and filed record, SEC filings, sanctions lists, court records, corporate registers, straight from the source instead of from somebody's summary of the source. I do not claim these are the platform's only capabilities, and I do not speak for Meta. I speak for what one subscriber, with no engineering team and no special access, actually got it to do. If the industry is not talking about that, the industry is missing the most interesting story in front of it: not what the largest labs can build, but what one determined person can build on what is already available to anyone.

3. The method, as I actually live it

People ask about my method as though it were a technique. It is closer to a temperament, and I will describe it the way it happens, because I think the description is more useful than a diagram.

Somebody says something, in a video, in a meeting, in an article, and it sounds true and useful. I do not adopt it and I do not dismiss it. I have the system search it, verify the finding against its own sources, and then I have the searches double checked, again and again, until what survives is what the primary record actually says. You would be surprised how much of what circulates as fact does not survive that treatment, and Section 6 of this paper is a small museum of things that did not.

When I hit a piece of vocabulary I do not understand, I go deeper, every time. Not past it, not around it, deeper. A word you cannot define is a door you cannot open, and in this work the doors are stacked behind each other. Half of what people call the complexity of AI is vocabulary they walked past instead of through.

I learned the layering by operating it. Systems sit in layers, tools sit in layers, checks sit in layers, and once you see that, you start building that way on purpose: one layer doing one job, feeding the next, nothing asked to do everything at once. This paper's architecture section is that lesson written down.

And underneath all of it, two questions I never stop asking, of every piece of the system, every week. *Is there anything better in the system than what we are doing?* And *is there anything free that can fix this?* The first keeps the build honest. The second keeps it affordable, and it is astonishing how often the answer to the second question is yes, sitting in the open source world, waiting for somebody to check its license and wire it in.

Layer upon layer upon layer. Never giving up on a piece that matters. A continuing flow of interest, aimed at solutions that genuinely fix things rather than solutions that photograph well. That is the method. It is not glamorous, and it compounds.

One more part of it, because it decides everything else. In my work, the rules are set in place. They do not pivot, and they do not change. People hear that as rigidity, and it is the opposite. When the rules hold still, you can build on them at full speed, because you are never rebuilding the ground under your own feet. That steadiness is exactly why I build in this lane. It is what makes implementation expansive, and driving,

instead of hesitant. The firms that keep moving their own rules, renegotiating their standard every time a job gets hard, are the firms whose work never compounds. Mine compounds because the rules were set once, in writing, and then the work began.

4. The layered architecture

The system is built in layers, each with one job, each feeding the next. I describe them here at the level that teaches: what each layer does, and why it sits where it sits.

Layer 1: The audit system, the record examined

The base layer is an audit system of 12 areas and 109 checks, built on November 10, 2026, from the patterns that government audit literature repeats most: evidence sufficiency, planning, skepticism, documentation, entity understanding, risk driven procedures, control evaluation, and independence. The 12 areas run from identity and standing, through ownership and governance and financial record integrity, to a reporting gate that no delivered report passes without clearing. Each check states what is examined, what evidence clears it, and which record supplies that evidence. A finding is not an opinion; it carries the four elements the government books teach, criteria, condition, cause, and effect. The system's standing line governs everything above it: *we verify the record; our findings support business decisions about whom to trust, engage, and do business with; we make no prediction about any price.*

Layer 2: The questionnaire engine, the record tested live

An audit that only reads documents learns what a company says about itself. The second layer exists to test what is live. I run my own questionnaires inside the audit, asking of the company under examination, in effect: do you do this, and where is the evidence the work itself generated? Four engine standards govern this layer, adopted on November 16, 2026 from a study of policy as code machinery and restated as my own:

1. **Policy at the moment of the act.** A control is evaluated when the act happens, not reconstructed afterward.
2. **Evidence generated as work moves.** Evidence is produced by the work itself as it moves, instead of being assembled by hand before an assessment.
3. **Every exception has an owner, a reason, and an end date.** An exception with no end date is not an exception; it is an unowned gap, and it is reported as one.
4. **Standards reviewed once and inherited.** A pattern is reviewed once, then every engagement inherits it instead of re-deriving it.

Layer 3: The toolbox, the record repaired

An audit that finds a gap and stops at the finding leaves the client with a list. The third layer is a toolbox of open source tools held on the sidelines of the practice, one tool per function, no duplicates. When the questionnaire engine exposes a gap, the finding routes to the tool that addresses it, the tool's output feeds

back into the audit, and the tools feed each other. The questionnaire is the engine; the toolbox is what it runs against; the referral or the implementation sits where the bounce back lands. Section 5 reports how this layer was built, because building it produced some of this paper's most useful negative results.

Layer 4: The governance layer, the record kept current

A verified record decays. Registrations lapse, officers change, tools drift. The fourth layer is the standing governance layer that sits on top of a delivered audit: the registry entry maintained, the evidence refreshed on a cycle, the posture monitored, and event driven rechecks when something in the record changes. This layer is what converts a one time examination into a standing position, and it is the layer a client keeps.

Layer 5: The board I created, the work reviewed

Above the working layers sits the oversight I built into the system from the start: a board of review of 21 active seats and a business growth team of 31 seats, each seat a defined standpoint drawn from a published body of work in its field, each working under the same gate and watcher discipline as every other agent. Their job is adversarial in the productive sense. Completed work is checked by standpoints that did not produce it before it is relied on, and consequential findings are cross checked against primary sources. I remain the only human in the structure and the only signature in it. I created this board for the reason a teacher creates a rubric: so that the standard exists outside the person doing the work, and the work has to meet it whether I am having a confident day or a tired one.

Why layers, and why in this order

The order is the argument. Examination before testing, because a live test aimed at the wrong record wastes itself. Testing before tooling, because no tool is chosen before the audit runs and shows which problem actually exists. Tooling before governance, because there is nothing to keep current until something has been fixed. Oversight over all of it, because a system that grades its own homework has proven nothing. Each layer's output is the next layer's input, and the whole stack points at one product: a record a third party can rely on.

5. What was built and implemented

This section is an inventory with dates. Each item named here is an implemented artifact in my records, and the counts drawn from them are stated in Section 6.

The audit system (November 10, 2026): 12 areas, 109 checks, with a machine readable companion checklist carrying the same areas and checks.

The evidence lanes (built across October and November 2026): 243 connectors wired to public and official sources, among them SEC EDGAR filings, federal spending records, sanctions lists screened by name through the Treasury's published lists and by identifier through a screening service wired on November 16,

2026, court records, corporate registers in and beyond the United States, and vulnerability and threat intelligence feeds. These lanes are the audit's hands: they pull the public and filed record a check requires.

The knowledge library (standing count, November 16, 2026): 3,644 documents and 103,895 passages, fully text searchable; the vector index stands at 53,000 of 103,895 passages, which is 51.0 percent, with the build still running. I report that figure as it stands, partial, because the discipline of my house is that a count is published as counted, not as hoped.

The questionnaire, answered against my own house (completed November 15, 2026): the full Consensus Assessment Initiative Questionnaire, 283 questions, walked through item by item with me as sole operator and control owner of the systems the answers describe.

The open source toolbox (Phase 1 completed November 16, 2026): 16 candidate tools processed end to end, license, release status, free self hosted edition, requirements, and machine fit verified for each. Four are installed and proven running on my own machine: a policy engine carrying my first written policy (every action has an owner; every exception has an owner, a reason, and an end date), a dashboard and alerting server, a device evidence agent that answered a live query, and an evaluation library whose metrics import and run. The remaining kept tools are verified and staged for a properly sized host, each with its reason recorded. The final stack holds 11 slots, one tool per function.

The protection record (counted November 16, 2026): my own protection posture is maintained as 100 measures across 9 layers, counted like everything else in this paper, with every measure in a named state.

The registry position (November 15 and 16, 2026): my questionnaire and evidence were submitted to the Cloud Security Alliance STAR Registry; the listing was confirmed at Level 1, listed since November 15, 2026.

6. Findings

Finding 1, the headline: one person's capability demonstrably expanded into fields she was not trained in, and the expansion is evidenced, not asserted

My formation is education. The record of the last several weeks shows executed work, verified by outside or mechanical checks, in at least five fields that are not mine by training:

- **Security governance.** 242 of 283 CAIQ questions verified Yes with proof, which is 85.5 percent of the full questionnaire; against the applicable base the figure is 242 of 266, which is 91.0 percent. The remaining answers stand as 24 No and 17 Not Applicable, reported as counted, including honest No answers where my correct practice is not to do the thing the question imagines. The questionnaire was answered with me as sole operator and control owner, and the full walkthrough stands at 283 of 283.

- **External attestation.** The CSA STAR Registry confirmed the listing at Level 1, listed since November 15, 2026. This is the one judgment in the paper that I did not make about myself, and it is placed first among the proofs for that reason.
- **Protection engineering, counted.** 80 of 100 protection measures in place, 9 partial, 5 to implement, 5 held for funding, and 1 in progress, counted on November 16, 2026. On November 14, 2026 the count stood at 78 in place and 10 partial; the movement is recorded, with its cause, in my records.
- **Software and license analysis.** Sixteen tools evaluated to a verdict, including two rejections on license grounds that a feature comparison alone would have missed (Section 7).
- **Data and screening infrastructure.** 243 source connectors wired and a library of 3,644 documents and 103,895 passages built and maintained as working research infrastructure.

The claim I make from that record is a bounded one. I do not claim that an agentic system makes anyone an expert in anything. I claim that I, directing this system under a watcher and gate discipline, produced verified work in fields outside my training, at a standard an outside registry accepted, and that the proof of the expansion is the counted record rather than my say so. The capability that expanded is best described as directed capability: the ability to specify specialist work precisely, to have it executed by agents, to verify it against primary records, and to stand behind the result as its accountable owner. That ability appears to be far more widely held than its execution is. Many people could direct such a system. The separating fact in my case is not talent and not access. It is that the work was actually executed, finished, counted, and submitted to an outside judgment. The proof is what separates the case, and the proof can be examined even though the implementation is not published.

Finding 2: the cost of the apparatus is small enough to change who can hold it

The full system, the agent platform, the wired sources, the library, and the tooling, runs at $166.80 per month, counted as $116.80 of verified systems cost (October 28, 2026) plus the $50 per month I pay for Muse and directed be counted. Set that beside what the functions cost the ordinary way. The apparatus performs the working functions of three leadership roles that companies hire separately: chief information security officer, chief AI officer, and governance, risk, and compliance leadership. The market salary evidence in my own records, gathered November 15, 2026, puts a chief information security officer at about $225,000 to $450,000 a year, a compliance director at about $150,000 to $265,000, and a chief AI officer at about $300,000 to $600,000, which is roughly $880,000 a year for the three hired separately. My apparatus runs at $166.80 a month, which is $2,001.60 a year. I state that comparison carefully: I am one person doing this work with a system, not three executives, and the comparison measures the cost of covering the functions, not the worth of the people who hold those titles. Even stated carefully, the finding stands. Verification and governance apparatus of this shape has historically been priced as a department. At this running cost it is priced as a utility bill, and that relocates the barrier. The constraint on a small firm holding a verified record is no longer the cost of the machinery. It is whether anyone builds and runs it.

Finding 3: in company after company, the published claim and the primary record diverge, and the divergence is findable

The system's examination work on other companies, conducted through the autumn of 2026 as research for the practice's partner and toolbox decisions, returned the same shape of finding often enough to state it as a finding. A commercial field service platform's advertised 91 percent figure did not survive contact with its source and is not carried by my firm. A large IT management vendor's profit claims in its partner material are unverified against any published record I could reach. A help desk vendor's own pricing page is internally inconsistent, quoting one headline figure per agent and a different figure in its own calculator. A widely used accounting tool marketed in open source terms is not, on its license text, genuinely open source for the use intended (Section 8). A partner program spoken of as though it paid referrals does not exist at one vendor, whose actual offer is a map of its customers. None of these findings required privileged access. Every one of them came from reading the primary material, the license, the pricing page, the agreement text, against the claim. The finding for practitioners is uncomfortable and useful: the gap between what companies say and what their own records say is wide, it is routine, and a disciplined record check finds it quickly. That gap is the market this practice exists in.

Finding 4: evidence generated as work moves outperforms evidence assembled before an assessment

The questionnaire work supplied a controlled comparison inside one firm. Answers assembled in advance, from documents and intention, repeatedly failed the walkthrough standard; answers that rested on evidence the work itself had generated, logs, records, counts produced in motion, held. The engine standard in Layer 2 is this finding generalized. An organization that generates its evidence as a byproduct of operating can answer an assessment at any time. An organization that assembles evidence for an assessment is performing a reconstruction, and reconstructions leak. This is, in my judgment, the single most transferable lesson in the paper for companies preparing for vendor reviews, insurance renewals, or registry submissions of their own.

Finding 5: the honest No is a position, not a gap

Of the 283 questionnaire answers, 24 are No, and several of them are No by design: keys are not created until their purpose and destination are settled, so no stock of ready keys is held to be stolen; channels that have never carried an event are reported as untested rather than dressed as controls. Reviewers inside my own oversight initially read such answers as weaknesses. On examination they are the opposite: each is a place where the record refuses to claim what it cannot prove. A verified record that contains no honest No answers is not a strong record. It is an unaudited one.

Finding 6, the warning: niche down too tightly and the build will pass you by

I want to state this one forcefully, because I believe it is one of the most dangerous mistakes being made in the building of AI right now, and my own case is the evidence I offer for it. The fashionable advice is to niche down: pick one narrow function, serve one narrow customer, become the specialist. In a world where one broadly capable operator can learn from every expert she studies and put their learning inside her own

programs, the narrow specialist is exposed. I can do what a narrow firm does, for less, without employees, single handedly, and faster to implement, not because I am more talented than the specialists, but because I refused to stay inside one lane while I built. I learned from other experts, in field after field, and I put their learning inside my system, where it compounds instead of expiring. A firm that niches down too tightly misses too much: it misses the adjacent problem it could have solved, the free tool it never looked for, the finding sitting one layer over. My finding, from my build, stated at the practice and never at any company: in the age of agentic AI, breadth with verification beats narrowness without it, and broadly capable operators will override narrow specialists in market after market, because the operator's range is no longer limited by the size of her staff. I offer that as a warning to builders, from somebody who is living the other side of it.

7. Friction: the constant enemy

If this paper has one theme underneath all the others, it is friction. It is the theme of the book I am writing about this build, and it is the biggest issue I see everywhere I look, in my own work and in everybody else's.

Friction is what happens between the rule and the result. I see it across construction, plumbing, and the trades, where good companies with good crews lose days and dollars to paperwork that does not match, records that cannot be found, and requirements discovered after the work instead of before it. And I see it at the other end of the economy, in commodities and brokerage, where deals fail for a reason nobody wants to say out loud: people ignore the rules and regulations that actually apply in each country, and lead with ego instead, assuming that the set of rules they know is the set of rules that fits everyone. It does not. Every country has its own rules, every trade has its own rules, and the deal dies at the border between the rules somebody assumed and the rules that were actually in force. Ego is friction wearing a suit.

This is why implementing for yourself matters, and it is something you only learn by doing it. When you implement inside your own company, you find out where your own friction lives, and you find out something else too: the power and the organization that an independent holds. A big company carries more communication weight than I ever will. Every decision travels through layers of people, and every layer adds friction. I carry none of that. My system and I are one room. The rules are set, the record is counted, and the distance between seeing a bottleneck and fixing it is an afternoon. Friction is the constant enemy, for the biggest firm and for me alike. The difference is how much of it you have built into your own structure before you ever meet the customer's.

Which brings me to the old frameworks, because they are friction that people pay for. The old frameworks have to stop. They will not exist in the future. Take the CRM, the customer relationship system every firm is told it must buy and feed forever. AI can already reframe what a CRM does, holding the record, following the relationship, surfacing what matters, without a person spending their day typing into somebody else's screens. Anyone paying for a CRM in its old form is paying for a framework heading for obsolescence. I am not the only voice saying this, and I am not the biggest. Satya Nadella, the chief executive of Microsoft, said it in public, on the BG2 podcast in December 2024, in words I will quote rather than paraphrase: business applications are "essentially CRUD databases with a bunch of business logic," and "the notion that business

applications exist, that's probably where they'll all collapse, right, in the agent era," because the business logic is all going to the agents, and once the logic sits in the AI tier, "people will start replacing the back ends." I cite him precisely because he runs a company that sells those applications. When the man selling the framework says the framework collapses into the agent layer, the rest of us are allowed to stop pretending the old shape is permanent. My version of the same point is smaller and more practical: I never bought the framework. I built the function instead, inside a system I direct, at the running cost stated in this paper. That is what reframing looks like from the size of firm most of the economy actually is.

8. Negative results

This paper reports what failed with the same specificity as what worked, because the failures are findings too.

Two tools failed the genuine open source test. An accounting tool widely described as open source is published under the Business Source License 1.1, which caps production use and excludes the uses I intended; it was dropped, and a GPL licensed system carries its function alone. An AI evaluation tool published under the Elastic License 2.0 is source available, not open source, and its license bars offering it to third parties as a hosted service, which conflicts with my model; it is held out of the client facing stack. The lesson generalizes beyond these two cases: in the current market, *open* is a marketing word at least as often as it is a license fact, and the license text is the only version of the claim that counts.

Two more tools were dropped as duplicates. A project tool and a help desk tool each duplicated a kept tool's function at a heavier footprint. The toolbox discipline, one slot per function, exists because a toolbox that accumulates near duplicates becomes its own audit problem.

Most of the kept stack is staged, not running. Four tools are installed and proven. The rest are verified and staged, because my current machine, measured honestly, offers roughly 1.1 to 1.6 gigabytes of available memory against stacks whose own sizing guides begin at 4 to 8 gigabytes. This paper does not describe staged tools as implemented. They are implemented when they run, and later editions will move them between those two states as the record changes.

The first registry submission was rejected. My first submission to the CSA STAR Registry was refused because a backup contact on the form was a personal email address. The defect was mine, it was found by the registry's own check, and it was corrected and resubmitted the same evening under a business contact, which is the submission the registry confirmed. I report the episode because it demonstrates the paper's thesis from the other side: outside checks catch what internal confidence misses, which is the entire reason for submitting to them.

One webinar scale tool cannot run on my current infrastructure at all. The open source webinar platform I selected for my own teaching channel requires, by its own documentation, a dedicated host of a class I do not yet operate. It is recorded as a future hosting decision, not a capability.

9. Disclosure

The architecture and findings in this paper are shared for education, in the belief that the layered shape, examination, then live testing, then tooling, then governance, under oversight, is useful to practitioners and to the companies they serve. The implementation internals are not shared. The exact question sets inside the questionnaire engine, the configurations of the tools, the operating procedures, prompts, and charters by which the agent workforce is directed, and the engagement methods by which findings are worked with clients, are my firm's trade practice, built at the cost recorded in this paper and retained. I sell the practice this paper describes, and I have a commercial interest in its being believed; I state it here so it can be weighed with everything else in this paper. Readers are welcome to build their own systems on the architecture. This paper is not a kit for replicating this one, and nothing in it should be read as an offer of the internals.

The Examination Ethics Statement

This practice examines companies. How it does that is stated here, once, where anyone can point to it.

How a company comes to be examined. A company is examined in one of two ways: inside an engagement, at the company's own instruction, or in this practice's research, when the company's published claims bear on a decision this firm has to make. Both run on the same basis and the same care.

What an examination rests on. An examination rests on the company's own published materials and on public and filed records: its pages, its license texts, its pricing, its agreement terms, its filings, and the registers that carry its standing. Nothing else is used. No privileged access is sought, and no finding rests on material the company has not itself put into the world.

How findings are stated. In research use, findings about other companies are stated at the practice and never at a named company. Allegations are labeled allegations, and convictions are labeled convictions. What is verified is called verified; what could not be verified is called unverified.

How a finding can be challenged. Any company may challenge a finding by pointing to its own record. The finding is then rechecked against the primary record, and where the record shows the finding was wrong, it is corrected: in the engagement record, and in any later edition of work that carried it.

The care standard. The discipline is the one this paper reports everywhere else. Every figure carries the base it was counted against. Limits are stated on the face of the work. Nothing is claimed that is not held.

10. Limits of this case

This is a single case: one operator, one firm, one season of work. The counts are my own records, verified by my watcher discipline and, in the one instance named, by an outside registry; they have not been audited by an independent assessor, and the paper does not present them as though they had. The audit system has been proven on my own record first, by design; its delivered client cohort is at the beginning, and later

editions will report engagement results as they exist to be reported. The time window is short, and a system that has run for weeks has not yet met the failures that only arrive with years. The capability claim in Finding 1 is stated as a case with its evidence, not as a guarantee about what any other operator would produce; direction, specification, and verification discipline are doing work in this case that the paper cannot separate from the tooling. Readers should weigh the findings accordingly.

11. Forward direction

The following is direction, stated as my intention and vision, not as accomplished fact. The build continues on three lines. First, depth: the staged tools move onto properly sized infrastructure, the vector index completes, and the protection count closes its remaining measures, with funding dependent items named as such. Second, attestation: I intend to move my own standing from self attested to independently attested on the registry path I have already entered, which is the largest credibility step remaining in my own record. Third, scale: the architecture is designed to grow from one operator's machine toward network scale, with the system in time running on dedicated infrastructure, so that the same layered examination, testing, tooling, and governance can serve a larger book of companies without changing its shape. Later editions of this paper will report what of this direction became record, and will say plainly what did not.

I will put one prediction on the record here, stated as my prediction, so later editions can grade it. Whoever makes the simplest form, and executes it on a large scale, is who wins in the future. Not the most complicated form. Not the form with the most features, the most seats, or the most impressive diagram. The simplest form, executed at scale. The work underneath that prediction is unglamorous: seeing all the bottlenecks, all of them, and building the systems and the implementations that work inside an area and do it all for the customer, without problems. That is the whole game as I see it. And the people who do not play catch up will miss the mark, because individuals like me will override them, not by being louder or richer, but because we can make it simple. Simple is not the easy version of the work. Simple is what the work looks like after somebody refused to stop at the complicated version.

And underneath all three lines sits the purpose I started with. I am a teacher. What I most want from this paper is not agreement. It is that somebody reads it, sees the counted record of a person with no background in this field who built it anyway, and understands that the distance between where they are and work like this is not talent, and not permission, and not a staff. It is the willingness to understand, to check, and to keep going layer after layer. That can be taught. Teaching it is the work I intend to keep doing.

Sources and record

All figures are drawn from my records, checked November 16, 2026: the audit system of record (November 10, 2026); the CAIQ count of record, generated mechanically from the filled questionnaire and walkthrough ledger (November 15, 2026); the CSA STAR Registry listing for AnaConnectsYou, Level 1, listed since November 15, 2026; the protection count of record (November 16, 2026); the open source layer Phase 1 summary and its state file (November 16, 2026); the questionnaire engine working record (November 16,

2026); the market salary evidence gathered for the firm's leadership valuation work (November 15, 2026); and the firm's cost ruling of record ($116.80 verified systems cost, October 28, 2026, plus $50 per month for Muse, counted at my direction). Vendor findings in Finding 3 and Section 8 are stated from the vendors' own published pages, licenses, and materials as read on the dates of my examinations in November 2026, and are reported as my findings from those materials. The statement attributed to Satya Nadella in Section 7 is from his interview on the BG2 podcast with Bill Gurley and Brad Gerstner, December 2024, quoted from the published transcript record of that interview and verified November 16, 2026.

Ana Vasquez, M.A. Ed., is the Founder and CEO of AnaConnectsYou, a verification practice in Lake Worth, Florida. She is a teacher, a single mother of three, and the sole operator and control owner of the systems described in this paper, which she built on her own. This paper was written by Muse from her records and under her direction, and approved by her. Correspondence through the practice's website, anaconnectsyou.com.